

# HOT TOPICS

A z/OS Newsletter

In this issue...

Migrating to z/OS

64-bit

ServerPac

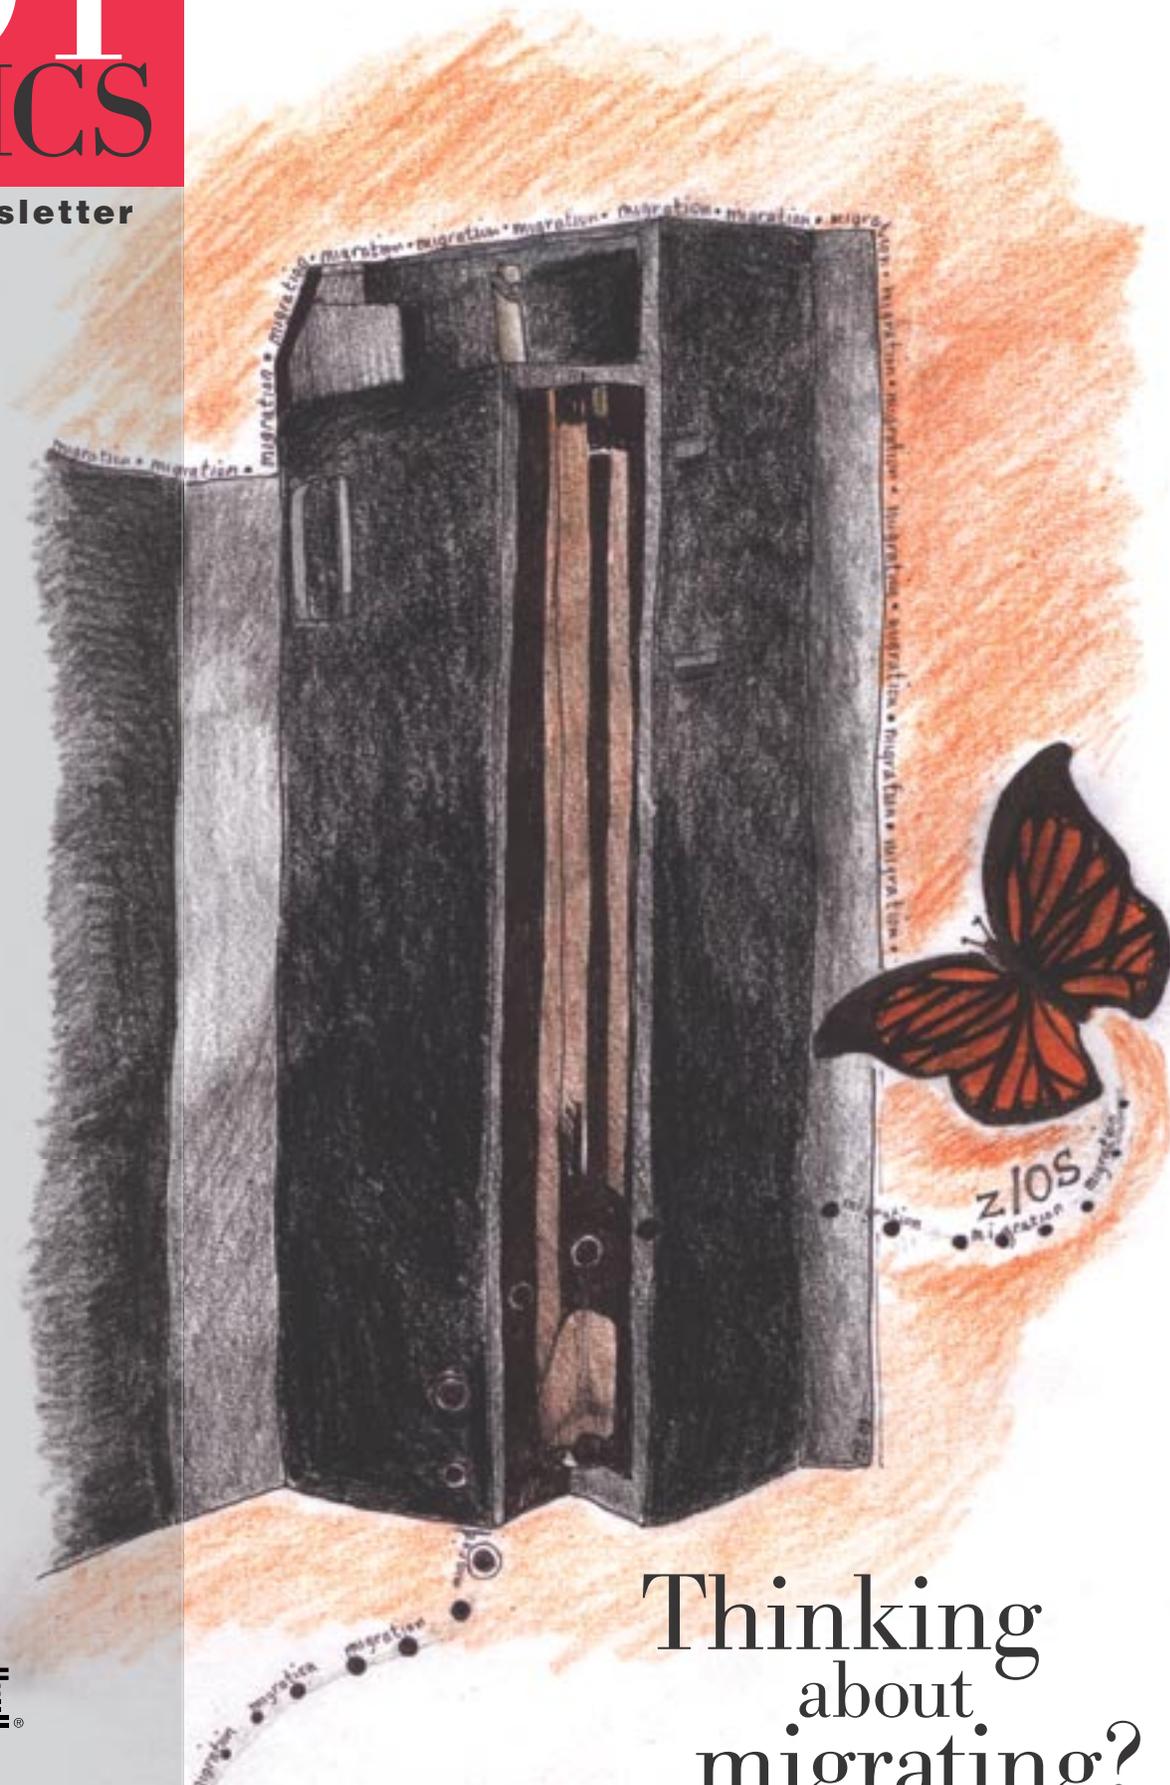
WLM

EIM

and much more!

**IBM**<sup>®</sup>

February 2003 - Issue 8



Thinking  
about  
migrating?

The logo consists of a red square containing the word "HOT" in white, serif, all-caps font, positioned above the word "TOPICS" in black, serif, all-caps font.

**A z/OS Newsletter**

**O N L I N E**

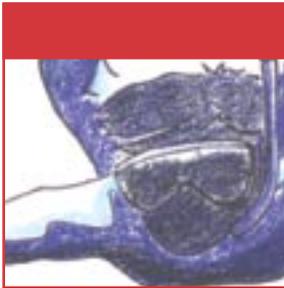
[http://www.ibm.com/servers/s390/os390/bkserv/hot\\_topics.html](http://www.ibm.com/servers/s390/os390/bkserv/hot_topics.html)

# Table of contents

## DEPARTMENTS

- 4 A letter from the editors
- 25 Crossword puzzle
- 27 z/Favorites CD
- 50 Our contributors

- 5 Migration to z/OS:  
Come on in, the water's fine  
BY GREG DAYNES AND  
JAMES BECKER



- 9 Migration to z/OS:  
Which path do I take?  
BY GREG DAYNES AND  
JAMES BECKER



- 11 The z/OS Bimodal  
Migration Accommodation:  
Separating myth from reality  
BY GREG DAYNES AND  
JAMES BECKER



- 13 Around the world in  
64-bit  
BY BOB RODGERS AND  
MARY MOORE

- 14 I'll migrate with a little help  
from my friends  
BY JUDY WASHO

- 18 Someone to watch  
over me:  
Making service  
safer and easier  
BY JUDY WASHO

- 20 Smooth sailing to  
e-business with z/OS 1.4  
Communications Server  
BY LESIA COX

- 23 ServerPac:  
It's a family affair  
BY JOHN EELLS

- 28 See how they run:  
Automatic balancing of  
batch workloads with  
WLM 1.4  
BY PETER BAEUERLE

- 30 More system  
automation magic!  
BY MARY ELLEN COLEMAN

- 31 Waste not, want not  
Reusing the softcopy  
collection interface  
BY SUE VAN PARYS

- 33 What on earth is  
transmogrification?  
Graphics, transmogrifica-  
tion, and the z/OS  
Collection  
BY SHIRLEY SWENSON

- 34 eee I'm being watched?  
... User identity mapping  
services for enabling  
secure heterogeneous  
applications  
BY FRANK HILDENBRAND,  
PEGGY LABELLE, AND  
JASON WILLIAMS

- 37 Line 'em up for Domino 6:  
Better performance,  
easier migration  
BY JOSEPH BOSTIAN

- 38 Availability:  
Getting better all the time  
BY ROBERT ABRAMS,  
ED MEKEEL, AND  
ROY MOEBUS

- 41 It takes two to make it right:  
System-Managed CF  
Structure Duplexing  
BY KELLY PUSHONG AND  
DAVID RAFTEN

- 43 Take control with SCLM:  
ISPF option 10  
BY MARSHA O'BRIEN

- 43 The IMS connection:  
Easing connectivity  
for z/OS  
BY BARBARA KLEIN

- 45 PhD not required:  
Accessing IMS data  
in z/OS 1.4  
BY JACK YUAN AND  
SANDY SHERRILL

- 48 This may be your lucky  
star: ATS STAR plus  
Automated Tape Allocation  
Manager for z/OS  
BY AMIR RAKHA

## Back cover

- Dude, where's the  
mainframe?  
BY BERNICE CASEY,  
PETER PAGEREY, AND  
DON SPANGLER

# A letter from the Editors



**Welcome** to the eighth issue of the Hot Topics Newsletter! We've recently scoured the self-help sections at our local bookstore and taken on the spirit of self-improvement. (Perhaps the IRD animation yogi inspired us.) Subsequently, we have a bunch of exciting new changes to tout and tell you about! Forgive us if we gush...

While our content standards and commitment to bringing you the information you crave hasn't changed at all, our hot topic of the moment has. This time around, we're concentrating on a topic that seemed most "hot" on your minds at the last SHARE conference—and practically everywhere else—migration!

Many of you know Greg Daynes, perhaps from the many sessions he presents at SHARE conferences every six months. Well, we roped him into writing a few articles for you on the topic of migration. Within, you can get your overall migration fix with "Migration to z/OS: Come in, the water's fine," determine your recommended migration scenario with "Migration to z/OS: Which path do I take?" and learn about bimodal migration with "The z/OS bimodal migration accommodation: Separating myth from reality." Bob Rogers and Mary Moore lend their expertise on migrating to 64-bit z/Architecture in "Around the world in 64-bit z/OS" and we snagged Judy Washo to tell you all about hands-on migration support in one of her two articles "I'll migrate with a little help from my friends."

As always, we also have an abundance of other articles that are sure to tickle your technical fancy. Beyond the migration concentration, this issue definitely packs a well-rounded punch of invaluable (if we do say so ourselves) morsels. You can read about everything from ServerPac ("ServerPac: It's a family affair") to Enterprise Identity Mapping ("eee I'm being watched?...") to availability ("Availability: Getting better all the time"), and should definitely check it all out.

By far, the biggest contribution to our quest for self-improvement is the newly revamped look and feel of the overall newsletter. We retired the old masthead for the snazzy new one we currently sport on the cover. Also, the visual layout of the whole newsletter has been updated, and the sections slightly rearranged. The index is now permanently out of the hardcopy newsletter, but you can still find it, tirelessly indexing the current issue and all past issues, on our Web site at [http://www.ibm.com/servers/s390/os390/bkserv/hot\\_topics.html](http://www.ibm.com/servers/s390/os390/bkserv/hot_topics.html).

Despite our thinking we're fabulous, we make all these changes for you, our readers. Our goal with all of this is to make the newsletter easier and more enjoyable for you to navigate and read, so let us know if we've done a good job or completely missed the sanity bus. In the meantime, enjoy your new issue!

-The Editors

---

EXECUTIVE EDITOR	SUE SHUMWAY
MANAGING EDITORS	TERRY MERRILL JOHN URBANIC
ART DIRECTOR	TRACY BONDI

---

# Migration to z/OS: Come on in, the water's fine

BY GREG DAYNES AND JAMES BECKER



Greg Daynes

Does diving into the cold water of a new release leave you a bit too invigorated? Then read about the migration enhancements we've made for z/OS™ 1.4 and future releases. They'll make your swim in the migration pool more pleasant than ever.

But first, just what do we mean by "migration"? *Migration* is the process of installing a new version or release of a program (in this case z/OS) to replace an earlier version or release. After a successful migration, the applications and resources on your new system function in the same (or similar) manner as they did on your old system. You are then in a position to do whatever customization and programming are necessary to make good use of (exploit) the enhancements available in your new system.

We advocate a fairly conservative approach to migration and recommend performing many of the migration actions in stages. This approach allows you to take advantage of incremental changes, while limiting the amount of change that occurs at any one time. Isolating changes makes it easier to plan for and implement upgrades, and identify and resolve problems if they occur.

For example, we recommend making processor changes separate from operating system upgrades. Also, we recommend (if possible) switching to z/Architecture™ separate from either of those changes. Furthermore, we recommend making as many changes as possible on your existing system. Again, these recommendations limit the amount of change being intro-

duced with the new software level. An added benefit of staging changes is that you can take advantage of new functions sooner, and you can use the latest "best practice" configuration values sooner.

While you are still using your existing system (even an OS/390® 2.10 system), the types of changes you can make include the following:

- Performing a required migration action, such as converting to workload manager (WLM) goal mode
- Specifying parameters that will become required in a future release, such as specifying CONSOLE names and defining console groups instead of using the ALTCONS parameter
- Updating parameters and options that will become obsolete in a future release, such as some C/C++ compiler options and the IP Services statements ASSORTEDPARMS and KEEPALIVEOPTIONS
- Finding alternatives for functions that will be removed in a future release, such as:
  - Elements or features that were removed (for example, LAN Server and SOMObjects ADE, which were removed from z/OS 1.2)
  - Functions within elements or features that were removed (for example, selected Library Management functions from ISPF, or the RDBM backend from LDAP Server)
  - Hardware (I/O devices as well as processors) that will not be supported on the new release

You need to evaluate how much change (and risk) you want to make (or accept) at any given upgrade. Some factors are:

- Your existing plans and commitments. Do you already have capacity requirements or need for new operating system functions? If you do, depending on the commitments, you might need to combine a hardware, software, and possibly architecture change. Minimally, the order in which you perform the migration tasks will be affected.
- Your ability to schedule multiple upgrade windows. If scheduling

upgrades is not a problem for your applications, you might want to minimize the amount of change (risk) and make several small hardware or software changes. However, if because of testing, application availability, or other factors you get only two upgrade opportunities a year, you might want to limit those to one hardware change and one software change.

- Your ability to back out from a given change. While it is normal to plan for backout from a processor change, architecture change, or a new operating system level, you might not want to back out from more than one of them at a time.
- How important it is to you to have similar test and production environments, and a similar disaster recovery environment (for example, in terms of hardware, software, and architecture levels). Disaster recovery environments should mirror as closely as possible their associated production environments. Therefore, plans for your disaster recovery systems should include upgrading the hardware to support z/Architecture once your production environment supports it.

## Migration enhancements

In response to your feedback, IBM® has made changes to ease migration to z/OS in the areas of software pricing, fallback to ESA/390 mode, migration complexity, and frequency of upgrades. Let's review the changes.

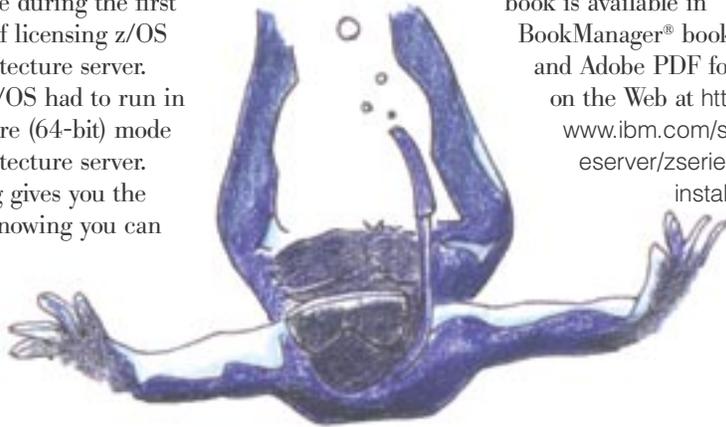
### Pricing initiatives:

In April 2002, IBM announced several new pricing initiatives that provide you with flexibility and price/performance benefits when pricing z/OS on a z/Architecture server. Workload License Charges (WLC) introduced the LPAR utilization model so that you could price z/OS at less than full capacity of a z900 server. This makes "white space" available to a much larger set of customers and allows you to price software at less than the machine's full capacity. Also, for the z900 server, you have the option of either remaining with your existing Parallel Sysplex License Charges (PSLC)

or moving to WLC. If you run WLC, you can take advantage of additional price/performance through the introduction of new pricing levels in excess of 875 MSUs. For the z/800 server, zSeries Entry Level License Charges (zELC) is an attractively-priced option for z/OS. And z/OS.e™ is a specially-priced offering to help you deploy new e-business applications.

### Bimodal Migration Accommodation:

In response to your requests to be able to fall back to ESA/390 (31-bit) mode during migration to z/Architecture (64-bit) mode while running z/OS, IBM is providing the z/OS 1.2/3/4 Bimodal Migration Accommodation. This software offering provides a safety net when first migrating to z/OS on a z/Architecture server (z800, z900, or equivalent) by allowing the alternative of running in ESA/390 (31-bit) mode during the first six months of licensing z/OS on a z/Architecture server. Previously, z/OS had to run in z/Architecture (64-bit) mode on a z/Architecture server. This offering gives you the security of knowing you can



return to 31-bit mode if there are any 64-bit problems during your migration. For more details on the Bimodal Migration Accommodation, see the article in this issue “The z/OS Bimodal Migration Accommodation: Separating myth from reality” on page 11.

### New z/OS migration book:

Migration from OS/390 2.10 to z/OS 1.4 is especially important because z/OS 1.4 is the base for exploiting future zSeries enhancements and because 2.10 is the last release of OS/390 from which a single-step migration to z/OS is possible. Therefore, IBM created a new migration book, *z/OS Version 1 Release 4 Migration* (GA22-7499). This book documents all migration considerations for a customer coming from an OS/390 2.10 system going to a z/OS 1.4 system, including general actions, element-specific actions, actions related to upgrading hardware, and actions related to migrating to z/Architecture. It consoli-

dates information from 34 z/OS books, five Redbooks™, five Washington Systems Center flashes, numerous APARs (including informational APARs), and a handful of IBM and user SHARE presentations.

Netting out all the scattered information into a single book that also includes information on enhancements you might want to exploit once your migration is complete, will certainly make planning and performing your migration easier. The book includes an appendix that describes the testing environment and procedures that we used to verify the information presented in the book. The book was created to accelerate the migration of OS/390 2.10 customers to z/OS 1.4. However, it can also be reviewed by customers migrating from z/OS releases to ensure that they haven't missed anything. The book is available in

BookManager® book format and Adobe PDF format on the Web at <http://www.ibm.com/servers/eserver/zseries/zos/installation/>.

It was originally made available on September 27, 2002 (concurrent with the general availability of z/OS 1.4) and has been updated since. The book is planned to be updated at least one more time in 2003.

The structure of the new z/OS migration book hopefully makes it easy to use. The book has separate chapters for each element and feature of z/OS, as well as one for general migration actions and one for hardware-related migration actions. Within each chapter, the information is categorized based on when the migration action should be performed:

- **Actions to perform before installing z/OS 1.4.** These are migration actions that you can perform today, that is, you can do them on your current OS/390 2.10 system. You don't need the z/OS 1.4 level of code to make these changes, and the changes don't require the z/OS 1.4 level of code to run once they are made.

- **Actions to perform before the first IPL of z/OS 1.4.** These are migration actions that you can perform after you've installed z/OS 1.4 but before the first time you IPL. These actions require the z/OS 1.4 level of code to be installed but don't require it to be active. That is, you need the z/OS 1.4 programs, utilities, and samples in order to perform the migration actions, but the z/OS 1.4 system does not have to be IPLed in order for the programs to run.
- **Actions to perform after the first IPL of z/OS 1.4.** These are migration actions that you can perform only after you've IPLed z/OS 1.4. That is, you need a running z/OS 1.4 system to perform these actions. Note that the term “first IPL” does not mean that you have to perform these actions after the very first IPL, but rather that you need z/OS 1.4 to be active to perform the task. You might perform the task quite a while after the first IPL. For example, in the case of activating z2 mode for JES2, it could be months after the first IPL.

Within each category is a separate topic for each functional change that caused a migration action. The functional change is briefly described, followed by a table that identifies:

- The z/OS release in which the change was introduced.
- Whether the migration action is required in all cases, whether it's required only if you are using a particular function, or whether it's not required but recommended for some particular reason.
- Any hardware required by the functional change. This includes processor and peripheral devices; drivers, engineering changes, or patches needed on the target system; and even specific hardware functions that must be active.
- Any software required by the functional change. This includes z/OS optional features, software products, and PTFs that are needed on the target system, as well as specific software functions that must be active.
- Any other requirements for systems that share resources with the new system. They include hardware or software that is required on other systems, and specific functions on other systems that must be active or must not be used.

- Any known limits or restrictions on how the function can be used.
- Any known impacts of using the function, such as increased storage or more time required to run.

Finally, the actual migration actions associated with the functional change are listed, followed by pointers to where more information about the functional change can be found.

This structure and consistency across all the elements and features makes it very easy for you to develop and execute an OS/390 2.10 to z/OS 1.4 migration, with the confidence that nothing will be missed and everything will go smoothly.

Because this is a new document, we'd like to know how we did. After you use the book, please fill out the brief (three questions) questionnaire at <http://www.ibm.com/eserver/zseries/zos/wizards/migguidesurvey/migguidesurvey.html>.

#### Annual release cycles:

The last change that affects current and future migrations is the frequency of upgrades. z/OS is switching from a six-month release cycle to an annual release cycle in response to your concerns. This gives you more time to order a release and to exploit the new functions provided in that release. The annual releases also allow IBM the time to do even more integration testing for z/OS and related middleware. This will include longer Early Support Programs that will allow participants the opportunity to test and stress z/OS releases in their environments for longer periods of time prior to availability.

In conjunction with annual releases, the coexistence, migration, and fallback policies are being aligned with the service policy. Starting with the first annual release planned for September 2004, whichever releases are currently in service will be able to coexist with the new release being introduced. As a result, you will no longer find yourself having to perform a two-step migration because the currently supported release that you are running doesn't coexist (or provide fallback with) with the release to which you want to migrate. There are some exceptions, which are documented in the *z/OS and z/OS.e Planning for Installation* book (GA22-7504).

The change to annual releases will enable more customers to migrate to every (or at least every other) release. Depending on when (or if) you freeze your systems, you can develop z/OS migration habits similar to migratory birds. Because annual releases will become available every September, you will be able to order and install a new level of z/OS every fall, just as birds fly south for the winter. If your business cycle demands that you make updates in the first half of the year, then you can migrate to z/OS every time the swallows fly back to Mission San Juan Capistrano. Migrating every year allows you to take advantage of the new z/OS features and enhancements, and puts you in a good position for maintaining operating system currency.

#### Why migrate to z/OS 1.4?

You should move forward to z/OS 1.4 because, in addition to the useful enhancements in z/OS, it positions you to take advantage of the latest middleware and zSeries technologies. It also is the base for exploiting future zSeries hardware enhancements. For example:

- The next version of DB2® will require z/OS 1.3 or higher in 64-bit mode.
- Lotus™ Domino™ 6 requires z/OS 1.2 or higher.
- The latest Java™ SDK, 1.4, requires z/OS 1.4 or higher.
- z/OS 1.4 is the minimum level required for the planned z/OS Console Availability feature.

Additionally, if you're on OS/390 2.10, we strongly encourage you to move to z/OS 1.4 as it offers you the last chance to migrate to z/OS in a single step.

Cheryl Watson of Watson & Walker, Inc., said that "z/OS 1.4 is the basis for future enhancements and it's definitely the place to be. Because there are no longer drawbacks of going to z/OS, I highly recommend that you make this move."

#### Where to start your migration

The typical steps to a successful migration are:

1. Learn about z/OS 1.4. Start with the *z/OS Version 1 Release 4 Introduction and Release Guide* and the z/OS home page: <http://www.ibm.com/servers/eserver/zseries/zos/>.
2. Perform as many of the migration actions as you can on your existing

OS/390 2.10 system. You can identify many of the tasks using the new z/OS migration book. Review the migration actions to determine which ones are required in your environment. Some depend on your configuration or workload, and are identified accordingly. These actions should be made to, or copied (cloned) to, all existing systems that will be migrated to z/OS 1.4.

3. Plan the installation of z/OS 1.4. Refer to the *z/OS and z/OS.e Planning for Installation* book. Ensure that all driving system, target system, and cross-system (coexistence) requirements are met. Contact your vendors and ensure that they support z/OS 1.4.
4. Order and install z/OS 1.4. If you use a ServerPac, read *ServerPac Installing Your Order*. If you use a CBPDO, read the z/OS 1.4 program directory.
5. Prepare target system hardware and software. If you are migrating from OS/390 2.10, you should use the z/OS migration book. If you are migrating from z/OS, you should follow the pointers in *z/OS and z/OS.e Planning for Installation* to the element and feature migration books for details on what is required. Again, not all of the actions are required. Some depend on your environment, configuration, and workload, and are identified accordingly.
6. IPL the new z/OS 1.4 system with your updated customization and configuration files.
7. Perform any migration actions that are required to be done after you IPL z/OS 1.4. Again, not all of the actions are required. Some depend on your environment, configuration, and workload, and are identified accordingly.
8. Test the new system with a representative sampling of applications and vendor products, as appropriate.
9. Deploy z/OS 1.4 to other systems within a sysplex, data center, or enterprise, as appropriate.

The migration is now complete.

10. When you are confident that a system, or in some cases all systems in a sysplex, are not going to fall back to OS/390 2.10, exploit the functions introduced in z/OS 1.4. Even if you did not migrate from OS/390 2.10, you can identify many of the new z/OS functions by reading the appendix

in the z/OS migration book that lists new functions to consider. Deploy this exploitation on other systems (again within a sysplex, data center, and eventually enterprise).

Depending on your situation, you might have to perform the steps above (except for 1 and 4) iteratively. For example, you might do them for an application development environment, again for a quality assurance environment, and again for your production environment. Also, although testing is mentioned in only Step 8, it is likely that you will perform some level of testing after each step.

### Ingredients in a successful migration

The following are key tasks in a successful z/OS 1.4 migration:

1. Determine if (or how) you are going to stage your migration. To assist you in determining the order in which you should upgrade your software, and possibly your processor and architecture mode, read the article “Migration to z/OS: Which path do I take?” on page 9.
2. Develop a migration plan identifying the order in which migration actions are to be performed, by whom they are to be performed, and when.
3. Verify that your installation plan includes the resolution of all target system requisites and cross-system (coexistence and fallback) requisites.
4. Ensure that either your installation plan or migration plan covers contacting your vendors and making the necessary upgrades. (Some people consider contacting their vendors an installation task and some consider it a migration task.)
5. Implement a software maintenance strategy. For information about IBM’s service strategy and recent enhancements to the service process, read the Red Paper “Improvements in Service” at: <http://www.redbooks.ibm.com/>.

The following are some element-specific migration actions. This is not meant to be a complete list but rather some of the more important migration actions that could affect you:

- Migrate from workload manager (WLM) compatibility mode to goal mode. (Goal mode was introduced in z/OS 1.3). You should try to migrate

to goal mode prior to starting your migration to z/OS 1.4. There are aids to assist you. For more information, see the WLM Web site at <http://www.ibm.com/eserver/zseries/zos/wlm/>.

- Switch to z/Architecture™ (introduced in OS/390 2.10). For information about when you should perform this migration action, see the article “Migrating to z/OS™: Which path do I take?” on page 9. In addition, you can find a detailed checklist of things to do when you convert to z/Architecture in the new z/OS migration book (available on the Web) and in Washington Systems Center flash 10185 (they both contain the same information).
- Use the new resolver component of Communications Server (introduced in z/OS 1.2). You might have to update your security definitions and your BPXPRMxx PARMLIB member. This needs to be performed before you IPL z/OS 1.2 (or later).
- Remove your dependency on the OS/390 2.10 level of the C/C++ compiler (introduced in z/OS 1.2). Because z/OS 1.2, 1.3, and 1.4 all ship both the OS/390 2.10 level and the z/OS 1.2 (ISO C++ 1998) level of the C/C++ compiler, you need to ensure that cataloged procedures, environment variables, and ISPF concatenations are set up properly for the level of the compiler you want to use. Once z/OS 1.2 (or later) is deployed, you should start developing a plan to migrate from the OS/390 2.10 level of the compiler because it will be removed in a future release of z/OS.
- Update programs that use Language Environment® because the non-CICS® default options were changed to favor 31-bit applications (introduced in z/OS 1.2).
- Migrate to JES2 z2 mode (introduced in z/OS 1.2). You should modify your exit routines to support both R4 mode and z2 mode. Field names were changed to reflect incompatible control block changes. Therefore, you might not be able to just reassemble your JES2 exit routines and continue on. You can’t turn on z2 mode until all systems sharing JES2 spool are running with z/OS 1.2 or later JES2.
- Migrate to a maximum of 999,999 job numbers in JES2 or JES3 (introduced in z/OS 1.2). If you choose

to raise the limit on the number of jobs to greater than 99,999, be aware that the JOBID format switches from JOBnnnnn to Jnnnnnnn, TSUnnnnnn to Tnnnnnnn, and STCnnnnn to Snnnnnnn. This might affect more than just your JES exit routines, for example, your automation products and security definitions might be affected. This is a migration action that will not affect you unless you choose to raise the limit.

- Reformat automatic restart management (ARM) couple data sets (introduced in z/OS 1.4). You must do this before you IPL z/OS 1.4.
- Reformat OMVS couple data sets (introduced in z/OS 1.4). You must do this before you IPL z/OS 1.4.

### Conclusion

We have introduced a number of changes that will make your migration to z/OS 1.4 and future releases easier and safer. The z/OS 1.2/3/4 Bimodal Migration Accommodation provides a safety net when first migrating to z/OS on a z/Architecture server (z800, z900, or equivalent). The annual release cycles maintain predictability while giving you more time to order, migrate, and exploit a release. The new z/OS migration book provides “one-stop shopping” for an OS/390 2.10 customer looking for migration actions.

There has never been a better time to migrate to z/OS. Take the plunge!



### Thinking about migrating to z/OS?

**A new book is available  
on migration...**

“z/OS Version 1 Release 4 Migration”

**Check it out at:**

**[http://www.ibm.com/  
servers/eserver/zseries/  
zos/installation/zos\\_  
migration.html](http://www.ibm.com/servers/eserver/zseries/zos/installation/zos_migration.html)**

# Migration to z/OS: Which path do I take?

BY GREG DAYNES AND JAMES BECKER

It's been said that all roads lead to Rome. While all roads do not lead to z/OS 14, there are enough migration paths that at least one should meet your needs. As you travel along the path you've chosen, expect to see changes. These changes can include an operating system change from OS/390 to z/OS, a processor change to a z/Architecture server (z800, z900, or equivalent), an architecture change from ESA/390 to z/Architecture, a performance tuning change from workload manager (WLM) compatibility mode to goal mode, and a change in your terms and conditions from Parallel Sysplex License Charges (PSLC) pricing to Workload License Charges (WLC) or zSeries Entry Level License Charges (zELC) pricing.

When you make the changes, you can do so one at a time, in a variety of orders, or all at once. As a result, many different migration paths are possible. The question then becomes, "Which is the best migration path for me?"

Robert Frost wrote that taking the road less traveled can lead to great adventure. But when it comes to migration, the objective is to make the journey as smooth as possible. So, unless you are a thrill seeker, you will want to take one of the well-traveled paths outlined below.

## General guidelines

While you have until OS/390 2.10 goes out of service to migrate to z/OS 14, there are three reasons why you shouldn't wait to start your migration. First, several of the migration paths have multiple steps, all of which must be complete before OS/390 goes out of service. Second, it might take you several weeks or even months to complete your deployment of z/OS. Third, you will not be able to exploit the functions of z/OS until your migration is complete.

As mentioned previously, you can make most, if not all, migration changes at the same time. In addition, you can make the changes necessary to exploit new functions in the new release at the same time. However, the more prudent approach is to migrate in stages, and then once the

migration is complete, begin exploiting new functions. This staged approach enables you to limit the amount of change that occurs at any one time. Isolating changes makes upgrades easier to plan for and implement, and makes problem determination and recovery easier. Staging migration



also allows you to make many changes before you even order or install z/OS 14. For example, anyone who is still running in WLM compatibility mode should migrate to WLM goal mode before upgrading to z/OS 14.

Your optimal migration path will depend on how much change (and risk) you are willing to make (or accept) at any given upgrade, the level of existing software, the level of existing hardware, the operating system functions currently used, the software and hardware levels of your test environment, the software and hardware levels of your disaster recovery environment, and the number of systems on which you will deploy the upgrades. As a result, different customers will choose different paths.

The following migration paths are organized by existing software and hardware levels. Find and read the path that is appropriate to your environment.

## OS/390 2.9 (or prior levels) on any processor

If you're running OS/390 R9 or earlier, you're running a release that is not supported for either service or migration. Contact your IBM representative to understand what alternatives are available.

## OS/390 2.10 on a processor that doesn't support z/OS

If your processor is an S/390 Parallel Enterprise Server R2, R3, G3, or G4; S/390 Multiprise® 2000; S/390 Integrated Server; PC Server S/390; or RS/6000 with S/390 Server-on-Board, you will have to upgrade your processor before you migrate to z/OS. Your current processor doesn't support z/OS. If your goal is to exploit the functionality of z/OS that requires z/Architecture, you will need to:

1. Migrate to a z/Architecture server (z800, z900, or equivalent).
2. Initially, run your OS/390® 2.10 system in ESA/390 (31-bit) mode.
3. Migrate to z/Architecture (64-bit) mode (while still running OS/390 2.10).
4. Migrate to z/OS 14, remaining in z/Architecture mode once the z/Architecture migration is complete.

**Note:** Alternatively, you can migrate to an S/390 Parallel Enterprise Server G5 or G6, or an S/390 Multiprise 3000, or equivalent. Once you've upgraded to this processor, you can migrate to z/OS 14. Because you are running on an ESA/390 processor, you will IPL in ESA/390 mode, not z/Architecture mode. Also, you will not be able to use z/OS functions that require z/Architecture.

## OS/390 2.10 on an ESA/390 processor that supports z/OS

If you are currently running OS/390 2.10 on an S/390 Parallel Enterprise Server G5 or G6, an S/390 Multiprise 3000, or an equivalent processor, then you can migrate either your software or processor first.

Migrate your *software first* if you want (or need) to exploit some of the many functions in z/OS that don't require a z/Architecture server (z800, z900, or equivalent). Also migrate your software first if you don't have plans to upgrade your processor until the second half of 2004, in order to complete your software migration prior to your hardware migration. To migrate software first:

1. Migrate to z/OS 1.4. Because you are running on an ESA/390 processor, you will IPL in ESA/390 (31-bit) mode, not z/Architecture (64-bit) mode.
2. Optionally (but recommended), download and install the z/OS 1.2/3/4 Bimodal Migration Accommodation. This will enable you to initially IPL z/OS 1.4 in ESA/390 mode when you upgrade your processor to a z/Architecture server (z800, z/900, or equivalent). For information about the Bimodal Migration Accommodation, see the article "The z/OS Bimodal Migration Accommodation: Separating myth from reality" on page 11.
3. Upgrade to a z/Architecture server.
4. Optionally (but recommended), initially IPL in ESA/390 mode by specifying ARCHLVL 1 in LOADxx (if you have downloaded and installed the Bimodal Migration Accommodation). You have six months from when you license z/OS on the z/Architecture server to migrate all z/OS LPARs on that server to z/Architecture mode.
5. Migrate to z/Architecture mode. (Note: If you have not downloaded and installed the Bimodal Migration Accommodation, you will be forced to IPL in z/Architecture mode—even your first IPL).

Migrate your *processor first* if you want (or need) a bigger, faster server. Also, migrate your processor first if you want to take more than six months to migrate to z/Architecture mode (once you upgrade to a z/Architecture server). Follow these steps:

1. Migrate to a z/Architecture server (z800, z900, or equivalent).
2. Initially run your OS/390 2.10 system in ESA/390 (31-bit) mode.
3. Migrate to z/Architecture (64-bit) mode (while still running OS/390 2.10).
4. Migrate to z/OS 1.4, remaining in z/Architecture mode once the z/Architecture migration is complete.

### **OS/390 2.10 on a z/Architecture-capable processor**

If you are currently running OS/390 2.10 on a z/Architecture server (z800, z900, or equivalent), you only have to upgrade your software.

1. Migrate to z/Architecture (64-bit) mode (while still running OS/390 2.10).
2. Migrate to z/OS 1.4, remaining in z/Architecture mode once the z/Architecture migration is complete.

### **z/OS 1.1, 1.2, or 1.3 on an ESA/390 processor that supports z/OS**

If you are currently running z/OS on an S/390 Parallel Enterprise Server G5 or G6, S/390 Multiprise 3000, or equivalent processor, then you can migrate either your software or processor first.

Migrate your *software first* if you want (or need) to exploit some of the many functions in z/OS that don't require a z/Architecture server. Also migrate software first if you don't have plans to upgrade your processor until the second half of 2004, in order to complete your software migration prior to your hardware migration. To migrate software first:

1. Migrate to z/OS 1.4. Because you are running on a ESA/390 processor, you will IPL in ESA/390 (31-bit) mode, not z/Architecture (64-bit) mode.
2. Optionally (but recommended), download and install the z/OS 1.2/3/4 Bimodal Migration Accommodation. This will enable you to initially IPL z/OS 1.4 in ESA/390 mode when you upgrade your processor to a z/Architecture server (z800, z/900, or equivalent). For information about the Bimodal Migration Accommodation, see the article "The z/OS Bimodal Migration Accommodation: Separating myth from reality" on page 11.
3. Upgrade to a z/Architecture server.
4. Optionally (but recommended), initially IPL in ESA/390 mode by specifying ARCHLVL 1 in LOADxx (if you have downloaded and installed the Bimodal Migration Accommodation). You have six months from when you license z/OS on the z/Architecture server to migrate all z/OS LPARs on that server to z/Architecture mode.

5. Migrate to z/Architecture mode. (Note: If you have not downloaded and installed the Bimodal Migration Accommodation, you will be forced to IPL in z/Architecture mode—even your first IPL).

Migrate your *processor first* if you want (or need) a bigger, faster server. Follow these steps:

1. Optionally (but recommended), download and install the z/OS 1.2/3/4 Bimodal Migration Accommodation. This will enable you to initially IPL z/OS 1.2, 1.3, or 1.4 in ESA/390 mode when you upgrade your processor to a z/Architecture server (z800, z/900, or equivalent). (Note: the Bimodal Migration Accommodation is not available if you are running z/OS 1.1) For information about the Bimodal Migration Accommodation, see the article "The z/OS Bimodal Migration Accommodation: Separating myth from reality" on page 11.
2. Upgrade to a z/Architecture server.
3. Optionally (but recommended), initially IPL in ESA/390 mode by specifying ARCHLVL 1 in LOADxx (if you have downloaded and installed the Bimodal Migration Accommodation). You have six months from when you license z/OS on the z/Architecture server to migrate all z/OS LPARs on that server to z/Architecture (64-bit) mode.
4. Migrate to z/Architecture mode. (Note: If you have not downloaded and installed the Bimodal Migration Accommodation, you will be forced to IPL in z/Architecture mode—even your first IPL).
5. Migrate to z/OS 1.4, remaining in z/Architecture mode once the z/Architecture migration is complete.

### **z/OS 1.1, 1.2, or 1.3 on a z/Architecture-capable processor**

If you are currently running z/OS on a zSeries server (z800, z900, or equivalent), your hardware is all right and you've already been running z/OS in z/Architecture mode. All you need to do is migrate to z/OS 1.4 (remaining in z/Architecture mode).

## Conclusion

You have a lot of flexibility in migrating to z/OS 1.4 on a zSeries server. By staging the various migration changes, you limit the amount of change that occurs at any one time, making it easier to plan for and implement your upgrade, and recover from any problems.

You can start on some of the migration tasks now, even before you order or install z/OS 1.4. For example, if you are currently running OS/390 2.10 on a z/Architecture server (z800, z900, or equivalent), you should migrate to z/Architecture now. This will reduce your workload when you migrate to z/OS 1.4.

There is a checklist to convert to z/Architecture in the new *z/OS Version 1 Release 4 Migration* book (GA22-7499, available on the Web). The checklist is also in Washington Systems Center flash 10185.

So choose your path and start migrating today!

# The z/OS Bimodal Migration Accommodation: Separating myth from reality

BY GREG DAYNES AND JAMES BECKER

On August 13, 2002, IBM announced the z/OS 1.2/3/4 Bimodal Migration Accommodation in response to customer requests to provide a safety net when first migrating to z/OS on a z/Architecture server (z800, z900, or equivalent). The Bimodal Migration Accommodation is software that allows you to run z/OS in ESA/390 (31-bit) mode during the first six months of licensing z/OS on a z/Architecture server. Previously, z/OS had to run in z/Architecture (64-bit) mode on a z/Architecture server. This offering gives you the security of knowing you can return to 31-bit mode if there are any 64-bit problems during your migration.

The Bimodal Migration Accommodation has unique terms and conditions. For example, it can only be used for a six-month period on each z/Architecture server. Workload License Charges (WLC) are not supported for pricing and licensing of IBM software on the server on which the accommodation is installed, and you can't use certain functions when IPLed in ESA/390 mode. These functions include, but are not limited to, Intelligent Resource Director (IRD), HiperSockets, 64-bit real storage, 64-bit virtual storage, and IPv6. In addition, z/OS.e is not supported by the accommodation.

There has been some confusion about the Bimodal Migration Accommodation since it was announced. This article will attempt to separate myth from reality.



### Myth 1

*The Bimodal Migration Accommodation is only available on z/OS 1.4.*

### Reality

The accommodation is available to z/OS 1.2, 1.3, and 1.4 customers.

### Myth 2

*The Bimodal Migration Accommodation is included in the base code or service stream of z/OS 1.4 (that is, it is integrated in z/OS 1.4 ServerPacs).*

### Reality

The accommodation is only available as a Web download. It is delivered as a small FMID installed on top of z/OS. There is a unique package (and FMID) for z/OS 1.2, 1.3, and 1.4 customers. In order to download the appropriate package, you must:

1. Go to the z/OS download Web page <http://www.ibm.com/eserver/zseries/zos/downloads/>.
2. Select z/OS 1.2-1.4 Bimodal Migration Accommodation offering.
3. Select the appropriate package: 1.2, 1.3, or 1.4.
4. Click **Continue**.
5. Sign in (or register if a first-time user).
6. Optionally, click **View License** to read the terms and conditions.

7. Click **Accept License**.

8. Download the package and the readme file.

### Myth 3

*You can't download and install the Bimodal Migration Accommodation while you're running z/OS 1.2, 1.3 or 1.4 on an ESA/390-architected server (for example, an S/390 Parallel Enterprise Server G5 or G6, or an S/390 Multiprise 3000 server).*

### Reality

If you are a z/OS 1.2, 1.3, or 1.4 customer, you can download and install the accommodation no matter what server you are currently using. In fact, we recommend that if you plan to upgrade to a z/Architecture-capable server, you should download and install the accommodation while you are on an ESA/390 server. See the article "Migration to z/OS: Which path do I take?" on page 9.

### Myth 4

*The six-month "Accommodation Period" starts as soon as you download the Bimodal Migration Accommodation.*

### Reality

The Bimodal Migration Accommodation is available for each z/OS license for six

## **z/OS Now! Migrating to z/OS**

BY CHERYL LOUGHLIN

Thinking about migrating to z/OS? Visit the **Migrating to z/OS** Web page and click on **z/OS Now!** in the right column to view a short, informative overview of z/OS 1.4. You will find out why this release is an especially good one to migrate to when you check out this slick animation.

### **Here is the Migrating to z/OS**

#### **Web page URL:**

**[http://www.ibm.com/servers/eserver/zseries/zos/installation/zos\\_migration.html](http://www.ibm.com/servers/eserver/zseries/zos/installation/zos_migration.html)**

months only (“Accommodation Period”). The Accommodation Period begins when z/OS (5694-A01) is licensed to a z/Architecture-capable server (z800, z900, or equivalent server), whether or not the Bimodal Migration Accommodation offering was downloaded or installed. There is no accommodation period for z/OS licenses while they are licensed to an S/390 Parallel Enterprise Server G5 or G6, or an S/390 Multiprise 3000 server.

#### **Myth 5**

*There is one six-month Accommodation Period for your entire enterprise.*

#### **Reality**

The six-month time period is unique for each z/Architecture-capable server (z800, z900, or equivalent server) in your enterprise. As stated previously, a separate six-month Accommodation Period begins when z/OS is licensed to a new server in your enterprise.

#### **Myth 6**

*The six-month Accommodation Period for a given server can actually span more than six months, if you only use the accommodation to run in ESA/390 mode infrequently.*

#### **Reality**

The Accommodation Period is for six consecutive calendar months. At the end of the Accommodation Period, z/OS will be supported in 64-bit mode only, and your authorization to use the accommodation on that Server will no longer be valid and may not be renewed.

#### **Myth 7**

*All z/OS images on the same server must be IPLed in the same architecture (ESA/390 or z/Architecture) during the Accommodation Period.*

#### **Reality**

During the Accommodation Period, each z/OS image can run in either ESA/390 (31-bit) mode or z/Architecture (64-bit) mode. Individual images (in LPARs) can be in different architectures on the same server at the same time. In fact, if the Bimodal Migration Accommodation is installed, the individual images can share the same SYSRES volumes and still IPL in different architectures (assuming that they have unique LOADxx members).

- To IPL an image in ESA/390 mode, specify ARCHLVL 1 in LOADxx. Make sure the LPAR is defined with both central and expanded storage.

If greater than 2 GB is specified for central storage, only 2 GB will be used and the remainder will not be available for other LPARs to use.

- To IPL an image in z/Architecture mode, either explicitly specify ARCHLVL 2 in LOADxx or allow it to default. Make sure that the LPAR is defined with all central (now called real) storage. If expanded storage is defined, it will not be used and it will not be available for other LPARs to use. At the end of the Accommodation Period, z/OS will be supported in z/Architecture mode only.

#### **Myth 8**

*There is an additional charge for the Bimodal Migration Accommodation.*

#### **Reality**

The accommodation is available at no charge.

#### **Myth 9**

*You can't IPL z/OS in ESA/390 mode as a guest on VM while running on a z/Architecture-capable server.*

#### **Reality**

During the Accommodation Period, you can IPL z/OS in ESA/390 mode as a guest under z/VM™, as well as in z/Architecture mode, when running on a z/Architecture server (z800, z900, or equivalent). The z/VM image must be a 64-bit CP image. At the end of the Accommodation Period, z/OS will be supported in 64-bit mode only, even as a guest under z/VM.

#### **Conclusion**

Now that we have dispelled the most common myths, you can see that in most cases reality is better than the myth. It provides more flexibility, covers more systems, and does not have any additional costs associated with it. You can now migrate to z/OS on a z/Architecture server with confidence that you can stage your z/Architecture migration, and that you can fall back to ESA/390 mode if you experience any problems.

# Around the world in 64-bit

BY BOB RODGERS AND MARY MOORE

The migration to z/OS is similar to other release-to-release migrations, with one exception: the move to the new 64-bit z/Architecture. When you make that switch and your operating system is running in z/Architecture mode, will your current business applications execute correctly? Will there be any problems with systems management tools and middleware when running under z/Architecture? The answers to these questions should give you added comfort in making the decision to move forward to z/OS, because the z/Architecture was designed to protect your current portfolio of 31-bit and even 24-bit applications!

Before we explain the impact of 64-bit to your applications, let's clarify some terms.

- ESA/390 architecture is what the S/390 servers are based on. These include 9672, MP2000 and MP3000 server. ESA/390 is a 31-bit mode architecture. OS/390 and even z/OS can run in this architecture level mode when you specify ARCHLVL 1 in LOADxx.
- z/Architecture is the base for the zSeries servers (z800 and z900). It includes 64-bit support and other new technology advancements. OS/390 V2R10 and all z/OS releases can take advantage of this architecture by specifying ARCHLVL 2 in LOADxx.

## Business applications

Business applications are written to execute business logic. They do not run in an authorized state and have been completely shielded from the changes in the architecture. The subset of the architecture that is used by business applications is identical in ESA/390 and z/Architecture. Since no programming interfaces were changed, applications perceive these architectures as compatible. This means that both 24-bit and 31-bit applications will operate correctly when running in z/Architecture mode. It is that simple.



For continued application development and maintenance, as far as the transition to z/Architecture is concerned, you can continue to use the same levels of assembler, compilers and language environment.

## Home-grown tools

If you have home-grown tools which execute in an authorized state, they will need to be reviewed. If the program does not deal with real storage addresses, then there is no impact when running in z/Architecture mode. You may proceed to the next section. The best clue that a tool deals with real storage addresses is that it issues the Load Real Address (LRA) instruction.

Now, if you are sitting with a program that is manipulating real storage addresses, a closer examination of the program is required, but there is still a good chance that this program will be compatible! The operating system by default will move data into real storage frames below 2GB as part of the page fixing operation. Therefore, even programs in this category will probably operate compatibly. If the program that creates the virtual storage that is *page fixed* does not use a new z/Architecture option to request that the page fix uses real storage above 2GB, there is no incompatibility.

One difference with z/Architecture—although it's not really an incompatibility—is that issuing an LRA instruction against virtual storage which is not page fixed will often lead to a program check (which becomes a 0D3 ABEND). Most cases of issuing LRA against pageable storage are

errors and it is good that they are caught under z/Architecture. For the other case, where LRA is being used simply to determine if a page is currently back in real storage, the TPROT instruction can be used to answer the same question without causing an error.

## IBM software (middleware and tools)

While there should be few business applications that are impacted by the introduction of 64-bit z/Architecture, this is not necessarily true for systems management tools and middleware. Some of these products may depend on real storage addresses (like a performance tool that examines memory activity), or they may actually issue their own I/O instructions.

The good news is these products have received extensive testing and have been in production in many customer environments since the introduction of z/Architecture in December 2000. As early as August 2000, hundreds of products from multiple software vendors were tested on z/Architecture servers. Now, over two years later, you can reap the benefits of this early testing and the efforts of our early support customers covering four releases of z/OS. Simply request the latest maintenance from your software vendors for your tools and products.

As with any release-to-release migration, check with your independent software vendor to ensure all of your products are supported on the new releases of z/OS. Some products may require maintenance or a new release in order to run on new releases of z/OS.

For IBM tools and middleware, you can obtain the latest maintenance for new releases of z/OS by requesting the PSP buckets for your specific server type. See "Migration to z/OS: Come on in the water's fine," on page 5 for information on other z/OS migration topics.

# I'll migrate with a little help from my friends

BY JUDY WASHO

Are you considering migrating to z/OS 1.4 sometime soon? With z/OS transitioning to annual releases, migrating now allows you to realize the benefits of the industry-leading z/OS 1.4 features and enhancements. It also allows you to quickly take advantage of upcoming releases of key middleware, such as DB2, Lotus Domino and Java™, and you may see a financial benefit

based on the new z/OS pricing options too.

For companies with z/OS, z/OS.e, or OS/390, migrating, installing, and maintaining their operating-system-related products and independent software vendor products can be time and resource intensive.

Due to factors such as current labor market conditions or operating strategies, many companies lack the staff necessary to perform these operations. And although some companies do have the in-house expertise to perform these jobs, such tasks can sidetrack key IT personnel from business-critical responsibilities.

Starting to worry? Don't! IBM Global Services in the United States and Canada has a new solution, SystemPac® Plus for z/OS and OS/390, that provides hands-on migration support for you and simplifies the way you create or upgrade a complete software system.

In a nutshell, an IBM specialist will assist you by researching and applying toleration/coexistence maintenance, upgrading you to a complete, customized z/OS system, including independent software vendor products, enabling an e-business application server environment, and completing setup for z/OS Communications Server, z/OS UNIX® System Services, z/OS Security Server, HTTP/WebSphere® Application Server, sysplex support for Shared HFS, plus installing and activating IBM Electronic Service Agent and enabling Distributed File Service (DFS) Server Message Block (SMB).

Sound good? You bet. So sit back and relax and let IBM do the work for you. Here's what IBM does for you when you order SystemPac Plus!

## 1. IBM conducts and documents your planning meeting!

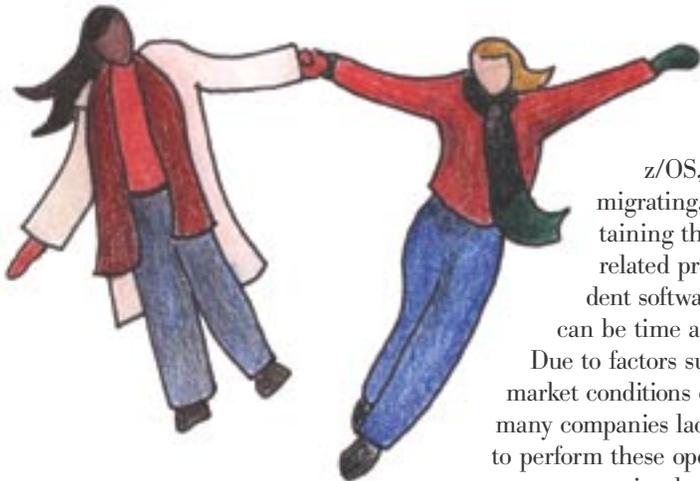
First, an IBM specialist meets with you to perform the necessary planning to implement the project, helps identify the priorities you want to be addressed, and documents the results of this meeting in a report to you.

## 2. IBM orders your toleration maintenance!

A key element in migrating to a new system is to ensure that you have all the toleration/coexistence maintenance necessary for your current system. To address this requirement, an IBM specialist gathers your SMP/E CSI (consolidated software inventory), and orders toleration/coexistence maintenance tailored to your specific environment, based on your SMP/E CSI. A maintenance package is built containing the toleration/coexistence PTFs and will be shipped back to you, where the IBM specialist performs the SMP/E RECEIVE. Your SMP/E CSI is retained at the IBM production center to be used to fill your subsequent orders for corrective and preventive service. For more information on getting system maintenance, refer to the article "Someone to watch over me" on page 18 in this issue!

## 3. IBM orders your full-volume dump SystemPac for z/OS, including subsystems!

The IBM specialist, having gathered all the required information to tailor build your z/OS system during the planning stage, places the order for your SystemPac and tracks its progress through delivery to your site.



## Not ready to migrate to z/OS 1.4 yet?

How about OS/390 2.10? No, it's not too late! Only SystemPac® and SystemPac Plus have extended the ordering window for OS/390 2.10 until June 2003. That's right!

If you are currently on OS/390 2.6, 2.7, 2.8, or 2.9, you can still migrate to OS/390 2.10 using SystemPac (available worldwide) or SystemPac Plus.

The IBM SystemPac you receive is a complete, customized, and fully integrated total operating system and software replacement package. It includes both IBM and independent software vendor products to ensure you the most reliable system possible.

You can select from a wide range of independent software vendor products to be installed on SystemPac, including choices from BMC Software, Inc., Candle Corporation, and Computer Associates International, Inc., to name a few. Where else will you find that?

Your SystemPac comes customized according to your specified parameters and is designed to exploit new technology and enable the software required to build an e-business Application Server environment, including e-business enabled in IBM default mode, initial setup and customization of WebSphere Application Server 4.0.1 and its prerequisites, Workload Manager goal-mode initial setup with a sample policy, middleware enabled and started at IPL, and enhanced customized systems maintenance with periodic refreshes.

Your SystemPac comes with the following activities already completed:

- \* Workload Manager (WLM) is configured in goal mode with a sample policy.
  - Monoplex environment is enabled.
  - z/OS Security Server is customized to authorize access to WLM policies.
  - WLM couple data sets are allocated.
  - Sample service definition is installed to the WLM couple data sets and activated.
- \* Resource Recovery Services (RRS) is implemented.
  - LOGR couple data sets are defined.
  - The RRS subsystem name is added to the IEFSSNxx parmlib member.
  - The RRS procedure is copied to CPAC.PROCLIB.
  - The RRS logstreams are defined.
  - RRS is set up to start at IPL.
- \* z/OS UNIX System Services is enabled in full-function mode.
  - A minimal systems-managed storage (SMS) environment is set up based on your specified storage class, storage group and SMS dedicated volume (volser, model/type, address).
  - z/OS Security Server is customized for z/OS UNIX with a RACF® database containing the necessary security profiles for z/OS UNIX. Both the started procedure table and the RACF STARTED class are updated with the required cataloged procedures. Customization of z/OS Security Server includes:
    1. RACF groups for z/OS UNIX
    2. A superuser definition
    3. Security levels for daemons.
  - Parmlib members are tailored for z/OS UNIX, including members: BPXPRM00, BPXPRMFS and COFVLFEX.
  - HFS file system is allocated and built.
- \* z/OS Communications Server is set up and enabled in IBM default mode.
  - Customized system parmlib members, including:
    1. APF-authorizing and linklist libraries in PROGxx
    2. Subsystem names TNF and VMCF defined in IEFSSNxx
    3. FILESYSTYPE defined in BPXPRMxx for the TCP/IP stack.
  - Constructed and customized TCP/IP SAMPPROF parmlib member, including:
    1. TELNET server profile (“BEGINVTAM” to “ENDVTAM”) to provide TELNET logon attributes
    2. Autolog definitions to start TCP/IP applications (FTPD, PORTMAP, MVS NFS, and WEBSRV)
    3. Ports numbers reserved for privileged applications to prevent accidental use by users
    4. Hardware configurations and home IP addresses
    5. Router profiles (static or dynamic).
- \* IBM HTTP Server is enabled.
  - Constructed procedure WEBSRV is in PROCLIB.
  - RACF security environment is established for WebSphere for z/OS.
  - Shell script is set up to create WebSphere environment in /etc directory (including: httpd.envars, httpd.config, and more).
  - HTTP server is customized (for example, httpd.config).
  - Websphere for z/OS is tailored using post-install jobs.
- \* InfoPrint® Server is customized and set up (optional). If the InfoPrint Server is included in your SystemPac order, the following are configured and customized:
  - InfoPrint Server startup procedure in PROCLIB.
  - System parmlib members including:
    1. APF-authorizing and link list libraries
- Constructed TCP/IP TCPDATA parmlib member and hosts local file.
- Updated TCPIP.HOSTS.LOCAL file with host IP address and host name.
- TSO MAKESITE command is used to create host information files: TCPIP.HOSTS.SITEINFO and TCPIP.HOSTS.ADDRINFO.
- TCP/IP daemon configurations are in the UNIX /etc directory including:
  1. SYSLOGD daemon configuration (syslog.conf)
  2. Updated /etc/rc set up to start daemons SYSLOGD and INETD when UNIX initializes.
- Constructed z/OS Communications Server procedures for TCPIP and FTPD
- Constructed VTAM® application definition TCPAPPL in VTAMLIST
- Established the RACF security environment for z/OS Communication Server
- Allocated and built TCP/IP IVP data sets (including: hosts local, translation tables, services and more)

2. Subsystem names TNF and VMCF in IEFSSNxx
  3. FILESYSTYPE in BPXPRMxx for UNIX TCP/IP stack
  4. MVS™ program properties table (SCHEDxx)
- Allocated and initialized IP PrintWay™ data sets (routing, options, transmission queue and message log).
  - IP PrintWay functional subsystem in JES2PARM.

This enables you to exploit InfoPrint Server to transmit output from your z/OS system to printers or servers in a TCP/IP network.

- \* z/OS Network File System is customized and setup. If NFS is included in your SystemPac order, a default configuration is supplied including:
  - Constructed startup procedures MVS NFS, PORTMAP and MVS NFS in PROCLIB with corresponding FACILITY class RACF profiles
  - Customized system parmlib members
  - Allocated and initialized NFS data sets (such as CNTL, FHD-BASE, LOG, and more)
  - Customized NFSATTR and EXPORT in “MVS NFS.CNTL”.

This allows you to access host data from your workstation or access remote data on other NFS servers.

- \* CICS Transaction Server for z/OS is set up and exploited (optional). If CICS Transaction Server for z/OS 1.3 or 2.2 is included in your SystemPac order, common data sets that are required for CICS initialization and CICS logstreams to MVS LOGR are defined:
  - CICS log manager is set up with DASD-only model log streams for the CICS system log.
  - CICS Web support (CWS) and 3270 Bridge are set up.
    1. System parmlib is customized to add the CICS libraries (SDFHDLL1 and SDFHEXCI) to MVS LNKLST.
    2. CICS library SDFHDLL1 is defined to RACF program control.

3. External link from /etc directory to the CICS DLL “plug-in” is defined and a “Service” directive is added to /etc/httpd.conf for identifying the CICS DLL “plug-in” to the Web server.
  4. TCPIP service (HTTPNSSL) and program resource definitions for the LE language interface modules are defined to the CICS CSD.
  5. CICS online IVP (DFHIVPOL) is defined for CWS and 3270 Bridge testing.
- CICS online IVP (DFHIVPOL) for CICS EXCI is set up by installing the sample group DFH\$EXCI and setting the initialization parameters IRCSTRT and ISC to YES. We provide a job DFHIVPXC to verify that the EXCI is established by running the sample assembler program DFH\$AXCC to issue EXCI CALLs with null parameters.

- \* DB2 Universal Database™ is set up and exploited (optional). If DB2 for z/OS and OS/390 6.1.0 or 7.1.0 is ordered, selected functions are configured:
  - DB2 WLM application environment
    1. DB2 sample WLM application environment procedure (DSN1WLM) is constructed.
    2. WLM application environment (WLMENV) is defined in the WLM sample policy.
    3. RACF security for the WLM application environment is established.
  - DB2 ODBC setup
    1. ODBC initialization file is updated with the correct subsystem ID (DSN1).
    2. Job JDB7717E is provided to bind CLI default packages and plan.
    3. A sample C application is run to verify the DB2 ODBC installation.
  - DB2 JDBC setup
    1. The db2jdbcclases.zipfile is added to the Java CLASSPATH.
    2. The JDBC ‘lib’ directory is added to the LIBPATH and LD\_LIBRARY\_PATH.

3. The JDBC ‘bin’ directory is added to the PATH setting.
  4. The JDBC driver properties file is modified to reflect your installation requirements.
  5. The JDBC profile is generated by executing the **db2genJDBC** command.
  6. A sample Java program (sample01) is run to verify the JDBC interfaces.
- \* WebSphere Application Server v4.0.1 for z/OS and OS/390 is enabled.
    - WLM sample policy is updated with workload management policies for the WebSphere for z/OS server environment
    - System parmlib and proclib are updated.
    - TCP/IP profiles are updated to reserve ports.
    - The WebSphere for z/OS error log stream is created.
    - RACF database is updated with the necessary security authorizations for WebSphere for z/OS.
    - The WebSphere for z/OS v4.0.1 installation variables that were tested by SystemPac are provided, which will save you time defining the ISFP customization and installation dialog.
    - Websphere for z/OS configuration HFS file and post-installation jobs are provided for creating the LDAP database, defining log streams, and RRS.
    - Jobs to create and grant DB2 tables for WebSphere for z/OS configurations are provided.
    - IVP jobs and instructions for J2EE and MOFW are given.
    - Instructions are also provided for these tasks:
      - Installing the administration application on your workstation.
      - Starting conversation to the BOOTSTRAP server
      - Updating or creating the Websphere for z/OS configuration file from the administration application workstation.

- \* So what kind of Websphere for z/OS environment do you get? When all these tasks are completed, the SystemPac you get delivered has Websphere for z/OS v4.0.1 set up and customized with all of its prerequisites are enabled:

- z/OS UNIX
- z/OS or OS/390 System Logger
- Resource Recovery Server (RRS)
- z/OS Security Server - RACF
- Java for z/OS and OS/390 1.3.0
- WorkLoad Manager - goal mode
- z/OS Communications Server
- IBM HTTP Server
- DB2 V7.1.0

*Now that's value-add!*

#### 4. IBM installs and verifies your SystemPac!

Now you have your own tailor-built SystemPac, but our SystemPac Plus work doesn't stop here. An IBM specialist performs the following activities:

- Installs the system
- Runs the installation job streams
- Completes the MVS SREL installation
- IPLs the system
- Enables WLM goal mode (for z/OS systems)
- Runs post-installation jobs after IPL
- Establishes z/OS UNIX
- Verifies the installation

#### 5. IBM tailors your system!

We're not done yet! There is much more that can be done to help you exploit new technology and build an e-business application server. This is when the IBM specialist assists you to further tailor your environment:

- Completes setup of z/OS Communications Server
- Completes setup of z/OS UNIX
- Completes setup of z/OS Security Server
- Customizes HTTP Server/WebSphere Application Server
- Implements z/OS UNIX sysplex support for shared HFS
- Activates z/OS and OS/390 Distributed File Service (DFS) Server Message Block (SMB)
- Installs and activates IBM Electronic Service Agent for zSeries™. If you are installing the SystemPac on a z/Series 800 or 900 model processor or a S/390® 9672 G5 or G6, order the free IBM Electronic Service Agent as part of the SystemPac to monitor, track, and capture software inventory, hardware error logs, and performance information!!

#### 6. IBM documents and turns over your new z/OS system!

You now have a complete customized z/OS system with an e-business Application Server environment! But we don't leave you just yet. The IBM specialist schedules a closeout meeting with your staff to review the project and its completion and leaves you with the necessary documentation for you to maintain your new system.

Hey, one more thing... You are entitled to follow-on maintenance packages called Selective Follow-on Service (SFS) which help to stabilize your installed system over the period of time you have specified. SFS packages are ordered at the same time as the initial SystemPac and contain critical service information, which is applicable to a package order that has become available since the SystemPac or previous SFS package was built.

Now wasn't that easy for you? IBM took the worry out of your migration and did most of the work for you. What could be simpler?

**For more information on SystemPac Plus - A total z/OS and OS/390 Migration Solution, visit our Web site at**

**<https://www.ibm.com/ca/custompac>**

**and select the United States as your country.**

**To place an order, contact your IBM specialist or call 888-426-4343.**

# Someone to watch over me: Making service safer and easier

BY JUDY WASHO

There is a risk of putting maintenance on a system, and there is a risk of not putting it on. The article called “z/OS platform maintenance just got easier”, in Issue 6 of the Hot Topics Newsletter, described IBM’s maintenance recommendation for z/OS and OS/390 software and the improved testing of service. With the improvements in service quality, delivery, and ease of installation, the risk of putting maintenance on a system has been significantly reduced. That’s not all. IBM Global Services provides even more value-add services for corrective and preventive service, as well as, high-impact or pervasive APARs (HIPERs) and PTFs in error (PEs). On top of that, now you can easily order and subsequently install, toleration and coexistence service, making migration to new hardware or software much easier and safer.

You can order a service package tailored to your specific environment, based on your SMP/E CSI (consolidated software inventory), and have it sent to you electronically within hours. Not only that, once submitted, your SMP/E CSI-based profile is maintained at the IBM production center and will be updated each time you place a new order for service. You can’t find that anyplace else!

So where is this extraordinary service found? It’s in the new and improved SoftwareXcel Enterprise Edition for zSeries (SoftwareXcel). SoftwareXcel provides a total package of support services for your zSeries system. It helps you simplify software maintenance, especially for migration, by reducing the time it takes for software problem resolution and providing tools to improve staff productivity. SoftwareXcel offers many support services. However, this article will only discuss the newest enhancements that are available in the service request and delivery (SRD) application at no additional charge when you have a SoftwareXcel contract.

If you already are a SoftwareXcel customer, all you need to do is access IBMLink via the internet at <http://www.ibm.com/ibmlink/link2>, logon to IBMLink, and go to the SRD application.

If you are not a SoftwareXcel customer, click on <http://www.ibm.com/services/its/us/mus19d1.html> and press the *call me now* button on the SoftwareXcel information screen to request that IBM contact you.

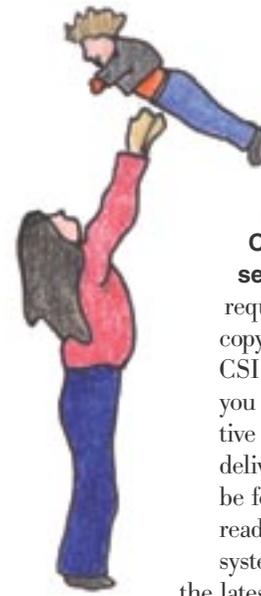
Once in the SRD application, select **Order OS/390**, z/OS from the left-hand navigation bar. The system displays the **Order OS/390, z/OS - Ordering options** page. This page contains the links to all of the above services. Now that you know how to get to these services, let’s take a closer look at each one.

**Corrective Service** is used to add individual PTFs to a system to fix known problems which you have encountered.

To request corrective service, select **Order corrective service PTF(s) based on CSI profile**. You can select single or multiple PTFs. You must specify the SMP/E CSI profile for which you want the corrective service configured. This is used to determine what service is needed, including all requisite service and resolution of any PE chains. If you do not have an existing SMP/E CSI profile at the IBM production center, you can select **Upload CSI profile** to get one. If you have forgotten your profile name, you can select **Request CSI profiles** to retrieve a list of all your existing profiles stored at the IBM production center.

You can track the status of your order, as well. All of the corrective service will be shipped to you electronically. So, all you need to do is download it from the mainframe or workstation. If your order is larger than the allowable limit (compressed), it will be shipped on the physical media of your choice. You can also resubmit your order to request less service and get it electronically! Don’t forget that your updated SMP/E CSI profile will be maintained by IBM for your use with subsequent orders.

**Preventive Service** for a system is used to add predefined collections of PTFs to a system to avoid problems which you have not yet encountered.



To request preventive service, select **Order preventive service**. You will be required to send in a copy of your SMP/E CSI profile every time you request preventive service so that the delivered service will be for all products already installed on your system to PUT-1 plus the latest recommended service upgrade (RSU) level.

(Remember, PUT is a monthly collection of all closed PTFs identified with a SOURCEID of PUTyyymm, where yy indicates the last two digits of the year and mm indicates the month. For example, PUT level 0010 indicates the collection of all PTFs that closed in October, 2000.) Additionally, HIPER and PE-fixing PTFs will be shipped, along with a customized installation guide and installation dialog jobs to install the service. Since your service order is based on your SMP/E CSI profile, it won’t include any extraneous service. So it can be delivered quickly and on the physical media you specify. Plus, the service is tested with an SMP/E APPLYCHECK and any return codes greater than 4 are resolved before the package is built and delivered to you! Now that’s value-add.

Order analysis has shown us that about 15% of outages affecting multiple systems or subsystems could have been avoided by better preventive maintenance practices. This means that the cause of the outage was fixed by a PTF that had been available for six months or more! This ordering option is just what you need.

**HIPER APARs and PE PTFs** are high impact or pervasive APARs and PTFs in error that could result in outages and/or impact end-user availability.

To request HIPER APARs and PE PTFs, select **Order HIPER/PE fixes or report, based on CSI profile**. Here you must specify your SMP/E CSI profile so that

the HIPER/PE PTFs missing from your installed inventory can be determined. Just as in corrective service ordering, you can **Upload CSI profile** or **Request CSI profiles** if you have forgotten the profile name. You have the option to request that just a report of the missing HIPER/PE PTFs be delivered electronically to your specified e-mail address and/or to get the actual HIPER/PE PTFs delivered to you electronically. Just as in corrective service, if your order is larger than the allowable limit (compressed), it will be shipped on the physical media of your choice.

**Toleration/coexistence service:** Toleration service allows a software product to tolerate, but not exploit, function introduced by a higher level of the same product. Coexistence service allows multiple levels of a software product to share resources on the same system or within a sysplex. This service also provides support for selected hardware devices.

To request Toleration/coexistence service, select **Order toleration/coexistence maintenance based on CSI profile**. Now here is an option that is only found in SRD and provides the maintenance you require when migrating to new hardware or new releases of software, on request. This option is very similar to **Order corrective service (PTFs) based on CSI profile**, except you get a list of toleration/coexistence hardware-base PTF's (2064, 3590 and more) or software-based (z/OS, OS/390, DB2, IMS, CICS, and more), using your SMP/E CSI profile. Once again, these are shipped to you electronically, except when over the allowable limit (compressed). Then it is shipped via the physical media you specify. If you are considering migrating, then this is a tool you do not want to be without!

IBM has added a new feature to its online software ordering tool, ShopzSeries at <http://www.software.ibm.com/ShopzSeries>, to provide information and links to request toleration and coexistence service for z/OS and OS/390. The links will direct you to the SoftwareXcel Web site. Additionally, ShopzSeries provides links to the CustomPac offerings for RefreshPac, ProductPac and SystemPac. SystemPac Plus, will be discussed in this issue's article, "I'll migrate with a little help from my friends" on page 14. Check it out! However, this article would not be complete without

mentioning RefreshPac and Selective Follow-on Services offered within the CustomPac offerings.

RefreshPac is a software preventive service package very similar to the **Order preventive service** option of SRD. The only difference is that RefreshPac can also include entitled follow-on maintenance packages called selective follow-on service (SFS). An SFS contains critical service information that is applicable to your RefreshPac that has become available since the RefreshPac or previous SFS package was built. These SFS packages help stabilize the system you have installed over a period of time, as specified by you during your order. You can select the number of SFS packages (from 0 to 3) and specify the intervals between them (30 to 60 days).

Each SFS package provides a report and includes SMP/E HOLD data, HIPER PTFs, PE PTFs, and preventive service. Your SFS packages are built according to a copy of your stored SMP/E CSI profile and can be installed using batch JCL or using the installation dialog. Although your CSI information for RefreshPac can be shipped either electronically or by physical tape, the finished package containing the actual service is shipped to you on physical media. For more information on RefreshPac, visit <http://www.ibm.com/ca/custompac>.

IBM recommends that you have a maintenance strategy that keeps you current and protects against known defects that could impact availability. Your maintenance strategy is key to your success. The information is intended to provide you with more options to improve your ordering and delivery of service, and make your service safer and easier to receive, apply and maintain.

## ShopzSeries

BY BOB TERWILLIGER

An easy, self-service way to plan and order zSeries software upgrades. Now includes VM and VSE so you can quickly order VM SDOs and VSE SIPOs! ShopzSeries checks co-reqs, pre-reqs, and incompatibility issues to ensure that your order is technically correct. Available in the United States and several countries in Europe. For more details (including geographic availability), visit the ShopzSeries Web site at:

**<http://www.ibm.com/software/ShopzSeries>**

# Smooth sailing to e-business with z/OS 1.4 Communications Server

BY LESIA COX

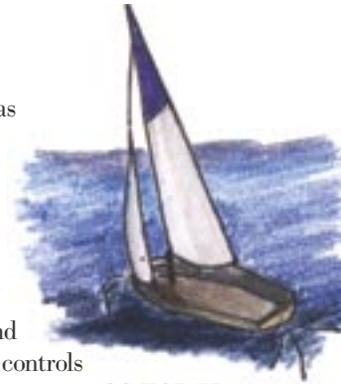
z/OS 1.4 Communications Server (CS) provides enhancements in areas which are critical to the success of your e-business. Enhanced security controls are available to address TCP/IP security concerns. Improvements have been made to TN3270E and Enterprise Extender to provide better integration of SNA and IP. Availability has also been enhanced by new Dynamic VIPA and Sysplex Distributor functions. In addition, the Policy Based Networking (PBN) Quality of Service (QoS) features and Intrusion Detection Services (IDS) allow self-optimization and self-protection for your data center, as shown in Figure 1.

The Communications Server protects data in the network by supporting a variety of cryptographic-based network security protocols such as IPSec, SSL/TLS, Kerberos™, and SNA Session Level Encryption. These security protocols ensure:

- Data origin authentication—the received data originated at the claimed sender
- Message integrity—the contents were unchanged in transit
- Data privacy—sensitive data is concealed using encryption

ally deploy, such as Intrusion Detection Services, to defend against attacks from the network.

These network and resource security controls enable a safe and secure z/OS TCP/IP deployment.



## SNA/IP integration

With the latest releases of Communications Server for OS/390 and z/OS Communications Server, S/390 and the IBM zSeries are world-class platforms for native e-business (TCP/IP-based) applications. However, conversion of existing SNA applications to TCP/IP-enabled applications can be economically impractical. Also, in many cases such conversions may even be technically impractical due to the lack of source code and adequate skills for the specific application. An additional complication is the wide variety of SNA-based endpoint devices, such as banking ATMs. So, how can we enable IP applications and preserve SNA application and endpoint investment, while converging toward a single network protocol?

Two key mechanisms, TN3270E and Enterprise Extender, provide SNA/IP integration.

## TN3270E

For 3270-based applications, the z/OS CS TN3270E server can be a key component to the solution, allowing TN3270E clients to access SNA applications through an IP network, and limiting the SNA network path to the inside of a single Communications Server image.

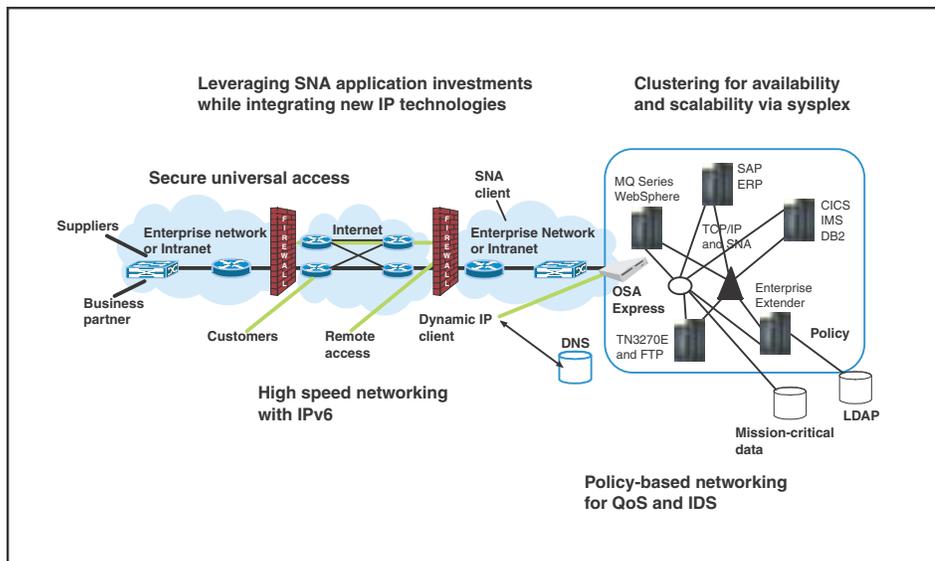


Figure 1 - z/OS Communications Server customer value

Read on as we describe these enhanced focus areas in more detail...

## Security

With increased use of TCP/IP and the Internet, enterprise security requirements have become more stringent and complex. The sources of many transactions are from “untrusted” networks such as the Internet, and sometimes from unknown users. Because of this, increased attention is paid to host and user authentication, data integrity and privacy in the network, as well as denial of service attacks. The z/OS Communications Server provides security controls to address these TCP/IP security concerns.

The Transport Layer Security (TLS) enablement for FTP and TN3270E functions allows you to specify this level of security without requiring end-to-end IPSec infrastructure to be in place.

The Communications Server also protects system resources and data from unauthorized access. Communications Server applications and the TCP/IP protocol stack protect data and resources on the system using standard SAF services.

The Communications Server safeguards the availability of the system by protecting against denial of service attacks from the network. The Communications Server has built-in defenses, and also provides services that an installation can option-

The already rich definitional capabilities of the z/OS TN3270E server are further enhanced in z/OS 1.4 CS with the following:

- Support for TLS secure TN3270E client connections.
- The ability to have multiple parameter sets differentiated by IP address or linkname is available.
- LU naming flexibility is provided by more flexible wildcarding capability and the provision for an LU naming exit.
- An option to preserve the client IP address to LU name association can now be specified for a period of time after disconnect.

### Enterprise Extender (EE)

EE provides an additional highly scalable and reliable component for SNA/IP integration strategies. It allows you to preserve your SNA application and device investment, while maintaining the session prioritization and availability characteristics of SNA and HPR. Furthermore, EE allows for a simplified network architecture that positions applications for exploitation of future advances in IP networking technologies.

z/OS 1.4 Communications Server enhances the dial processing for EE connections by providing a number of controls that allow for automatic recovery of EE connections after failures. Scalability is improved since EE lines and PU's can now utilize high-order element addresses. Diagnostics is also improved with a new RTP display which includes a number of additional statistics concerning the EE connection.

### Availability

z/OS 1.4 introduces new capabilities that allow a single sysplex-wide source VIPA address to be used for outbound TCP connections by all images in the sysplex resulting in single system image capabilities for both inbound and outbound connections.

For clients outside a Parallel Sysplex, Sysplex Distributor provides a single-IP-address appearance to application instances spread across the sysplex. It also distributes incoming work among the various instances. Many applications are part of a cooperative network of applications, and the sysplex applications that serve as clients to end users may also have to initiate (client-like) outbound connection requests to cooperating applications.

The SOURCEVIPA feature available in and prior to z/OS 1.4 allows applications to attain independence of any physical adapter; however, SOURCEVIPA is limited to statically defined VIPAs within a stack. Different instances of the same application using Sysplex Distributor, and thus having a single IP address for inbound connection requests, will use different IP addresses for their outbound connection requests.

These problems are resolved in z/OS 1.4 by allowing a Dynamic VIPA to be used as the source IP address for TCP applications and to have the Sysplex stacks collaborate on assigning ephemeral ports to prevent duplicate connection "4-tuples" when the same Distributed DVIPA is used as the source address on multiple stacks. (The term 4-tuples here refers to the source IP address, the source port, the destination IP address, and the destination port). This allows for outbound connections to be managed across the entire sysplex, such that for a particular distributed DVIPA a particular port value is assigned to a socket on only one TCP stack in the sysplex. This will ensure that inbound connection data can always be uniquely routed to the correct application instance, whether the connection was initiated by the client or by the sysplex application instances.

### Policy Based Networking (PBN)

Following the tradition of delivering the best and most advanced features for z/OS, the PBN support of z/OS CS allows customers to

- Differentiate e-business applications
- Optimize their e-business infrastructure
- Protect and defend it from intrusion that can disrupt e-business transactions.

Therefore, PBN plays an important role in enabling a self-managing autonomic computing system to achieve a set of desired objectives in performance, availability, and security. The PBN functions for network Quality of Service (QoS), and Intrusion Detection Services (IDS) greatly enhance the self-optimizing and self-protecting qualities of z/OS.

The PBN QoS feature in z/OS supports an extensive set of condition filters that allow administrators different ways to identify and differentiate service levels for

- Different hosts (clients/servers)
- Applications (for example, the Web, MQ, and FTP)
- Times (including time of day, days of week, months etc.)
- Application specific data (for example, URL/URI for Web traffic, and message queue priority for MQ).

Once a class of traffic is identified by the condition, a corresponding QoS service level is assigned (often identified as gold, silver, bronze, etc.) in the action. A set of QoS service level parameters can be used in the action for service differentiation ranging from setting the ToS/DSCP (including the VLAN priorities in VLAN header) of the corresponding packets to controlling TCP throughput and connections, to limiting bandwidth of a traffic class by using the token bucket mechanism, to denying traffic from utilizing network resources. PBN in z/OS provides network SLA Management Information Base (SLA MIB), which can be used to monitor the performance of respective network QoS policies. Thresholds can be defined for different QoS performance attributes (for example, delay and throughput) such that when a deviation occurs, traps will be sent to the attention of the right SNMP manager for appropriate actions. The **NETSTAT** command can also be used to retrieve policy performance data locally on the system.

The focus of IDS is self-detection and self-protection. The IDS is policy-driven and its intrusion detection points are integrated into the z/OS Communications Server TCP/IP system on the data path. The IDS is viewed as complementary to external network-based intrusion detection sensors. By exploiting its position as a communications endpoint, the IDS broadens the set of detectable intrusions and can detect intrusions that otherwise could go undetected. By integrating this function in the target server, the IDS can efficiently evaluate data that is unavailable to external network intrusion detection sensors. For example, data that has been encrypted end-to-end by IPSec can be examined in cleartext after decryption. The IDS also has access to internal information such as system resource usage, connection state information, and internal thresholds and counters that can further determine whether an intrusion event is surfaced.

The z/OS Communications Server IDS can also implement real-time online defensive policies against attacking packets that cannot be implemented by external network-based intrusion detection sensors. Using this approach to intrusion detection, the z/OS IDS has the capability to detect intrusions not yet invented, without requiring an update to policy.

An IDS report program, TRMDSTAT, provides summary and detailed reporting of IDS events and statistics. IDS messages can drive NetView™ message automation. Possible message automation messages include routing messages to a NetView console, notifying a system administrator via email or pager, or running TRMDSTAT reports and attaching the reports to an email to the system administrator.

### **Preparing for the next generation of IP**

In addition to providing the enhancements discussed thus far, z/OS 1.4 Communications Server lays the groundwork for the future by providing the initial support for Internet Protocol Version 6 (IPv6), which allows application development and testing for the next generation of the Internet.

IPv6 is designed to replace the current version, Internet Protocol Version 4 (IPv4). Most of today's internets use IPv4, which is approximately 20 years old and is approaching the end of its physical limits. The most significant issue surrounding IPv4 is the growing shortage of IPv4 addresses. With the continued growth of the Internet, exhaustion of IPv4 addresses early in the 21st century is inevitable.

IPv6 fixes a number of problems in IPv4, such as the limited number of available IPv4 addresses. IPv6 uses 128-bit addresses, an address space large enough to last for the foreseeable future. It also adds many improvements to IPv4 in areas such as routing and network autoconfiguration. IPv6 is expected to gradually replace IPv4, with the two coexisting for a number of years during a transition period. z/OS 1.4 Communications Server contains the first phase of IPv6 support.

Many companies have begun to announce IPv6 support in their products. Although many North American companies already have sufficient IPv4 addresses, other geographies will be moving to IPv6 more rapidly. In order to continue business-to-business communication with these areas, American companies will also have to support IPv6. The migration to IPv6 cannot be done overnight, so it is important to get started on this transition. The IPv6 support provided in the z/OS 1.4 release enables customers to build an IPv6 network, start using IPv6 specific applications and, more importantly, convert existing IPv4 applications to the IPv6 API without losing any functionality. If you are an application developer, take advantage of the IPv6 support available in z/OS 1.4 to start updating your applications for IPv6. z/OS 1.4 Communications Server can be an IPv4-only stack or a dual-mode stack. Dual mode stack refers to a single TCP/IP stack supporting both IPv4 and IPv6 protocols at the same time. The same source code can be used for IPv6-enabled stacks and traditional IPv4-only stacks. New protocol-independent APIs which support both IPv4 and IPv6 communication are available. These APIs are similar to what is already used for IPv4 communication (in most cases they are identical, with only the parameters being passed changing) and the changes necessary to support IPv6 are minor. Once updated to use the new APIs, applications on a dual stack node can simultaneously communicate with IPv4 and IPv6 peers using the same socket. To understand the details of what IPv6 support is provided in z/OS 1.4 Communications Server, refer to the *z/OS Communications Server: IPv6 Network and Application Design Guide* (SC31-885).

Your e-business is important. z/OS 1.4 Communications Server brings new function and enhancements to areas that are important to the success of your business ... So make your move to z/OS 1.4 Communications Server today!

# ServerPac: It's a family affair

BY JOHN EELLS

## Select-O-JES

All right, all right! Ours is still a somewhat conservative company. So we really call this “Selecting a JES element for installation.” But before, you had to install both JES2 and JES3. Then, like most people, you probably deleted the unwanted JES. Next, if you wanted the JES and BCP SMP/E zones merged, you had to do that. Finally, you could reclaim the space taken up by the JES element you deleted and give it back to your storage administrators.

No more! ServerPacs built after September show you a new panel when you create a configuration. It appears right after the “Choose an Installation Type” panel (where you select either Full-System Replacement or Software Upgrade). See Figure 1.

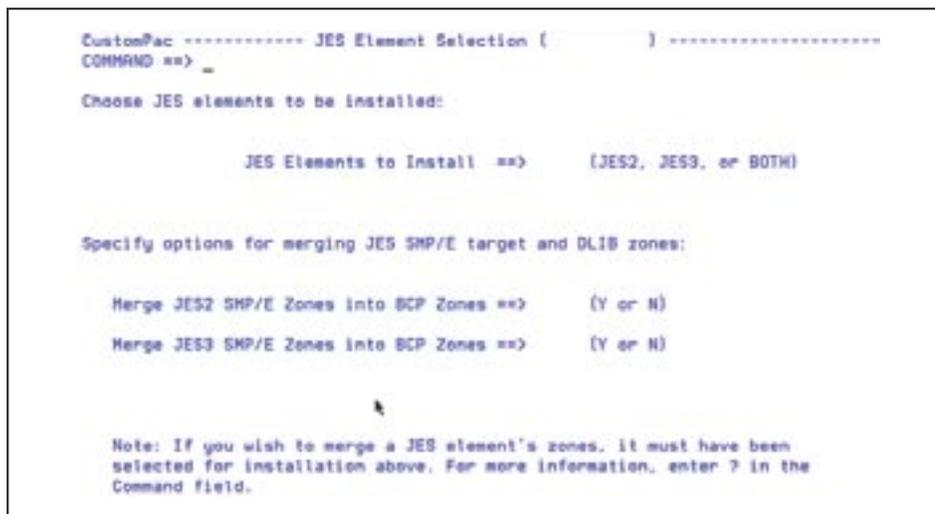


Figure 1 - New ServerPac panel.

Here, you can select the JES element you want to install—or even both of them, if you want to, and choose whether to merge JES and BCP SMP/E zones. (When you pick JES2, SDSF comes with it. Of course!)

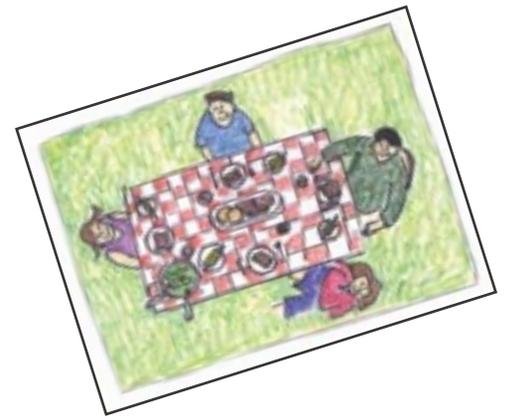
Only the data sets and zones for your selected JES show up in the dialog. The installation jobs will load only the data sets and zones for JES elements you want, and if their zones are to be merged with the BCP zone, it's done when they are loaded. The jobs you once ran to delete an unused JES element and merge the zones

are gone! And we took this opportunity to consolidate two of the SMP/E update jobs into one.

## More dialog changes

You will find better data set merge rules that prevent merging of incompatible load libraries. Data sets must now have compatible LPA and link list attributes to be merged. For example, you cannot merge a data set required in LPA with one that is not eligible for LPA. In addition to preventing unworkable configurations, this reduces the size of the change candidate lists, which should save you time.

You will also see a new **CHANGE MCAT** command that lets you override the “required in master catalog” flag for any data set. So, if you build maintenance systems



## Installation job changes

You will also find new and improved installation jobs. The ALLOCDS data set allocation job has changed. A lot! The old one—a traditional IEFBR14 allocation job—had some problems. For one thing, an IEFBR14 job's return code is *always* zero—whether data set allocations succeed or not. No sane person ever bothered to read all the output to make sure it ended without error. And worse, some problems were not obvious until something far later in the installation job stream failed, and fixing them sometimes meant starting over from the beginning!

Of course, most of the time, ALLOCDS allocated data sets without incident. But even then, there was a problem. Data set allocation done in JCL runs at the initiator address space's priority (very high) rather than the system's batch priority (typically much lower). Allocating thousands of data sets this way sometimes caused problems with online system response time.

The new ALLOCDS is different. First, it checks first to see whether data sets and catalog entries exist that would cause the job to fail. Then, it allocates most data sets using IDCAMS. If an allocation fails, the job stops with a non-zero return code and there are error messages about what happened. And the new job is restartable without editing; if you follow the restart instructions after fixing a problem, it will simply pick up where it left off and allocate the remaining data sets. Since most allocations are done by the jobstep program now, the work falls under the system's batch address space priority—which you control.

We also changed the RESTFS job that loads z/OS UNIX hierarchical file system data sets. (Level 2 told us this one job was

the single largest generator of service calls in the history of the planet. I think they might have been exaggerating...at least, a little.)

What changed? You no longer need to have UID(0) to run this job; you can run it with access to BPX.SUPERUSER. There are checks for things that could cause it to fail and messages to tell you what is wrong, too. Usually without having to look up arcane return and reason codes.

You will also find a minor but welcome change to the ALTCAT job. We took advantage of IDCAMS support to rename the data components of pagespaces. Now, the data set names you will see in VTOC and catalog listings for the same pagespace will be identical.

Starting with z/OS 1.3, you must use Workload Manager's goal mode... but, if you are still using compatibility mode (an IPS and ICS), a new job called IWMINSTL will help you load and activate the sample WLM service definition. It can be displayed and changed in WLM's ISPF panels immediately. This should help make it a bit easier to get to goal mode if you are not there already. Another new job helps you set up basic functions of the Integrated Cryptographic Services Facility (ICSF), paving the way for you to use SMP/E's new **RECEIVE FROM NETWORK** command.

To help balance the scales, we also got rid of some jobs. J2MERC, J2DELETE, J3MERC, J3DELETE, UPDCSI, FOMIS-CHO, and BPXISETS are all history. (We don't think you'll miss them.)

### Subsystem parmlibs

If you install CICS, DB2, IMS, or NCP using ServerPac, you will find they now come with their very own parmlib data sets. IEASVCxx, IEFSSNxx, LPALSTxx, and PROGxx members are built as needed in these new subsystem parmlibs. By default, each one is named using the subsystem's name, and they can be added to your parmlib concatenations to make migration—especially future migration—easier. For example, a parmlib concatenation might contain the data sets in the following example on a system used for CICS and DB2 workloads.

```
SYS1.PARMLIB          (your own parmlib)
SYS1.CPAC.PARMLIB    (ServerPac z/OS parmlib)
SYS1.IBM.PARMLIB     (SMP/E-managed parmlib)
CPAC.CICS.PARMLIB    (ServerPac CICS parmlib)
CPAC.DB2.PARMLIB     (ServerPac DB2 parmlib)
```

And, in IEASYSxx, you might specify:

```
LPA=(00,C1),
PROG=(00,C1,DB),      /* 00 for z/OS, C1 for CICS, DB for DB2 */
SSN=(00,C1,DB),
SVC=(00,C1)
```

Combined with symbolic aliases and system symbols, you can use subsystem parmlibs with mixed levels of subsystem software and shared master catalogs, too.

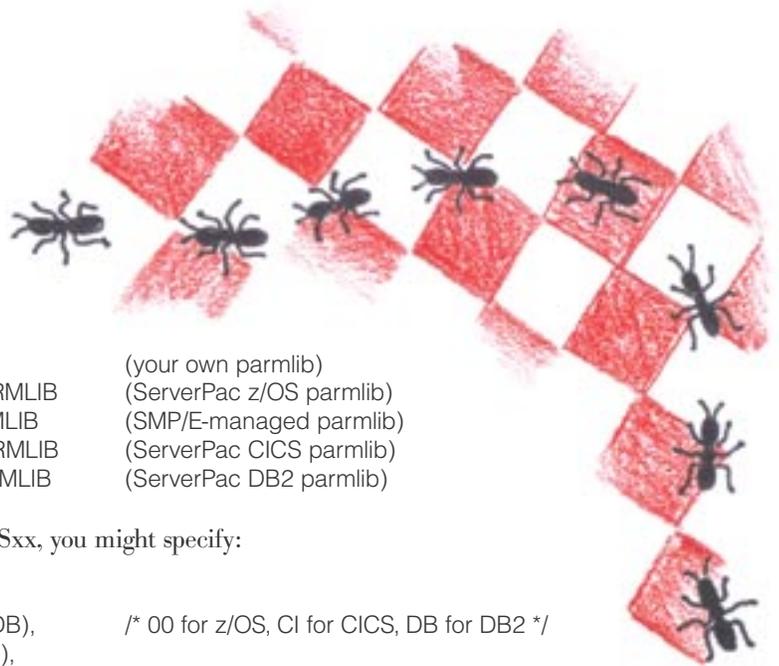
### msys for Operations

An analysis by z/OS Development showed that you can avoid certain outages using automation. To help you improve availability, msys for Operations was added to z/OS 1.2 and recently made available on OS/390 2.10. (msys is short for "Managed Systems Infrastructure".) It avoids and recovers from impending problems like console buffer and auxiliary storage shortages. But if you are like most people, you haven't set it up yet.

Starting with z/OS 1.4, ServerPac helps you set up msys for Operations. So you can get the benefits of automation more quickly. msys for Operations contains a subset of the function of the Tivoli NetView and System Automation for OS/390 products.

### 3590 support

ServerPac is available on 3590 tape now. If you order this way, expect just three tapes. That means there are fewer tapes for you to inventory, check into the tape library, and—probably, most importantly—fewer to mount while running the installation jobs.



### Watch this space!

What should you expect from ServerPac in the future? Well, we're thinking about some things like automatic block size setting (bye, bye, OPTIBLOCK?) and retrieving device characteristics based on volume serial, automatically. We're thinking about some of the little things too, like displaying data set space in tracks; changing most **Attrib** and **Assign** line commands to **Select** to be more consistent with other ISPF applications; taking the **SIZE** calculation out of **IODFLOAD**; and even spelling out most command names. Quick! What does L, F, N, P, or FC stand for? Gotcha!

### Requirements, requirements

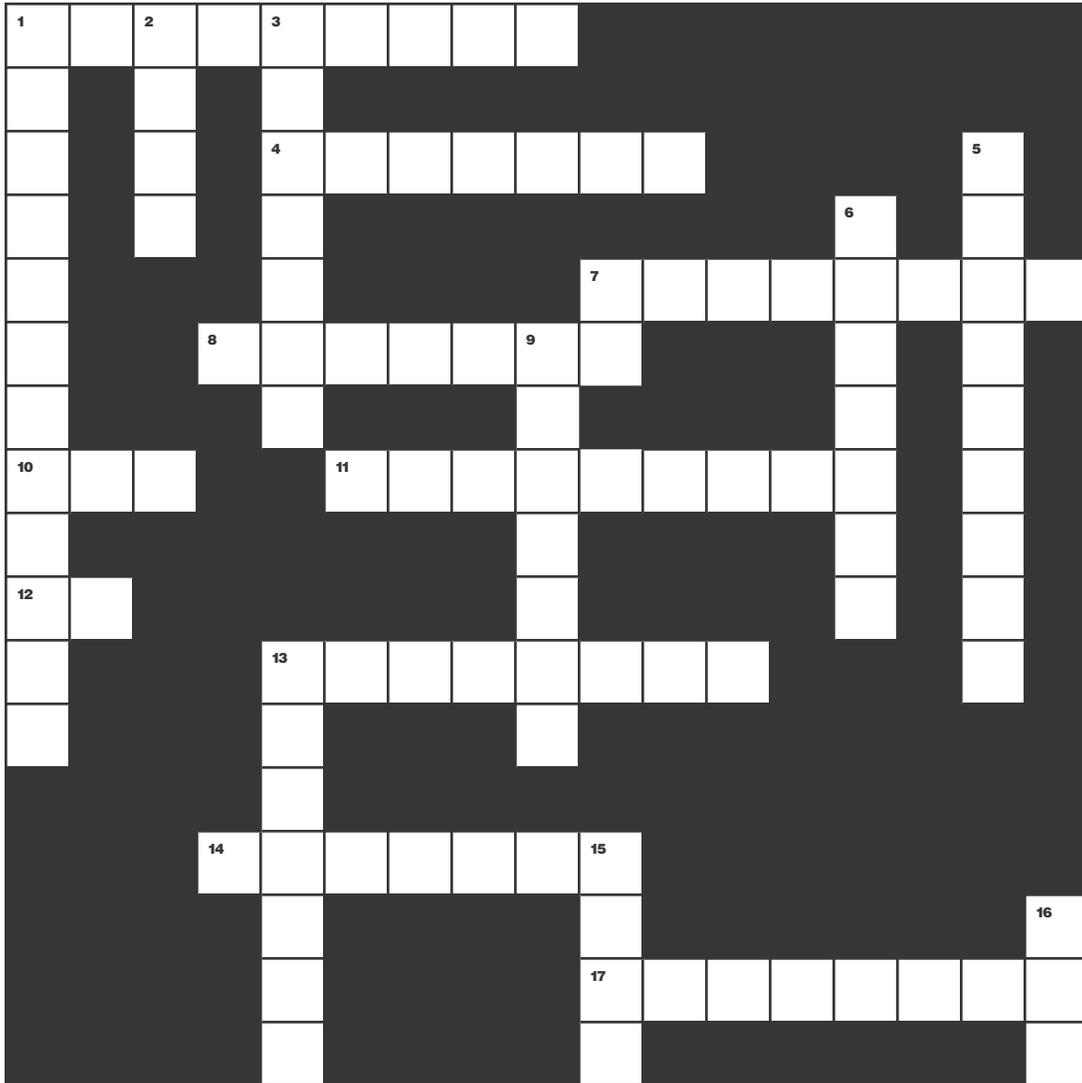
We hope you find the changes you just read about helpful, and hope they will save time during installation. With z/OS 1.3 and 1.4, we satisfied these customer-submitted requirements:

```
REQ00059273
REQ00061180
REQ00067676 (partially satisfied)
REQ00072235
REQ00076344
MR040901452
```

As always, we also encourage you to submit formal requirements for the changes you would like to see. They are the *best* way to tell us what you want!

# Crossword fun!

BY TRACY BONDI AND JOHN URBANIC



## Across

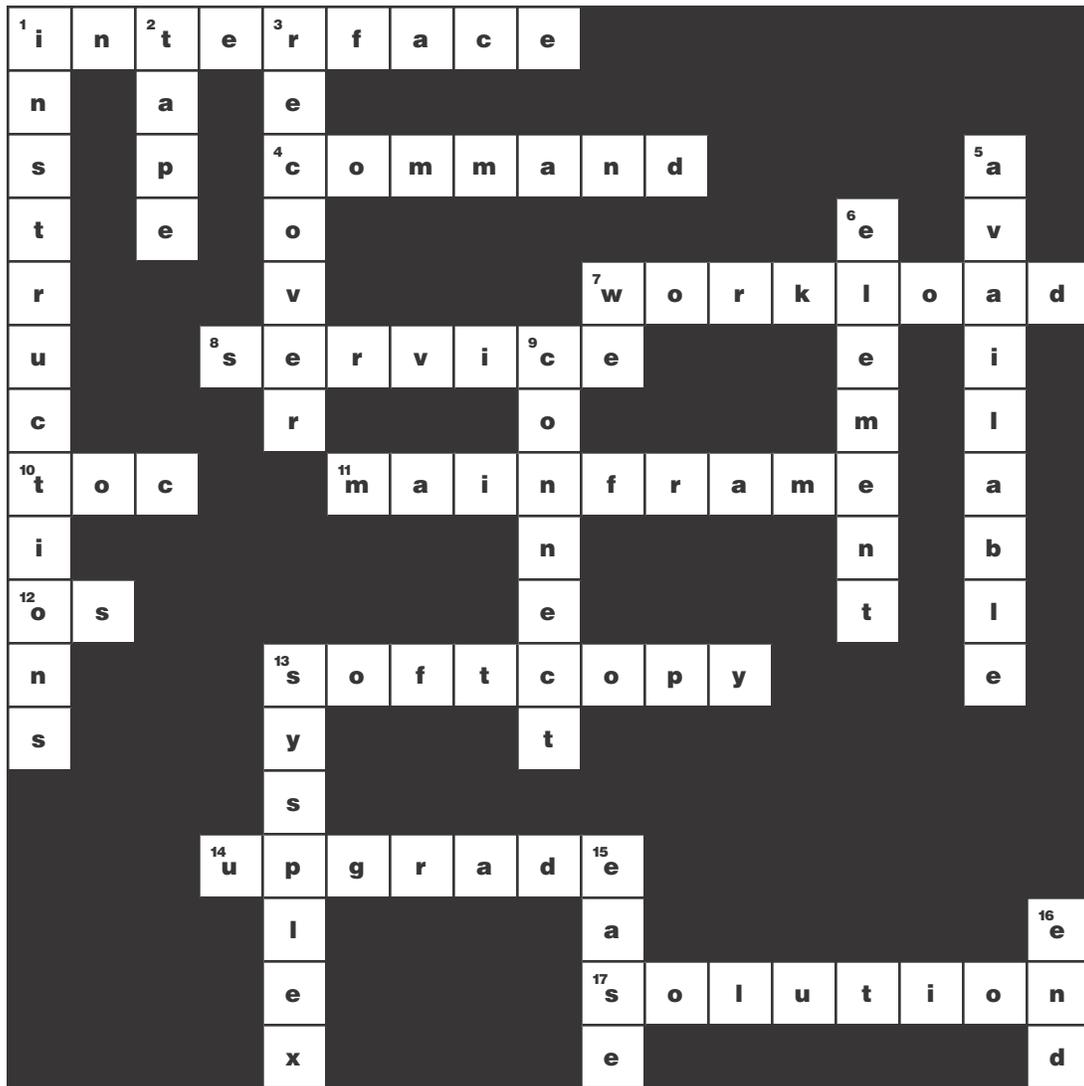
1. Go-between
4. Order
7. Job list
8. Fixes
10. Table of contents(abbr.)
11. Server, old style
12. Operation system(abbr.)
13. Online info
14. Improve
17. Answer

## Down

1. Manual
2. Back-up medium
3. Get back
5. Ready
6. z/OS piece
7. "Here \_\_\_ go!"
13. Many systems
15. "\_\_\_ of use"
16. Finish

*Crossword answers on  
next page...*

# Crossword fun ANSWERS!



## Across

1. Go-between
4. Order
7. Job list
8. Fixes
10. Table of contents(abbr.)
11. Server, old style
12. Operation system(abbr.)
13. Online info
14. Improve
17. Answer

## Down

1. Manual
2. Back-up medium
3. Get back
5. Ready
6. z/OS piece
7. "Here \_\_ go!"
13. Many systems
15. "\_\_\_ of use"
16. Finish



To use the CD: Insert it in any standard CD-ROM and it should start automatically.

If it does not, click on the **Start** button, choose **Run...** and then type **x:\index.htm** (where x is your CD-ROM drive letter) and press **Enter**.

### **z/Favorites!**

It's the zFavorites for zSeries credit card CD! You're gonna love this. It has all sorts of helpful Web links, like for:

- **Hardcopy**
- **Operating systems**
- **Software**
- **Language and tools**
- **ISV development and applications**
- **Product documentation**
- **Marketing information**
- **Education**
- **Support**
- **Links to FREE downloads**

Additional copies of zFavorites CD-ROM (GK3T-4331-02) are now separately orderable.

# See how they run: Automatic balancing of batch workloads with WLM 1.4

BY PETER BAEUERLE

In OS/390 2.4, the management of batch initiators, based on the goals of WLM goal mode, was introduced. In z/OS 1.4, this is improved with a new WLM functionality that extends the scope of batch management from mainly targeting the achievement of defined goals to the target of optimizing the distribution of batch workload over a sysplex—the batch initiator balancing functionality.

## How does WLM batch initiator management work?

It can be decided, per job class, if the initiators should be managed by JES or by WLM. If JES-managed, there is a fixed number of initiators available, each one assigned to one or more job classes. Incoming jobs can be started as long as there are idle initiators available.

In the WLM-managed case, the number of initiators is controlled by WLM and the initiators are assigned to service classes. WLM has several criteria for starting and stopping initiators:

- New initiators are started for a service class if they help to achieve the goals of that service class. That means, WLM takes the queue delay time of jobs waiting for execution into account when it checks in its 10-second policy adjustment cycle and it considers which resource the workload, classified to one service class, needs to fulfill its goal. If the reduction of queue delay time is found to have the best effect, new initiators for that service class are started. And those initiators are started first on systems with low load.
- Another criteria is independent of the goals. If there are jobs waiting in the queue to be started and a system has enough unused capacity—in terms of available CPU and memory—to take at least one of those jobs, an additional initiator is started on that system.
- WLM stops initiators if one of the following situations occurs:
  - if there are, over time, significantly more initiators available for a service class as are really used, or

- if a system is so heavily loaded, that available CPU or memory capacity have become short. In any case, WLM keeps at least one initiator per service class and system for at least one hour.
- Besides those major points there are some additional criteria for starting initiators. For example, if jobs are waiting that have system affinity to one or more systems, but no initiators are available for them, one initiator is started for those jobs.

## Where have the limitations been?

The description above shows that there is already some balancing of batch workload between the systems of a sysplex. When starting new initiators, WLM tends to prefer low loaded systems and tries to avoid heavily loaded ones.

But in a situation where there are enough initiators available and the load distribution might have changed, in comes a mechanism that works independently of the initiators being WLM- or JES-managed. A job can be started on any system, where it finds an idle initiator (assuming there is no system affinity that restricts it). Again and again, the situation is created where there are many initiators on constrained systems. As soon as one of them



becomes idle, because the job running under it has ended, it is taken over by a job from the queue that has just started on that system. As a consequence, many jobs can start on constrained systems, even if there are other systems that have much more available capacity. This is especially bad for jobs whose service classes have low importance or discretionary goals. They might start on constrained systems and rot there, because they cannot get much resource from the more important workload, while they would have been treated 'like VIPs' and run much faster on other, unconstrained systems.

## What is the improvement in z/OS 1.4?

The new functionality in z/OS 1.4 attacks this situation. The basic mechanism is that if there is a constrained system that has initiators of the service class for the waiting batch work, and there is at least one other system that has enough available CPU and memory to take over a waiting job of the average size of that service class, initiators are moved from the constrained system to the unconstrained one. See Figure 1.

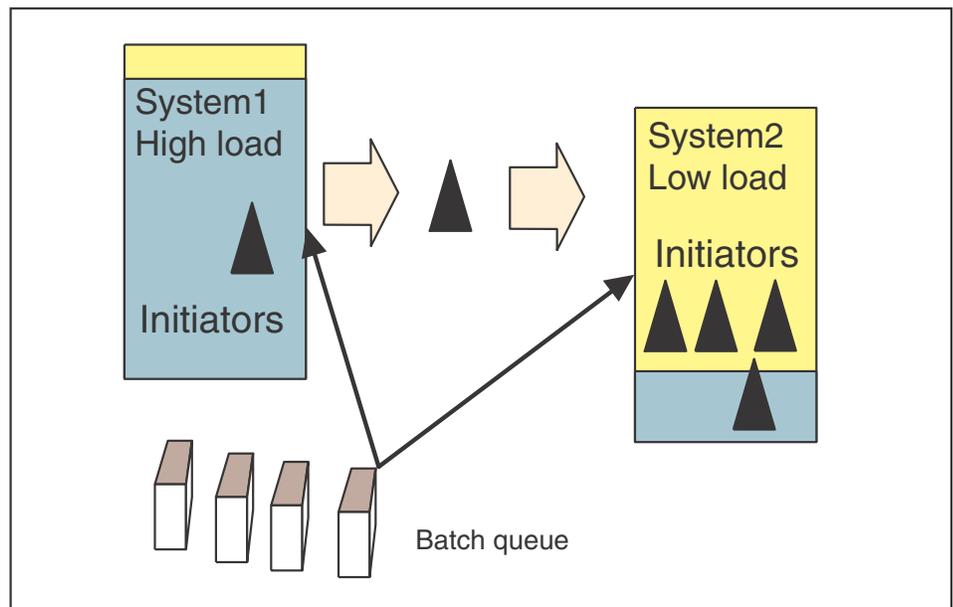


Figure 1: Balancing initiators to bring more batch to low-use systems

How it works in more detail:

At each 10-second interval, the following is done under the assumption that there are jobs waiting in the queue:

- Check if a system is constrained. A system is considered constrained when the CPU consumption is more than 95% based on a one-minute average. If one is found, the service class with the lowest importance, that has waiting jobs, and more than one initiator assigned to it, is identified as candidate for stopping an initiator.
- In this case, the other systems are scanned. If at least one is found that has enough memory and its CPU usage is so low that it would stay unconstrained after taking over a job of the identified service class, one initiator on the constrained system is stopped.
- If more than one system is over 95% CPU usage, only the most constrained system is a candidate for balancing an initiator down per interval.
- On unconstrained systems, up to five new initiators can be started per 10-second interval if there is enough CPU and memory available. (In previous releases, this maximum was one!)
- Special care is taken if there are jobs waiting that have system affinity for constrained systems and thus cannot be directed to unconstrained systems. In this case, the reduction of initiators is limited to maintain enough initiators for those kind of jobs.

This shows that the balancing mechanism is not a one-to-one kind in the sense that for each initiator stopped on one system, exactly one would be started on another one. It is a combination of asynchronously linked processes of reducing initiators on constrained systems and starting more on unconstrained ones. When we say “an initiator is stopped”, it is usually not really stopped immediately but the stop is pending until the job running under it has finished.

The initiator balancing functionality becomes automatically active under z/OS 1.4 with JES2 1.4 if WLM-managed initiators are used. If there are mixed releases in a sysplex, it balances only between those systems that are on z/OS 1.4 and JES2 1.4 or higher. Under JES3, it will be available on a later release.

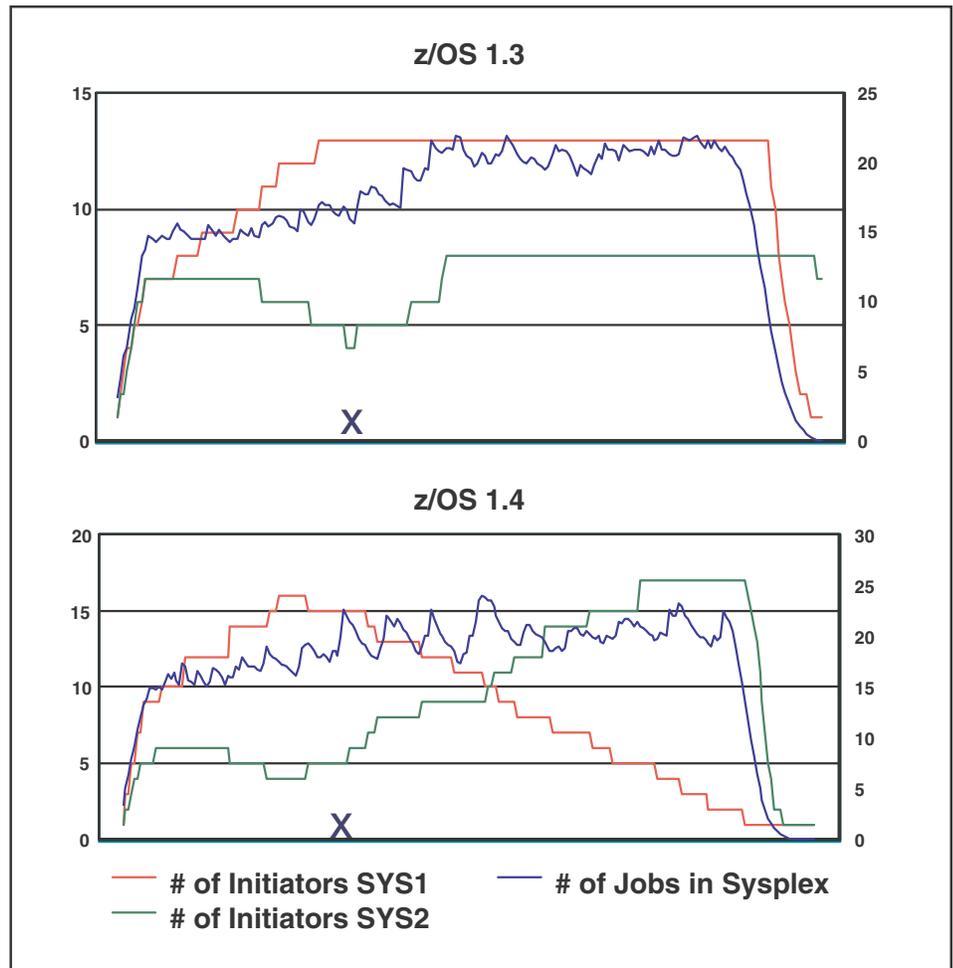


Figure 2 - Difference between z/OS 1.3 and 1.4 as SYS1 becomes constrained

Figure 2 shows an example how the old behavior compares to the new one. Two systems start with equal load, then one system, SYS1, becomes heavily loaded by other, non-batch, workload. The red and green lines are the number of initiators on SYS1 with respect to SYS2. The blue line shows the total number of batch jobs on the sysplex, either waiting or executing. (The charts do not show the total system loads).

While both runs behave similarly in the first phase, you see the major difference after SYS1 has become constrained at point “x”. In the z/OS 1.3 sysplex, the number of initiators stays on a constant level until the total number of batch jobs goes down. On the z/OS 1.4 sysplex, the SYS1 initiators are reduced step-by-step while they increase on the low loaded SYS2.

# More system automation magic!

BY MARY ELLEN COLEMAN

Alacazam! The wizards at IBM have cast another spell to make implementing your z/OS system easier. Their newest concoction is the Intelligent Resource Director (IRD) wizard. You may already know about IRD, and how it automatically shifts processing and I/O resources to the workloads that need them most. (See *z/OS Hot Topics Newsletter #4, February 2001, GA22-7501-00*.) If you haven't implemented this z/OS function yet, the new IRD wizard can help you.

If you've used other z/OS wizards, you'll find the IRD wizard's interface familiar. It has a series of four interviews that guide you through the information gathering process. The interview topics are:

1. Can you benefit from IRD?
2. Do you have the necessary hardware and software installed?
3. Are you running the correct microcode and software levels?
4. Implementing IRD

And, like the other wizards, it generates a report with step-by-step implementation instructions tailored for your installation. These instructions will get you up and running without requiring extensive knowledge of IRD. After you're running, you can refer back to the hardcopy documentation to learn how to adjust these basic settings to enhance your system performance. (See *z/OS Intelligent Resource Director Redbook, SG24-5952-00*.)

## Some neat features of this wizard

The IRD wizard includes a couple of special features. The interviews are arranged so that you can answer the first questions without having to do much research, helping you determine the value of IRD to your enterprise. After you finish answering all the questions in an interview, the wizard immediately takes you to the output report with your results. If IRD is valuable to you, you can then move on to the next interviews, which require more detailed information.

To assist you in your research, the IRD wizard includes an input worksheet that helps you gather the information you will

need for the interviews prior to actually using the wizard. It can be very annoying to get half way through an interview only to discover that there is information that you don't have, and you must go off and find it before you can continue. The worksheet is printable, and includes the interview questions along with space for your answers.

Now, let's take a look at the information the IRD wizard asks for.



## Can you benefit from IRD?

This topic asks basic questions about your system utilization, channel utilization, and channel assignments. You may be able to answer these questions without doing any research at all, and quickly get a feeling for whether you would benefit from implementing IRD or not.

## Do you have the necessary hardware and software installed?

This topic lists the hardware and software capital you need to run IRD. If you are missing something, the report will tell you. If you have everything, you can move on to the next interview.

## Are you running the correct microcode and software levels?

This topic lists the fixes you must install on your system before IRD can run. These questions are also included in the input worksheet. You may want to ask your IBM representative for help identifying the fixes already installed on your system. If you are missing anything, the report will tell you.

If you have everything, you can move on to the last interview.

## Implementing IRD

By this point, you have determined that you would benefit from implementing IRD, and that your systems are at the proper levels. This topic first asks you what IRD functions you want to implement—channel subsystem priority queueing (CSSPQ), WLM CPU management, and dynamic channel path management (DCM)—and then asks for detailed information about your sysplexes, LPARs, and CPCs. These questions are also included in the input worksheet, and may require the most research to answer.

After you fill in everything, the wizard generates a report for each IRD function you want to implement. For CSSPQ and WLM CPU management, the report provides implementation instructions, plus links to online documents for further information. For DCM, the report currently only provides links to online documents.

## A final thought

The IRD wizard offers one more link in a chain of IBM wizards that help simplify your task of planning for and setting up z/OS. They can save considerable time and effort, even if you decide you only want to use their output as a guide to doing your own setup. They can also help you start using a z/OS function by providing a basic setup strategy that you can modify as you gain more experience.

Customers who helped us test the IRD wizard found the wizard valuable, and liked the approach of answering simple questions first, before investing a lot of time in finding detailed answers.

You can find the IRD wizard along with the other z/OS wizards at the following Web site:

<http://www.ibm.com/servers/eserver/zseries/zos/wizards/>

Let us know what automated spells you cast with the IRD wizard! As always, we welcome your comments and suggestions.

# Waste not, want not: Reusing the softcopy collection interface

BY SUE VAN PARYS

Are you sitting with a repository of online books (especially a mix of BookManager and PDF files) and wondering how to give your users easy access to all the files? With real titles instead of just filenames, and books listed by product? Are you tempted to reuse the softcopy collection interface? Be our guest! Here are some guidelines to help you make sure the interface works in your environment. As a bonus, we've tossed in some tips to help you provide your users with the softcopy tools they may need to work with the online books.

## Reusing the softcopy collection interface for your own repository

After transferring online books from a softcopy collection to your own repository, you might want to provide the same access to these files as you have when the softcopy collection index (SCINDEX.HTM) is loaded within the SCFRAME.HTM frameset on a collection disc. The files can be:

- On your desktop
- On a local network (such as a network drive or a LAN server)
- Served by an HTTP server

Here are some tips for using the softcopy collection index with a softcopy repository.

- **Recommended browser:** To make full use of the support provided in SCFRAME.HTM, use Internet Explorer 5.0 or later.
- **For viewing BookManager files:** Install IBM Softcopy Reader 2.3.8 or later. This level of the product provides MIME-type information to simplify access to BookManager files through a browser interface.
- **For files on a network drive:** Map the network drive to a local drive letter. This ensures that the browser has access to the files. To provide access to the BookManager files on that drive, open Bookshelf Organizer and add the local drive letter and path to the list of user-defined directories (Edit >> User Defined Directories).
- **For files served by an HTTP server (Netscape users only):** To launch BookManager books and bookshelves with Softcopy Reader from the

Netscape browser, the HTTP server administrator must modify the HTTP server configuration.

Specifically, a MIME content type/subtype of application/boo must be assigned to the .boo file extension, and application/bks must be assigned to the .bks file extension. These assignments correspond to the content type values specified by the Softcopy Reader installation program in the .boo and .bks helper application information at the client. The method to make these assignments depends on the HTTP server. Here are some examples:

- For the Apache HTTP server, create AddType directives in the httpd.conf file.
  - For the Microsoft™ IIS HTTP server, create MimeMap registry key entries.
  - For the Lotus Go HTTP server, use the Configuration and Administration function to add new entries to the Mime Types form.
- Copy the contents of the root directory and “\AUTOHTML” directory from a numbered disc in the desired collection, using the same directory structure as on the collection.

For example, you could create a directory called “\IBMBOOKS” and nest the copied directories under it. Copy the files from the root directory of the collection into this top-level directory, and copy the contents of the “\AUTOHTML” directory of the collection disc into the “\AUTOHTML” directory under your top-level directory.

You should end up with a directory structure similar to this:

```
IBMBOOKS
├── BOOKS
├── SHELVES
├── PDF
└── AUTOHTML
```

- **For files not served by an HTTP server:** Delete the DISC.JS and CHKDISC.JS files from the “\AUTOHTML” directory on your system, and delete DISC.HTM from the top-level directory (such as IBMBOOKS). Deleting these files enables the SCINDEX.HTM file to work as if the softcopy files it links to are all on one disc.
- The links in SCINDEX.HTM assume that the book and bookshelf files are in the “\BOOKS” and “\SHELVES” directories, and the PDF files are in the “\PDF” directory. If you moved the files into other directories, you must edit the links in SCINDEX.HTM to point to those directories.
- The links assume that you are using SCINDEX.HTM and the files it links to (books, bookshelves, PDFs) are from the same drive. There might be situations when you cannot do this. For instance, if you are linking to books that are accessible on a network server, you must change the links to match the drive letter on which the linked-to files are mounted.
- To create a desktop shortcut for launching SCFRAME.HTM from within your top-level directory (such as “\IBMBOOKS”):
  1. Right-click SCFRAME.HTM.
  2. Drag it to your desktop.
  3. Select the option to Create Shortcut Here.
- Each time the collection is updated, refresh your repository from the contents of the root directory and “\AUTOHTML” directory of the collection discs. The interface is constantly being improved, and this is the only way to be assured of getting the latest code.

## Helping users install softcopy tools from a network drive

If you are providing a group of users with access to the files from softcopy collections, you might also want to set up a network drive with the softcopy tools these users might need. There are several way to do this:

- The simplest way is to copy the desired executables from the collection to the desired directory on the network drive. For example, to enable users to install IBM Softcopy Reader, you could copy ILRJAVA.EXE from the “\ILRJAVA” directory on the tools disc of the collection to a directory on your network drive. Users could then run ILRJAVA.EXE from the network drive to install IBM Softcopy Reader.
- To run the complete SCTOOLS.EXE interface from a network drive, you need to copy the complete directory structure. All the paths are relative. (The files and directories listed below are the only ones required for SCTOOLS.EXE to install tools. The other files and directories on the tools disc provide background information.)

<b>SCTOOLS.EXE</b>	-	program for installing softcopy tools
<b>\ACROBAT</b>	-	contains Adobe Acrobat Reader
<b>\AUTOHTML</b>	-	contains products.txt
<b>\ILRJAVA</b>	-	contains IBM Softcopy Reader
<b>\J2RE130 or \J2RE140</b>	-	contains Sun Microsystems JRE
<b>\SCL</b>	-	contains IBM SoftCopy Librarian

### Reminders:

1. SCTOOLS.EXE is written as a local application. It cannot run from an HTTP server and interact with users.
2. To ensure that you have the latest levels of the softcopy tools, run SCTOOLS.EXE from the tools disc each time you receive a new collection.
3. Some of the tools (especially IBM Softcopy Reader) are updated more frequently than the collections are. To check for new levels of the following tools between updates to collections, go to the associated URL:

IBM SoftCopy Librarian:	<a href="http://www.ibm.com/software/office/bkmgr/librarian.html">http://www.ibm.com/software/office/bkmgr/librarian.html</a>
IBM Softcopy Reader:	<a href="http://www.ibm.com/software/office/bkmgr/softcopyread.html">http://www.ibm.com/software/office/bkmgr/softcopyread.html</a>
Adobe Acrobat Reader:	<a href="http://www.adobe.com">http://www.adobe.com</a>
Sun Microsystems JRE:	<a href="http://java.sun.com/j2se/1.3/jre/download-windows.html">http://java.sun.com/j2se/1.3/jre/download-windows.html</a> <a href="http://java.sun.com/j2se/1.4/download.html">http://java.sun.com/j2se/1.4/download.html</a>



# What on earth is transmogrification?

## Graphics, transmogrification, and the z/OS Collection

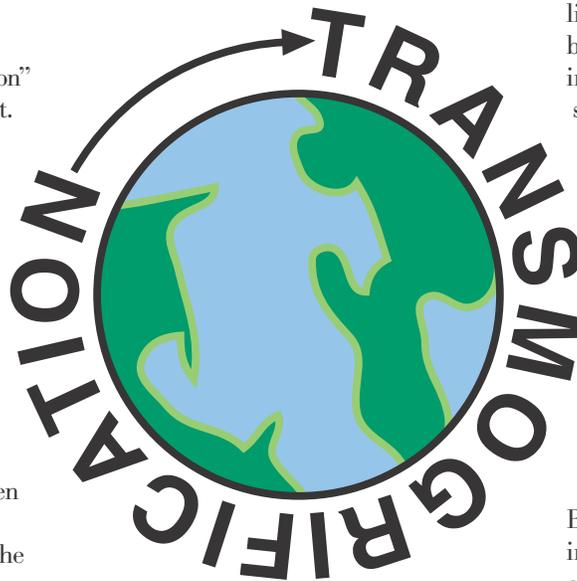
BY SHIRLEY SWENSON

What do these three things have in common? What on earth is transmogrification, and what's happened to the z/OS Collection?

**Graphics:** Beginning in September 2002, IBM began to upgrade the graphics in BookManager BOOK files to improve viewing with IBM BookManager BookServer on the Internet and with IBM Softcopy Reader. Previously you might have encountered an "Under construction" image for these graphics on the Internet. The upgraded graphics primarily affect BOOK files containing graphic figures and illustrations that were built with host versions of the BookManager Build products. The original graphics were copied, converted into a Web-ready format, and inserted into the BOOK files for use with BookServer and Softcopy Reader. The original graphics still remain in the files for viewing on host platforms. To benefit from the upgraded graphics when viewed with IBM Softcopy Reader, you should use Version 2.3.6 or later. Visit the BookManager Web site at <http://www.ibm.com/software/office/bkmgr/> to get the latest version.

**Transmogrification:** The tool used to convert the host-built graphics to a Web-ready format is the IBM Book Graphics Conversion Utility (also known as the TransMogrifier). This free tool can be downloaded from the BookManager Web site and used to "transmogrify" the graphics contained in any Version 1.3 IBM BookManager BOOK. The graphics in a BOOK must be converted for the Linux™ version of BookServer 2.3 to render the graphics. For the AIX® version of BookServer 2.3, the conversion is only necessary if the Version 1.3 BOOK contains graphics in vector format.

Converting the graphics and inserting them into a BOOK file that already contains the host-built graphics obviously increases the size of the affected BookManager files, sometimes as much as 20 percent or more depending on how many graphics are included.



The increased file size caused disc overruns for some collections, so we had to reorganize the files across the discs.

**z/OS Collection:** The collection affected most was the *z/OS Version 1 Release x Collection*, SK3T-4269. For September 2002, we reorganized this collection with product libraries sorted alphabetically by title across the four discs and with both the BookManager and PDF formats of each book on the same disc. Early feedback from users, who could no longer search across the z/OS libraries that they use frequently, caused us to put the BookManager files together for the December 2002 edition, albeit spread across two discs because of the increased file sizes. This also resulted in the PDF files spread alphabetically across the remaining two discs.

Unfortunately, we can no longer fit all, or even most, of the BOOK files for a single release of z/OS on a single disc. For those users who have DVD drives, there is some relief. The *z/OS Version 1 Release x and Software Products DVD Collection*, SK3T-4271, includes the entire z/OS release library, all the z/OS software product libraries, and selected IBM zSeries Redbooks on a single disc! And, we built and included a special all-z/OS (release) bookshelf and index so you can search across the entire z/OS library for a release when using the DVD. This same all-z/OS bookshelf is used on the z/OS Internet library at <http://www.ibm.com/servers/eserver/zseries/zos/bkserv/>.

### Transmogrification forever?

Following each regular update cycle for the collections, we are transmogrifying, or upgrading, the graphics that were built with host versions of the BookManager Build products and replacing the BOOK files for use on future collections and the Internet. In 2002 we upgraded several thousand BOOK files with improved graphics. Eventually we plan to include the transmogrifier in the host build products or in the submission process for BOOKs.

So, enjoy those improved graphics that you're seeing on the Internet library and on your workstation. Graphics in Web-ready format are especially important to users who are relying more and more on the Internet library or accessing books on a company intranet using IBM BookServer.

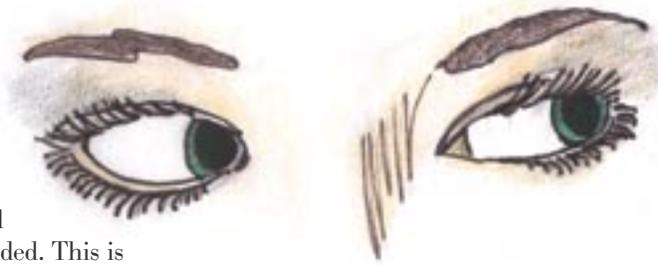
# e<sup>e</sup> E<sup>E</sup> I'm being watched? User identity mapping services for enabling secure heterogeneous applications

BY FRANK HILDENBRAND, PEGGY LABELLE, AND JASON WILLIAMS

As computing environments expand and services execute across platforms, the capability for systems to self-optimize, heal, configure, and protect information is needed. This is what started IBM's autonomic computing initiative. In support of this, IBM eServer is providing the base infrastructure and programming model to enable products and applications to run on both IBM and non-IBM systems. This will help them to function and interoperate efficiently in this heterogeneous computing environment. To address some of the security aspects of this heterogeneous environment, particularly the problem of dealing with multiple user registries, IBM delivered Enterprise Identity Mapping (EIM) on all IBM server platforms (z/OS, OS/400®, AIX, Linux on Intel, and Windows™ 2000) in 2002.

EIM enables application providers (such as vendors, customers, business partners) to address the problem of multiple user registries more easily and with less cost than previously possible. Customers can write multi-tier applications that never have to ask the user for a platform-specific user identity and password when accessing legacy data. This capability eliminates the need to cache or send passwords over the network as is typically done with these types of solutions. Furthermore, the data will continue to be protected by IBM's highly respected, platform-specific security services.

Many security vendors are focused on solving the problem of having multiple user identities defined in different user registries (such as SAP, Lotus Notes™, RACF, AS/400®, Kerberos, PKI, and LDAP) using some form of single sign-on technology. Current single sign-on products are focused on hiding (but not eliminating) the complexity of having a user log on to many systems. This does not address the growth in distributed applications where the application spans multiple servers which need to relate the end-user to local user identities at the servers. EIM is working to



solve this problem by providing common system services to allow the association of one person to multiple user identities in existing security registries. Instead of an application needing to create side files that define these associations, the mapping associations are stored in one centralized place, the EIM domain controller. These associations will allow the multi-tiered applications and servers to understand who the end user is in their own terms, and allow access control and auditing to be done with the local identity. This technology is built on a LDAP-accessible directory service which is now shipped in all eServer operating systems. Using a directory service allows the associations to be used and managed in a distributed environment.

## The EIM approach

EIM is an open architecture that administrators can use to represent identity mapping relationships for any registry. EIM offers a new approach to enable inexpen-

sive solutions to easily manage multiple user registries and user identities in an enterprise. The ability to map between a user's identities in different user registries provides multiple benefits. Primarily, it means that applications can have the flexibility of using one user registry for authentication while using an entirely different user registry for authorization, with no code changes required to existing user registries. The administrator does not need to have mappings for all identities in a user registry because:

- EIM allows one-to-many mappings (in other words, a single user with more than one user identity in a single user registry).
- EIM allows many-to-one mappings (in other words, multiple users sharing a single user identity in a single user registry, which is supported but not advised).

An administrator can represent any user registry of any type in EIM.

## EIM concepts

A conceptual understanding of how EIM works is necessary to fully understand how you can use it in your enterprise. Although the configuration and implementation of EIM APIs can differ among server platforms, EIM concepts are common across IBM eServers.

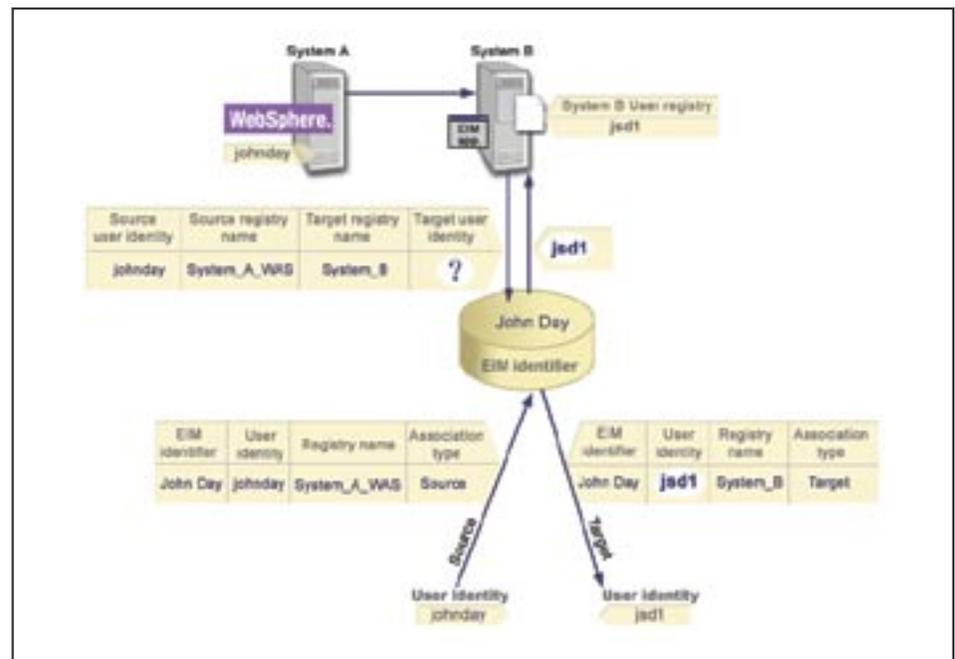


Figure 1

Figure 1 provides an EIM implementation example in an enterprise. A server acts as an EIM client and contains an EIM-enabled application that requests EIM data using lookup operations. The domain controller stores information about the EIM domain, which includes an EIM identifier, associations between these EIM identifiers and user identities, and EIM registry definitions. The domain controller may reside on the same system or a different one.

**EIM domain controller:** The EIM domain controller is an LDAP server that is configured to manage at least one EIM domain. A minimum of one EIM domain controller must exist in the enterprise. Currently, you can configure a subset of IBM eServer platforms to act as an EIM domain controller see table on page 37. Any system that supports the EIM APIs can participate as a client in the domain. These client systems use EIM APIs to contact an EIM domain controller to perform EIM lookup operations.

**EIM domains:** An EIM domain is a directory within an LDAP server that contains EIM data for an enterprise. An EIM domain is the collection of all the EIM identifiers, EIM associations, and user registries that are defined in that domain. Systems (EIM clients) participate in the domain by using the domain data for EIM lookup operations.

An EIM domain is different from a user registry. A user registry defines a set of user identities known to and trusted by a particular instance of an operating system or application. A user registry also contains the information needed to authenticate the user of the identity. Additionally, a user registry often contains other attributes such as user preferences, system privileges, or personal information for that identity. In contrast, an EIM domain refers to user identities that are defined in user registries. An EIM domain only contains information about the relationship between identities in various user registries (user name, registry type, and registry instance) and the actual people or entities that these identities represent (an EIM identifier).

**EIM applications:** EIM applications look up mappings in an EIM domain. For example, a basic client application that uses EIM mappings might authenticate the end-user to system A and then send a request to a server on system B with that user identity. The server on system B calls EIM services to find the mapping between the user identity from system A and the local user identity, and accesses application resources using the local user identity.

In order to make this work, an EIM administrator needs to perform the following actions in the EIM domain:

1. Create EIM identifiers that represent people or entities using the application.
2. Create EIM registry definitions that describe the user registries.
3. Define the relationships between the user identities in those registries and the EIM identifiers:
  - a. Source associations between the user identities in system A and the appropriate EIM identifiers
  - b. Target associations between the EIM identifiers and the user identities on system B.

OS/400 uses this approach to support Kerberos identities in its telnet server. The telnet client authenticates the user on Windows 2000 and sends a Kerberos ticket to the telnet server running on the OS/400. The telnet server uses EIM to find the OS/400 user ID for the end user.

Some other possible applications of EIM are:

- Storing mappings between user identities defined in different z/OS security management products
- Using mappings to track the activities of an enterprise user by analyzing audit logs generated on different enterprise platforms
- Tracking which user identities assigned to someone within an enterprise

## Customer interest and quotes

Several customers have publicly expressed their interest in Enterprise Identity Mapping (EIM). Some companies have gone as far as to publish statements of direction for developing applications that exploit this new technology. Some of the companies that have expressed their intentions to deploy EIM applications include:

- Kim Greene and Associates has a special announcement that highlights plans for EIM:  
<http://www.kimgreene.com/>
- Tango/04 has a statement of direction (SOD) for integrating EIM  
<http://www.tango04.com/news/newsdetail.php?id=153>

Customer comments on EIM include:

- “Security is on all our minds today. We at Danske Bank understand that our data and applications are only as secure as the weakest link in a distributed heterogeneous environment. Managing a consistent set of security related information and policies is a “weak spot” of any infrastructure. IBM’s Enterprise Identity Mapping technology facilitates a consistent user identity across our whole infrastructure, which enables us to centrally define and manage users and user policies thus strengthening our security chain.”  
- *Claus Torp Jensen, Vice President of Systems Architecture at Danske Bank and z/OS customer.*
- “OS/400 V5R2 products like Enterprise Identity Mapping (EIM) on iSeries™ that allow for a single logon across multi-platforms...will reduce wasted Help line calls significantly—perhaps by 25% or more.”  
- *Stan Wilkins of Realsolutions*

## EIM availability and component details

The following table identifies the operating system and component release information for EIM.

Platform	Version	EIM domain controller	EIM client
OS/400	5.2 available 8/30/2002	Included	Included in OS/400 release V5R2
z/OS	1.4 available 9/27/2002	z/OS 1.4 Security Server LDAP - included	z/OS 1.4 Security Server Enterprise Identity Mapping (with SPE OW57137)
AIX	V5R2 available 10/25/2002	To be announced	Included in AIX release 5.2: utilizes IBM Directory 4.1 client
Windows 2000	Available 11/15/2002 via Web download: <a href="http://wcs.haw.ibm.com/servers/eserver/security/">http://wcs.haw.ibm.com/servers/eserver/security/</a>	To be announced	Included in deliverable: utilizes IBM Directory 4.1 client
Linux	Available 11/15/2002 via Web download: <a href="http://wcs.haw.ibm.com/servers/eserver/security/">http://wcs.haw.ibm.com/servers/eserver/security/</a>	To be announced	Included in deliverable: utilizes Open LDAP v2.0.23 and IBM Directory Server 4.1

# Line 'em up for Domino 6: Better performance, easier migration

BY JOSEPH BOSTIAN



Domino 6 for z/OS is the most recent version of Lotus' flagship groupware product. It provides a long list of improvements over Domino R5, and features substantially enhanced performance as well as an install and setup process that simplifies migration a great deal. These improvements were achieved through a combination of core Domino enhancements, and the exploitation of z/OS-specific platform features such as XPLink. We will briefly outline these improvements, and show how much simpler it is to migrate from Domino R5 to Domino 6 than it was from Domino R4 to R5.

## Performance and scale improvements

One of the most compelling features of Domino 6 on z/OS is performance. It's the first release to exploit XPLink—the new high performance linkage that substantially reduces the call/return path length between one C/C++ subroutine and another. The large majority of Domino is written in C and C++, and it is sensitive to such overhead, because like many middleware products, it has been designed to run on multiple platforms. Platform independence is a strength of the product that requires well-defined internal interfaces, and modularity that would not be necessary for software that runs on a single platform.

This modularity comes at some cost, because a structured, modular design necessarily will insert some call overhead to isolate data, and provide clear component interfaces. The default call linkage provided by z/OS adds to this pathlength for similar reasons—it provides a well structured, relatively simple, and very robust means to pass control between subroutines. This linkage architecture is what often allows for detailed problem determination as well, because it makes parameter passing explicit, and call stacks easier to locate in a dump.

XPLink has required substantial changes to the z/OS binder, LE runtime library, and z/OS compiler, and we provided important feedback to the development groups for these components as XPLink was being implemented. Domino for z/OS

required changes to a handful of performance-critical assembler routines, and the process used to build the product, as well as the procedures and tools we use for problem determination. XPLink did not require us to make any significant changes to our C or C++ source.

In addition to the changes required for XPLink, enhancements were made to the compiler to generate more optimal code, and exploit the new zSeries hardware architecture. These changes provided additional path length reduction, which has further improved Domino 6 server performance relative to R5.

Our first indications of the performance gains we could expect from XPLink and associated changes came in late 2000 and early 2001, when we ran an R5 Domino server on an OS/390 2.10 ssystem with early versions of the binder, LE, and the compiler. The results indicated that by simply building for XPLink, and using an XPLink'ed runtime library, we would get an 18% CPU reduction. About this same time our efforts on Domino 6 began to ramp up, and we began focusing our attention on that release. We also kept up with the continuous improvements that were being made to z/OS relative to XPLink, particularly to the compiler and runtime library. By June of 2002 we were seeing Domino 6 servers under a Notes client mail workload consistently using 25% less CPU than an R5 server under an identical load. When we used an iNotes™ Web Access workload, this gap grew to 35%. This growth in savings from 25 to 35% was primarily due to the redesign of the Domino HTTP server. These results have held up in both the lab and in IBM production environments. Our experiences with XPLink have exceeded our expectations.

## Migration, installation and setup

Migration to Domino 6 from any previous release is substantially simplified through two improvements in the install and setup process. The first of these is provided by the multi-version install feature, which allows an LPAR to support one server running at R5 or earlier, and multiple Domino 6 servers simultaneously.

This feature is made possible by changes to the Domino 6 server, which allow the product executables to be installed at a configurable location, and not forced to the standard Lotus install directory (/usr/lpp/lotus). Domino 6 also exploits the new shared library support of z/OS, which allows multiple versions of common Domino executables to be fetched into shared library regions simultaneously. z/OS shared library not only enables this binary coexistence, it simplifies installation by eliminating the need to put Domino executables into the link pack area (LPA).

Multi-version installation eases the migration from Domino R5, by allowing Domino 6 to live on the same LPAR as the existing server. When the migration to Domino 6 is complete, the R5 server can be shut down, and production can begin. Without multi-version install, the administrator is required to carve out another LPAR for the new Domino server instance.

In addition to eliminating the new LPAR requirement for installation, Domino 6 includes a new Java-based, remote server setup interface that is common to all platforms. This common interface is much improved, and will be familiar to Domino administrators who may not have a strong zSeries background.

## Other improvements and enhancements

There are many other enhancements to the Domino 6 server that can't all be listed here. Some of these are specific to z/OS—particularly in the area of recoverability, availability, and serviceability (RAS). Many other improvements have been made to the core Domino product available on all platforms. For a complete list of these improvements, see the *Domino 6 Technical Overview*, available at the Lotus Developer Domain, <http://www.lotus.com/ldd> (formerly Notes.Net). For Domino information specific to z/OS, including service level prerequisites, online documentation, and performance information, see the Lotus Domino for z/OS home page, <http://www.ibm.com/server/eserver/zseries/software/domino/>.

# Availability: Getting better all the time

BY: ROBERT ABRAMS, ED MEKEEL, AND ROY MOEBUS

High availability has become fundamental to success in today's marketplace, especially with the growing investments in system configurations to meet the escalating demands of e-business and the development and execution of applications for managing business-critical customer information systems. Over the years, MVS and OS/390 have delivered the software foundation for highly available systems. With that strong foundation, we have amassed significant new functionality and have implemented an aggressive process to continuously raise the bar. Many of these new functions have been the result of scenarios you've pointed out to us.

In addition, as part of IBM's commitment to strengthen the availability characteristics of z/OS, we have implemented an elaborate evaluation process for every multisystem outage reported to IBM across the entire z/OS platform. Each outage is analyzed to eliminate the error that triggered the event as well as strengthen the RAS attributes of the system, permitting recovery from similar errors. As a result, a significant, ongoing stream of enhancements have been, and continue to be, implemented exclusively on z/OS. The intent of this article is to highlight many of these enhancements. Areas of focus include hang detection and analysis, automated system and sysplex recovery, operations usability, preventing the depletion of system resources when handling extremely high system activity, and minimizing the need for planned outages by supporting dynamic changes to system resource definitions. Read on for examples of the function that has been delivered in these areas.

## **Help has arrived...**

### **Operations usability and recovery**

Operations assistance is probably one of the most important areas related to avoiding outages. It is very important to have the system recover properly from problems and automatically avoid outages, but there are still lots of things that an operations

staff needs to deal with, no matter how well-automated the environment is. Issues we've seen range from mundane typing errors, to very complex situations where much knowledge is expected of the operations staff, and simplicity is long overdue. Operations-level recovery tasks in z/OS vary, but many of these tasks are addressed by Managed System Infrastructure for Operations (msys for Ops).



An element of z/OS, msys for Ops increases the availability of your systems by helping to reduce the number of unnecessary IPLs. It accomplishes this through improved problem detection and analysis, and by initiating recovery. msys for Ops uses a panel interface to provide a consolidated view of selected system and sysplex resources. The panels enable you to display detailed information and to initiate actions directly and easily. Panels are provided for displaying system details, displaying console information, managing coupling facilities and couple data sets, displaying IPL information (for example, parmlib members), and specifying options for initiating a system dump. The panel interface makes it much easier to deal with multiple systems, and more importantly, it is a more effective way to present data that can affect availability and to subsequently take actions to correct or recover from potential problems.

The msys for Ops interface greatly reduces the complexity associated with issuing sysplex or system commands, providing an action invocation by either a PF key or single-character command. msys for Ops also provides a sound foundation for automating a wide variety of recovery actions, based on system messages. The automation functionality in msys for Ops also avoids potential errors introduced when recovery requires

a series of tasks that increases the opportunity for errors. Some of the automation related to availability and recovery provided with msys for Ops includes

- WTO /WTOR buffer shortage recovery, detection and elimination of long ENQs, dynamic creation of local page data sets, canceling of jobs that are rapidly increasing storage use, and partitioning a system out of the sysplex.

Along with important availability functions delivered in msys for Ops, we've enhanced other system components to help in the operations arena.

For JES2 we've delivered a resource monitor which can detect why JES2 is not responding to JES2 commands and why key JES2 functions aren't functioning. It can expose impacts to the JES2 main task and look at the level of resource usage. The JES2 monitor function will take an appropriate action based on the duration of the impacting event. Additionally in JES2 we have enhanced checkpoint recovery processing to protect from starting jobs when the wrong spool or checkpoint data sets are accidentally used.

In the consoles area, we've created a restartable command address space to prevent outages caused by command flooding in areas like automation programs. In addition, we enhanced SYSLOG to unconditionally close even though OPERLOG or printer consoles are not available. This prevents the WTO storage from filling while waiting for a hardcopy console to become available.

DFSMSdftp™ provides the ability to separate different data sets in a predefined set from each other by allocating them on separate physical DASD control units. This allows the installation to ensure that an outage on one physical control unit does not affect both the primary and alternate versions of key system or user data sets.

The z/OS Communication Server Sysplex Distributor is enhanced to improve the efficiency of its load balancing algorithms to improve detection of connection failures. In addition, enhancements are included to monitor Communication Server's use of common (CSA) storage.

### **Who's blocking traffic?... Hang detection and analysis**

A software hang is usually defined as some function that cannot complete because it is waiting for some other process to complete first. The first place that we think of when locating and resolving hangs focuses on identifying the work contending for some system resource that has been obtained, or serialized, by another program and is waiting until the current owner releases it. Resource serialization mechanisms include ENQs, latches, reserves and internal locks, each with varying levels of impact.

msys for Ops helps here also with functions to identify the cause of a long blocking task and terminating it. Policy is provided to allow the installation to identify batch jobs and address spaces that are important to enable msys for Ops to take action when deemed appropriate. The ENQ analysis employs the same function that backs the **DISPLAY GRS,ANALYZE** command. msys for Ops also provides for hung command recovery. This function detects critical ENQ contention which if held for extended periods will cause commands to hang. Similar to Active ECA recovery, the degree of disruptive recovery is based on policy settings coded by the installation. Eligible hung command processes are terminated after a dump has been captured.

WLM also has functions to help in this area. WLM ENQ management integrates knowledge of the business importance of the waiters for the resource into the decision to reallocate CPU to the resource holder so that it can finish and release the resource. Servers processing heavy updates concurrently with on-line workloads, like those found in business intelligence

environments, would derive the most value from WLM's sensitivity to ENQ management.

The ability to recognize bottlenecks and be able to take action without affecting all other work is critical to continuous availability. The System Logger component provides the ability to identify an offload activity being inhibited (either hung, or not operating for some reason) and allows for the log offload instance to be stopped so work can commence on other systems.

### **Get well soon... Automatic system and sysplex recovery**

z/OS reliability and availability is ensured through an elaborate system recovery framework to ensure that errors can be recovered and erroneous work is isolated from the rest of the system. However in a sysplex, recovery extends to ensuring that system resources are managed with integrity on each affected system. This includes speedy management of devices and shared data structures.

In z/OS UNIX System Services, we provide recovery options for customers to choose which system takes over mounts that have been dropped when the current mount owner leaves the sysplex. This provides for improved control over the mount file locations in a sysplex.

msys for Ops has functions to help in managing sysplexes. For example, it provides for failed system isolation. This function removes failed systems from a Parallel Sysplex®. Trigger points are IXC102A and IXC402D messages. This recovery requires that the failed system image first be hardware-deactivated before replying 'DOWN' to the trigger message.

We have also enhanced our GDPS offering with our new hyperswap support. The actual hyperswap function performs a switch from primary DASD to secondary DASD in the event of a recovery action. The "hyper" means this swap is performed at speeds in orders of magnitude faster than in the previous DASD swap functionality. We are also considering further enhancements to this function in the future.

### **It's a corrupt world we live in... Preventing corruption of system resources**

Proper Parallel Sysplex setup and configuration are critical success factors in realizing highly available systems. The most common issues are associated with management of the sysplex topology. Considerations include sufficient redundancy, the placement of critical data sets on shared DASD and the management of message flow between systems.

We've answered several requirements with page data set protection which prevents two systems from accidentally using the same physical data set. When this occurs it's very difficult to detect and diagnose, and can cause data integrity and/or system failures. This support will prevent the same data set from being used when possible, and where not possible it will protect the integrity of the system by causing system termination when it's detected.

msys for Ops ensures that the primary and alternate couple data sets (CDS) are on different volumes to prevent slowdowns. Checks are performed to ascertain whether there are single points of failure with regard to couple data sets being allocated on the same volume or behind the same storage director. When discovered these situations are highlighted on the CDS management panel displays. Detection of message IXC253I (primary or alternate CDS removal) causes CDS switching and also results in the dynamic creation of a new alternate. During this activity, msys for Ops ensures that separation between the primary and alternate CDS placement is observed.

### **Don't make me IPL... Minimizing planned outages**

Much of this article is focused on eliminating the causes of unplanned outages that IBM has observed from analysis and have been reported as customer requirements. But minimizing your planned outage window, or avoiding an IPL to make such changes, is equally important.

In this area, Logger features enhancements to allow LOGR policy attributes to be changed without requiring an IPL, provided that the log stream is not in use at the time of the change. When setting up the System Logger, it is generally required that VSAM SHAREOPTIONS(3,3) be

specified to ensure the appropriate control characteristics of the log data sets. This option is now enforced to avoid the possibilities of log data browse failures, log stream recovery failures, and even possibly loss of data.

Since the addition of z/OS UNIX as part of the base, customers have been looking for the ability to recycle OMVS without requiring a full system IPL. UNIX-based applications causing system errors will be able to start up quicker by not having to wait for a scheduled system IPL.

### **Summing up and looking ahead**

As you can see from what we've described here, we've enhanced several very important areas of z/OS. This article covers functions delivered over several releases. We've really only been able to hit the highlights here and everything mentioned in the preceding sections, and more, is available with z/OS 1.4. We're pleased to say there's more to come. z/OS is the exclusive base for future high availability functions within IBM's major subsystems. In addition, we continue to plan z/OS releases to include capabilities aimed at helping to reduce the number of things that could impact your availability windows.

For example, keep your eyes open for more refinements in the consoles area. We plan to enhance consoles so that there's always a console available to receive critical messages during system initialization. Use of the hardware system console would be disabled only when MCS, SMCS or EMCS consoles become available. In fact, IBM will be embarking on a strategy to reorganize the consoles component over time to ensure that it is not an area of constraint as your sysplexes continue to grow. The major areas of focus will be to improve the performance of message delivery and consumption, synchronize the console state information across the sysplex and to remove the limit of 99 MCS, SMCS and subsystem consoles.

We'll also be looking to streamline VARY OFFLINE processing to handle situations where multiple devices are varied offline across multiple systems.

You can continue to look for enhancements in our automated operations element, msys for Ops, to remove potential impacts to sysplex high availability.

So, hopefully it's clear to you after reading this article that we here in "z/OS Land" are serious about bulletproofing and building upon our solid foundation for high availability processor environments. We sincerely hope that as you formalize your z/OS migration plans, you have the added confidence that you will benefit from these additional z/OS availability items and that z/OS 1.4 is your target release.

# It takes two to make it right: System-Managed CF Structure Duplexing

BY: KELLY PUSHONG AND DAVID RAFTEN

The Parallel Sysplex technology provides many benefits to the zSeries environment, including high availability, workload balancing, scalable growth, reduced cost of computing, ease of use, and investment protection of current applications. System-Managed Coupling Facility (CF) Structure Duplexing enhances this technology further by providing a general-purpose, hardware-assisted mechanism for duplexing CF structure data. Available on z/OS 1.2 and higher via an enabling APAR, System-Managed CF Structure Duplexing provides an easily exploited common framework for duplexing the structure data contained in CF structures, with installation control over which structures are or are not duplexed. When System-Managed CF Structure Duplexing is enabled, a duplexed copy of the CF structure is created prior to any failure. In failure scenarios, rapid failover to the unaffected structure instance of the duplexed pair is achieved with very little disruption to the ongoing execution of work by the application. The benefits of System-Managed CF Structure Duplexing include improved application availability, improved manageability and usability of a Parallel Sysplex environment, possible configuration simplification, and potential financial benefits. Interested? Read on...

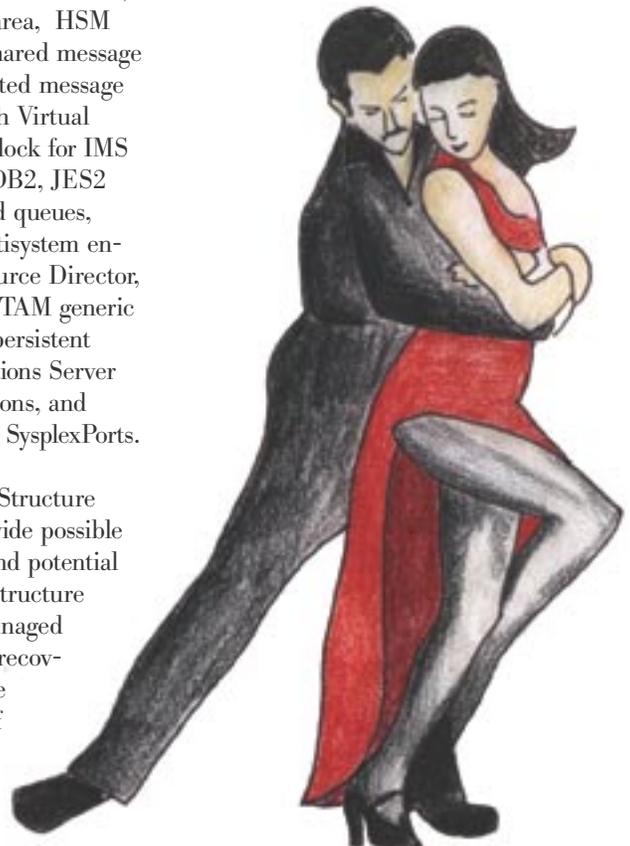
One of the major benefits of System-Managed CF Structure Duplexing is its ability to rapidly recover from situations where there is a loss of connectivity to a coupling facility or a coupling facility failure. This rapid recovery therefore results in improved application availability. At an IBM test facility, a test system was configured to simulate two separate zSeries servers. Dual coupling links were configured for availability, while dedicated coupling links were configured to avoid any possible performance bottlenecks. A CICS/DB2 workload was run with 24 structures spread across the two coupling facilities. A loss of connectivity to one of the CFs resulted in all structures in that coupling facility being rebuilt or recovered using duplexing. In a configuration where all structures were duplexed using System-Managed CF Duplexing (with the exception of the DB2 group buffer pool structure which was duplexed

using user-managed duplexing), the time to recover was 28 seconds. In a configuration where no structures were duplexed, the time to recover was 117 seconds, roughly 4 times slower than the duplexed configuration. This experiment did not include the case where the structures had to be rebuilt from logs. Experience suggests that the time for log recovery would be about 10 times the time to recover the non-duplexed configuration, or 1170 seconds. This recovery time is approximately 40 times slower than the recovery time for the duplexed configuration!

Another significant benefit of System-Managed CF Structure Duplexing is that it helps manageability and usability of a Parallel Sysplex environment. System-Managed CF Structure Duplexing provides a consistent procedure to set up and manage coupling facility structure recovery across multiple structures. A wide range of IBM CF structure exploitation has been updated to take advantage of System-Managed CF Structure Duplexing. These structures include: Batchpipes<sup>®</sup>, CICS shared temporary storage queues, CICS coupling facility data tables, CICS named counters, DB2 system communication area, HSM common recall queue, IMS shared message queues, IMS Fast Path expedited message handler (EMH), IMS Fast Path Virtual Storage Option (VSO), IRLM lock for IMS data sharing, IRLM lock for DB2, JES2 checkpoint, MQSeries<sup>®</sup> shared queues, System Logger, WLM for multisystem enclaves, WLM Intelligent Resource Director, VSAM record level sharing, VTAM generic resources, VTAM multinode persistent structures, z/OS Communications Server sysplex-wide security associations, and z/OS Communications Server SysplexPorts.

Finally, System-Managed CF Structure Duplexing is designed to provide possible configuration simplification and potential financial benefits. For some structure exploiters, prior to System-Managed CF Structure Duplexing, the recovery protocol for rebuilding the CF structure required each of the application instances actively connected to the structure to provide some

portion of the data in order to reconstruct the complete contents of the CF structure that was lost. When the CF structure was lost in conjunction with one or more of the application instances actively connected to the structure, it was either impossible to recover or an elongated recovery process for the application instances to reconstruct the CF structure data was required. Such structures required “failure isolation” or the separation of a CF structure and its actively connected application instances into different servers. Oftentimes, satisfying the failure isolation prerequisites required the presence of one or more standalone CFs in the configuration. This issue of failure isolation—of being able to efficiently repopulate a CF structure if a structure on a CF image together with one of the actively connected application instances should be lost due to an (unlikely) hardware failure—currently limited internal CF usage for certain data sharing structures. Now, System-Managed CF Structure Duplexing resolves these failure isolation issues and enables the use of internal CF-only configurations for all resource sharing and data sharing environments,



thereby possibly eliminating the need for a standalone model coupling facility.

As you would expect, there are costs required to attain such benefits. These costs are highly dependent upon which structures are duplexed, how those structures are accessed by applications executing, and the particular Parallel Sysplex environment itself. Costs of System-Managed CF Structure Duplexing can include increased z/OS CPU utilization, increased coupling facility CPU utilization, and increased coupling facility link utilization.

In order to understand the potential costs, it is helpful to have an understanding of the workings of the System-Managed CF Structure Duplexing technology. Much of the cost of System-Managed CF Structure Duplexing is incurred by the mechanism required to keep two copies of the structure in complete synchronization while at the same time continuing to allow structure access and updates by applications. When z/OS receives a coupling facility update request for a duplexed structure (for example, obtaining a lock in the coupling facility), z/OS will first split the request into two distinct CF commands, one destined for each of the two structure instances. z/OS tries to launch these two commands at as close to the same time as possible so that the CF commands can execute at the two CFs with the maximum amount of parallelism. Once the two commands arrive at their respective CFs, the CFs coordinate the command execution so that the update to the duplexed structure is synchronized between the two structure instances. In effect, the commands must be made to execute in “lockstep” with one another to preserve the logical appearance to the CF exploiter that there is a single consistent copy of the data. To do this, the CFs will exchange signals with one another over a CF-to-CF coupling link, a new configuration requirement for System-Man-

aged CF Structure Duplexing. Once both of the duplexed commands have exchanged all the required signals and completed execution, they each return their individual command responses to the z/OS system that originated them. z/OS inspects each of the responses, validates that the results are consistent between the two CFs, and then merges the results of the operations into a single consolidated request response. It is this consolidated response that is presented back to the CF exploiter, whose mainline request processing is therefore completely unaware of whether the structure being used is operating in simplex or duplex mode. It is important to note that the above algorithm of sending a command to each of the CFs, synchronizing the execution of those commands at the CFs, and receiving and reconciling the results of both commands is only required for those operations that modify a structure. Read requests will be processed by sending one command to one structure and so no additional cost is incurred. For example, a cache structure may have 80% read accesses. Therefore, the System-Managed CF Structure Duplexing costs will apply to only the other 20% of its accesses that update the structure.

So, the costs of System-Managed CF Structure Duplexing vary depending upon which structures are being duplexed and how they are being accessed by the applications using the structures. The costs attributed to those CF operations that update a duplexed structure include increased z/OS CPU utilization, increased coupling facility CPU utilization, and increased coupling facility link usage. It should be noted that CF storage requirements need not increase. Although an additional structure is now required on the second CF, this space should have already been planned to accommodate the CF rebuild situation.

System-Managed CF Structure Duplexing is the next evolutionary step in the Parallel Sysplex architecture. While there is value in System-Managed CF Structure Duplexing, IBM does not recommend its use in all situations, nor should it necessarily be used for all of the structures that support it. A cost/benefit evaluation must first take place. IBM Technical Paper “System-Managed CF Structure Duplexing” (GM13-0103) presents detailed background information and a methodology for performing such a cost/benefit analysis. This paper can be found on the IBM Parallel Sysplex Web page: <http://www.ibm.com/servers/eserver/zseries/pso>. The potential costs have to be weighed against the benefits of improved application availability through faster structure recovery, improved manageability and usability of a Parallel Sysplex environment from a consistent recovery mechanism across a wide range of CF structure exploitation, possible configuration simplification and potential financial benefits through the use of internal CFs instead of standalone CFs.

# The IMS connection: Easing connectivity for z/OS

BY BARBARA KLEIN

IBM's Information Management System (IMS™) is the world's premier e-business database and transaction management system, offering unsurpassed performance in database availability and processing speed for the z/OS environment. IMS, along with its strategic communications product, IMS Connect, have been providing and enhancing high performance paths for access to mission-critical IMS applications, data, and operations in a z/OS environment. These paths take advantage of the latest z/OS Communications and Security Server enhancements to maximize your investment in this platform. Together these products have been improving manageability in workload balancing and autonomic computing across the Parallel Sysplex. And they have been providing integration with ease of application development and connectivity through WebSphere.

## **z/OS Communication Server support:**

z/OS 1.4 Communications Server Internet Protocol Version 6 (IPV6) support is used by IMS Connect 1.2 through service stream enhancements to improve integration, manageability and scalability in TCP/IP access to IMS applications and data. IMS Connect, a separately orderable product running under z/OS, provides enhanced TCP/IP and local option access to IMS applications and data. Early performance data resulted in nearly 6000 transactions per second with a single IMS and a single IMS Connect. Together with its IMS Connector for Java client code and Websphere tooling, Java applications can be developed to run under WebSphere with access to IMS applications through IMS Connect. Recent IMS Connect 1.2 service stream enhancements provide J2EE Connector Architecture two-phase commit support for local environments where Websphere, IMS Connect, and IMS are in the same z/OS. IMS Connect also supports a new GUI workstation single



point of operations control, called the IMS Control Center and integrated with the DB2 Control Center. This Control Center communicates with the IMS V8 Operations Manager through IMS Connect for distributed IMS operations management.

IMS V8, generally available October 25, 2002, provides new Parallel Sysplex terminal management capability, which enables you to resume terminal and user status on a different IMS. In the event of failure, or if you logoff or signoff, you can log onto a different IMS using VTAM generic resources and resume your terminal and user status without having to wait for IMS to restart and without having to logon to the same IMS you did before. Enhanced usability is achieved by using this in conjunction with the z/OS 1.2 Communications Server.

IMS V8 advanced program-to-program communication (APPC) and Open Transaction Manager Access (OTMA) support for synchronous shared queues enhances the distribution of workloads in an IMS Parallel Sysplex environment for Send\_then\_commit (CM1) transactions. Providing access across the MVS Cross Coupling Facility (XCF), and taking advantage of z/OS 1.2 Resource Recovery Services (RRS) support for multi-system cascaded transactions, this IMS support allows users to run synchronous transactions entered from APPC or OTMA on any IMS system in the shared queues group. Synchronous transactions can run on any back-end system from any front-end system with output sent back to the client from the front-end system. Additional

# Take control with SCLM: ISPF option 10

BY MARSHA O'BRIEN

What are you paying for a software configuration management tool? If you are a current user of z/OS, then you already have a full function SCM product and may not even know it! The IBM Software Configuration and Library Manager (SCLM) is a feature of Interactive System Productivity Facility (ISPF) and comes with z/OS. SCLM offers application developers and project administrators the tools they need to develop and manage software projects.

SCLM's library management capabilities allow you to manage your source code and perform versioning and auditing.

SCLM's configuration management capabilities track how all of the pieces of your application fit together and ensure that your inputs and outputs are always in synch. It also manages the controlled movement of your application from one stage in the application life cycle to the next.

SCLM is tightly integrated with ISPF. If you are using ISPF today, then you already know how to use the key functions of SCLM!

If you need an enterprise SCM solution including the management of mixed applications (for example, those containing e-business™ and distributed components as well as zSeries components), you can add the IBM z/OS SCLM Suite.

The SCLM Suite is a set of add-on products that extend SCLM's capabilities beyond the management of traditional zSeries application components. For more information visit our Web site at <http://www.ibm.com/software/ad/sclmsuite>.

APPC enhancements also provide the ability to dynamically add and delete LU 6.2 descriptors, increasing availability by not needing to bring the system down for this; to use another logical unit (LU) as an outbound LU, increasing availability caused by problems if the only outbound LU were disabled; and to use a CPU-time limit for CPI-C (explicit) transactions, increasing availability by helping to prevent looping programs tying up resources.

### **z/OS Security Server support**

#### **z/OS 1.4 RACROUTE**

REQUEST=VERIFY(X) support for the suppression of RACF® abends is used by the IMS V8 RACF-abend-suppression enhancement to prevent a RACF 483-4C abend from causing the abnormal termination of the IMS system. This event can occur during a user signon (or equivalent request) when there is an invalid index block in the RACF data base. When this invalid block is found, a signon request fails with a message and a new reason code, rather than in terminating IMS. This enhancement improves the availability and serviceability of the IMS system.

IMS V8 OTMA enhancements also provide for enhanced security and manageability. Customers are provided enhanced message control over DFS2082 and MQSeries non-persistent messages. The use of the OTMA ACEE hash table for the AUTH call is also being provided.

IMS Connect 1.2 also provides security enhancements, which include PassTicket support, time-out granularity at a message, and relief of limitations on the use of the User Message exits with the limit moving from 15 to 254 user exits. IMS Connect's Connector for Java, providing client connectivity from Websphere to IMS, has been providing security enhancements with container-managed sign-on support for local option environments.

### **Parallel Sysplex manageability**

IMS V8 has been structured to enhance Parallel Sysplex manageability and help address customer requirements for autonomic computing and workload balancing. IMS V8 introduces the Common Service Layer, which is a set of new IMS address spaces that provide the infrastructure needed for systems management tasks in a sysplex. This structure provides a single system image across a sysplex for ease of use. It

enables terminals and users to resume status on another IMS in a sysplex. It supports coordinated online change of resources across all of the IMS systems in a sysplex. It supports enhanced recovery, such as notifying all IMS systems of the loss of a recovery control (RECON) data set in order to reconfigure immediately. And it improves system management of operations and automation by introducing a new IMS sysplex command syntax, a new direction for handling asynchronous command responses, and a single point of control. z/OS 1.2 or higher is recommended for the new resource management and coordinated online change of resources capability.

### **Easing application development and connectivity through WebSphere**

IMS V8 and IMS Connect are providing Java and XML support. Along with WebSphere tooling, they provide WebSphere Java application development and access to IMS applications. They also provide Java application development of new IMS applications, and JDBC access to IMS databases from Java applications and Enterprise Java Beans (EJBs) in z/OS subsystems (e.g. from Websphere, DB2, and CICS).

IMS and IMS Connect, along with other IBM tooling, also provide support for XML, critical for future transparent application integration. IMS documents can be processed in new IMS Java and C++ applications using the XML parser and/or access existing IMS applications using MQSeries. IMS Cobol and PL/I XML application capability for parsing is available using the IBM Enterprise Cobol and PL/I compilers, which allow development of new or modification of existing IMS applications using XML support for COBOL and PL/I. WebSphere tooling also enables existing IMS Cobol and C applications as Web services by connecting the new XML document type—Simple Object Access Protocol (SOAP)—and EJBs to IMS.

### **Summary**

The world of e-business has evolved into advanced e-transaction processing with demands for high levels of integration, manageability and scalability. These demands include continuous up-time for applications, high levels of security, end-to-end business transaction integrity, real-time data currency, and management of unpredictable volumes. In this world

of e-transaction processing, IMS, along with the zSeries and z/OS, stands out to address these demands and integrate new applications with existing data to harness the global opportunity of enterprise-critical e-business.

IMS Connect is part of the overall restructure of IMS for the 21st century, architected as the base for future IMS connectivity. It is provided as a separate entity, while developed alongside of IMS DB and IMS TM, under z/OS, to ensure it can maintain consistency and exploit new functions in an optimized timely manner, ensuring ultra-high performance and availability.

For additional information, visit the IMS Web site at <http://www.ibm.com/ims>.

# PhD not required: Accessing IMS data in z/OS 1.4

BY JACK YUAN AND SANDY SHERRILL

Are you a z/OS customer using IMS? Are you looking for easy access to IMS transactions and commands from other z/OS subsystems? IMS's Open Transaction Manager Access (OTMA) callable interface (CI) might be just what you're looking for!

## What is OTMA CI?

OTMA CI was introduced with IMS Version 6 to help customers more easily write application programs into IMS using OTMA. OTMA is a client/server protocol that uses z/OS XCF to communicate with other z/OS applications. The level of complexity within such a client/server environment often made it difficult to write application programs into IMS. So a focus on ease-of-use was needed. OTMA CI is a high-level interface that simplifies access to IMS transactions and commands. With OTMA CI, you no longer need to be an expert on OTMA's protocol or on XCF macros!

## What can OTMA CI do for me?

OTMA CI can provide a significant amount of function to your z/OS installation. OTMA CI:

- Supports execution of IMS transactions and commands
- Supports unsolicited messages from IMS
- Supports multi-segment messages in both directions
- Supports segment length up to 32K. However, no limit to the size of message.
- Passes RRS context tokens to IMS
- Connects to multiple IMS systems
- Supports calls from both authorized and unauthorized libraries
- Supports preemptive aborts of uncompleted transactions

## What can I do with OTMA CI?

After OTMA CI is installed on z/OS, it can be used for many different applications.

Examples are:

- As a monitor program to send messages to multiple IMS systems
- As a tool application to send IMS commands across a sysplex environment
- As a z/OS adapter to communicate with IMS
- For Web serving from a z/OS Web server
- To test volume without TPNS
- Like APPC to run a transaction on another IMS system without an MSC link

## What APIs does OTMA CI support?

OTMA CI provides several APIs to help you access IMS. These APIs are the heart of OTMA's CI. They are:

**otma\_create:** allocates storage for XCF and IMS communications

**otma\_open:** connects to IMS

**otma\_openx:** also connects to IMS, but in addition you can specify an OTMA exit routine

**otma\_alloc:** creates an independent session to exchange messages

**otma\_send\_receive:** sends a transaction or command to IMS and receive output from IMS

**otma\_send\_receivex:** initiates a message exchange with IMS, but also allows you to pass OTMA user data

**otma\_free:** frees an independent session created by **otma\_alloc**

**otma\_send\_async:** sends a transaction or command to IMS

**otma\_receive\_async:** receives an IMS output message or an unsolicited message

**otma\_close:** frees storage for CI and ends the IMS communications



## How do I use the APIs?

The following are some invocation samples:

- The basic method for sending a transaction to IMS and receiving IMS output:

```
otma_create(...)  
Check the RC  
otma_open(...)  
Wait on ECB  
Check the RC
```

```
otma_alloc (...)  
Check the RC  
otma_send_receive(...)  
Wait on ECB  
Check the RC  
otma_free (...)  
Check the RC
```

```
otma_close (...)  
Check the RC
```

The blue section can be repeated to send additional transactions or commands to IMS.

- Using the OTMA CI to connect to multiple IMS systems:

```
otma_open (IMS1...)  
Wait on ECB  
Check the RC  
otma_open (IMS2...)  
Wait on ECB  
Check the RC
```

```
//Sending transaction  
// and receive output...
```

```
otma_close (IMS1...)  
Check the RC  
otma_close (IMS2...)  
Check the RC
```

- Sending an IMS conversational transaction:

```

otma_open (...)

otma_alloc (...)

//first iteration
otma_send_receive (...)

//second iteration
otma_send_receive (...)

//third iteration
otma_send_receive (...)

otma_free (...)

otma_close (...)

```

The **otma\_open** and **otma\_send\_receive** API need to wait on the ECB.

- Sending a transaction input to IMS and receiving output. Also, receiving unsolicited output from IMS.

```

otma_open (...)

otma_alloc (...)
otma_send_receive (...)
otma_free (...)

//for unsolicited output
otma_receive_async (...)

otma_close (...)

```

- Sending a send-only transaction to IMS

```

otma_open (...)

otma_send_async (...)

otma_close (...)

```

**otma\_send\_async** needs to wait on the ECB.

- Receiving multiple IMS queued output, unsolicited output, or ALT-PCB output:

```

otma_open (...)

Loop begin here:

otma_receive_async (...)
//wait for output

Loop end;

otma_close (...)

```

**otma\_receive\_async** API also needs to wait on the ECB.

### Language wrappers for non-C or non-C++ callers

The CI can be invoked directly from any C or C++ program running in a z/OS environment. The C header file, DFSYCO, needs to be included in the calling program; it declares each API invocation and the parameters used. For caller programs other than C or C++ , you must create a wrapper. For example, the Java wrapper for CI can be created to contain the thin Java/JNI routine, which wraps around the CI. The Java wrapper provides access to IMS for Java applications and servlets, with OTMA CI capabilities and semantics preserved. For callers from assembler and COBOL programs, you would create an assembler macro or COBOL procedure to map each API call to an entry point routine in the OTMA CI object stub, DFSYCRET. See the DFSYCO header file for detailed entry point information.

### Are there security considerations?

For callers from a z/OS authorized library, all of the input parameters, including RACF user ID and group name, are trusted and passed to IMS for transaction or command authorization checking. CI does not perform any security checking. For callers from a z/OS non-authorized library, CI ignores the input RACF user ID and group name. CI invokes the RACF routine to extract the UTOKEN for you and passes it to IMS for transaction and command authorization checking. CI does not perform authentication checking, so password is not a supported parameter for any call. Also, to protect access to the XCF group from any unauthorized callers, the RACF facility resource class IMSXCF.OTMACI can be defined. When the CI connection call, **otma\_open**, is invoked, the resource class, if defined, can be used to reject the request.

### Tips for z/OS 1.4

CI needs to be installed and initialized in the z/OS environment before it can be invoked. If CI is not installed, an F92 abend occurs when **otma\_create** or **otma\_open** is issued. If CI is not properly installed, a DFS3911E error message occurs.

- **otma\_create** and **otma\_open** calls have the same parameters. **otma\_create** can be ignored if **otma\_open** is invoked with input anchor initialized with 0.
- **otma\_open**, **otma\_openx**, **otma\_send\_receive**, **otma\_send\_receivex**, **otma\_send\_async**, and **otma\_receive\_async** each have an ECB parameter. This ECB is posted by the function or by an SRB routine that the function precipitates. The caller must check the ECB and wait for it to be posted before inspecting the return code and output data. Be sure to initialize the ECB with 0 before passing to the CI call.
- Each **otma\_alloc** call creates an independent session for the subsequent **otma\_send\_receive** call. One of the **otma\_alloc** calls can be used to specify the name of IMS transaction or IMS command to be sent to IMS. The maximum length of the transaction name is 8 characters. If no transaction name or command is specified in the **otma\_alloc** call, the command or the transaction name, followed by one or more blanks, needs to be specified in the beginning of the send buffer of the **otma\_send\_receive** call. After each **otma\_send\_receive** call, **otma\_free** is required, except for the IMS conversation transaction. See the invocation sample C for sending a conversation transaction.
- **otma\_alloc** and **otma\_free** are not needed for any async API calls (**otma\_send\_async** and **otma\_receive\_async**)
- CI builds the standard LLZZ part of IMS application data format. You do not need to worry about the LLZZ at all.

- To send a multi-segment message to IMS, the send segment list of the **otma\_send\_receive** call must identify the length of each input segment. The first element in the segment list specifies the number of the segment. The first element is then followed by the length of segment 1, the length of segment 2, and so on.
- When a multi-segment output message is received, an output segment list is provided for the **otma\_send\_receive** call. The first element in the output segment list contains the number of the output segment. The first element is then followed by the length of output segment 1, the length of output segment 2, and so on.
- Sample programs (DFSYCSMP) are shipped with IMS. Some sample programs are also documented in the *IMS/ESA V6 OTMA Guide* (SC26-8743).
- CI can be used to send a protected transaction to IMS by passing a context token to **otma\_send\_receive** call.
- Because some of the CI calls require the calling program to wait, implementing the time-out routine in the calling program is highly recommended to avoid long running transactions in IMS and the internal CI hang.
- To run the CI application efficiently, limit the number of **otma\_open** and **otma\_close** calls in the application. Also, for all **otma\_open** and **otma\_create** calls, try to use the same member name rather than generating a different member name for each call.
- If the size of the output receive buffer specified in the **otma\_send\_receive** call is too short, the actual data returned is limited by the size of the receive buffer. The output can be rejected if a special option, SyncLevel1, is specified in the **otma\_alloc** call. However, if the size of the output receive buffer is too short for the **otma\_receive\_async** call, CI always rejects the output.
- CI could return a bad return code to inform the caller about an abnormal condition. Logging or saving the bad return code for debugging purpose is recommended.
- Here are some additional tips if you have IMS OTMA knowledge:
  1. The **otma\_send\_receive** call sends an OTMA send-then-commit message with synclevel=none to IMS. With the V7 APAR PQ51471, the caller was given an option to set a synclevel=confirm using the **otma\_alloc** call.
  2. When an input RRS context token is given in the **otma\_send\_receive** call, the synclevel is then changed to SYNCPT to support the protected transaction.
  3. The **otma\_send\_async** call sends an OTMA commit-then-send message to IMS.
  4. The **otma\_receive\_async** call receives an OTMA commit-then-send output message from IMS.
  5. CI does not support either the OTMA resync protocol or the OTMA security PROFILE option.

### Summary

Adding IMS OTMA CI to your z/OS 1.4 environment can enhance your access to IMS transaction and commands. What was complex is now simple; you no longer need a diploma in OTMA and XCF! You can find more information about IMS OTMA CI in *IMS/ESA V6 OTMA Guide*.

# This may be your lucky star: ATS STAR plus Automated Tape Allocation Manager for z/OS

BY AMIR RAKHA

Very few of us have extra time these days, especially spare time to wait around for a tape mount. When execution time is critical, every moment counts. Maximizing availability of tape devices and reducing the need for operator intervention help business operations run more smoothly and efficiently. IBM has designed two solutions to improve automatic tape switching and to simplify tape management: ATS STAR and IBM Automated Tape Allocation Manager for z/OS (5697-H62). Depending on your systems environment, IBM has the right choice suited to meet your needs.

## ATS STAR

Shown in Figure 1, ATS STAR is designed for sharing tape devices between systems participating within a single sysplex. Introduced in z/OS 1.2, the ATS STAR function replaces the previous IEFAUTOS coupling facility structure implementation. It has been rewritten to improve robustness and availability. The new design utilizes global resource serialization (GRS) and XCF services to maintain serialization when allocating shared tape devices. This z/OS function maintains tape device information in the ALLOCAS address space.



ATS STAR is available at no additional charge within z/OS. Choose ATS STAR when migrating to z/OS and if your environment is a single sysplex.

## ATAM for z/OS

ATAM provides additional functionality beyond ATS STAR; think of it as “ATS STAR Plus.” Also shown in Figure 1, ATAM is the ideal solution for sharing tape devices in an xplex environment (that is, an environment consisting of multiple sysplexes and/or non-plexes).

ATAM for z/OS reserves tape units as “idle” and immediately makes them available to images contending for tape storage resources—without the need for operator intervention. Unlike older technologies used by some ISVs that rely upon constant updates and monitoring of a shared control file, ATAM operates at the hardware level and exploits changes which occur in tape drive unit control blocks. This “hardware-centric” approach eliminates the single point of failure.

Interaction with ATAM is easy through standard MVS console operator commands or through ATAM’s TSO/ISPF interface. With ATAM, you can define the scope and

level of granularity by system (system ID), range of tape units or by single units to be included or excluded from ATAM’s control. ATAM can be dynamically adjusted without downtime to support your changing workloads and priorities. You can:

- Select continuous operation or request-driven activation
- Alter offline/online polling intervals
- Vary online wait time or modify maximum attempts allowed.

Intuitive navigation through ISPF panels provides data on all recorded tape unit activity, including:

- Tape activity by job and step
- Usage by tape unit
- A record of allocation activity including failed automatic allocations and wait times for tape drive access.

ATAM allows sorting of data by event, date, time, unit, job, step, DD name and data set name. ATAM’s intelligent Virtual Tape Server (VTS) load balancing feature provides unique library identification and enhanced allocation usage in z/OS environments that share two or more DFSMS-managed VTS systems. There are no requirements for sysplex, coupling facility, channel-to-channel (CTC) communications or XCF in order to utilize ATAM in these environments. ATAM operates in any z/OS, OS/390, or MVS environment—including environments that contain all three! If you have been waiting for an IBM alternative to simplify and improve automated tape sharing in an xplex environment, then choose Automated Tape Allocation Manager for z/OS.

## The right choice

Managing storage is central to managing today’s enterprise. IBM has the right tape management solutions to satisfy your needs: ATS STAR and ATAM. By leveraging tape unit availability across multiple images, improving job throughput and minimizing backlogs of job requests, IBM is enabling z/OS enterprises to achieve a better return on investment (ROI) from existing resources. For more information visit:

<http://www.ibm.com/software/ad/atam/>.

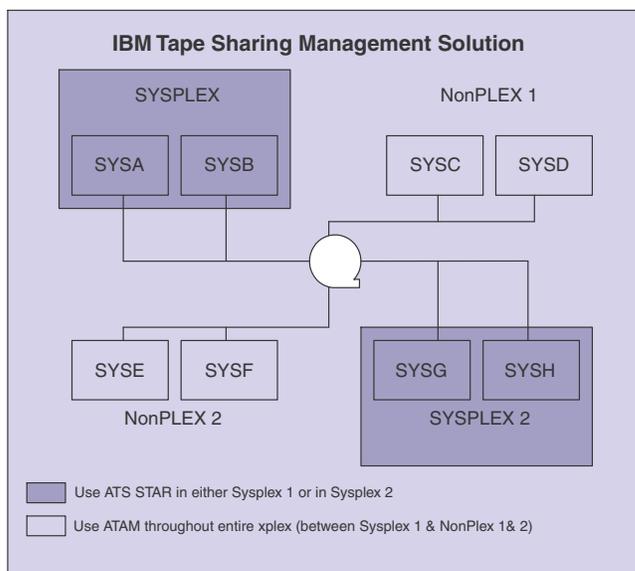


Figure 1



#### continued from back cover

All of the customers that we interviewed expressed their interest in continuing to work with IBM on developing solutions to the skills shortage problem.

#### Solutions to address the skills crisis

What should be done to tackle the skills crisis? We heard suggestions from those who took our online survey and from the IT managers we interviewed. A “Birds of a Feather” session on skills at the August 2002 SHARE conference generated many creative ideas. And IBMers who have heard the results of our skills study have also come up with exciting new proposals.

Here are just two examples:

- To change young people’s perception of the mainframe: Develop an interactive game that gets the word that zSeries servers are not only alive and well, but strategic. The world relies on them to be up and running 24 hours a day, 365 days a year. Mainframes run the latest, hottest technologies; they’re cutting edge.
- To ensure that we have an adequate supply of skilled mainframe professionals: Team businesses with colleges and universities (especially those that offer 2- or 4-year degrees in IT disciplines). The schools would provide mainframe education, and the businesses would provide mentoring and internships. IBM could also partner with these schools, offering internships at IBM sites with a focus on system management training.

In addition to exploring solutions that address perception and training, IBM is also considering solutions to address platform complexity. The theory is that if we can simplify aspects of system management that are currently complex, it will be possible for less experienced people to master them. When we discussed this approach with the IT managers, we received suggestions for areas to tackle, such as problem determination. Basically, many told us, anything that we can do to automate complex tasks would be helpful.

#### The next steps

Our challenge now is to examine the solutions that have been suggested, and implement the ones that promise the greatest return on investment. IBM will host a 2-day workshop with a small set of customers and service providers to delve more deeply into the training issue; namely, how to develop a sufficient supply of IT professionals to meet future demands. Solution proposals and implementation plans will follow.

Other activities in the works include:

- A listserv on mainframe skills named “learnzOS Skills,” which is being set up at the request of those who took part in the SHARE BOF on skills. Listserv subscribers will have an opportunity to review solution proposals and share new ideas.
- A Distance Learning z/OS certification program offered by Marist College, which utilizes the Knowledge Center of the IBM Scholars zSeries Program. The URL for the program is <http://www.academic.marist.edu/S390>.

Bookmark the z/OS Ease of Use Web site <http://www.ibm.com/servers/eserver/zseries/zos/eou/> for more information and announcements about the z/OS Skills Study and related projects. And look for updates in future issues of z/OS Hot Topics.

## Want to help solve the skills shortage?

As a system programmer, you must overcome challenges with creative solutions every day. No wonder there is a skills shortage! Join us at the “Greying of the System Programmer” birds-of-a-feather (BOF) session at SHARE in Dallas (February 2003) for a role playing game in which you help us figure out how to attract new recruits to the mainframe world.

At the BOF session, we will explore the idea of using a video game to recruit new professionals to MVS. What mainframe strengths would such a game help demonstrate? What scenarios could we use to show off the profession? And, what is so darned cool about mainframe system programming anyway? Come to the BOF session and help us to design the ultimate “computer simulation” game.

# Our contributors

**Bob Abrams** is a Senior Software Engineer in the zSeries Software System Design, focusing on availability and problem determination issues. Bob is a member of the zSeries Business Leaders Council (zBLC), zSeries Software Design Council (SDC) and the IBM Problem Determination Council. His background includes 10 years in MVS design and development, 4 years as a JES2/JES3 manager and 10 years in systems management and system design areas related to problem management and RAS.

**Peter Baeuerle** works in design and development of the Workload Manager of z/OS at the IBM Lab in Boeblingen, Germany. He has been with IBM for 21 years and gathered experience in areas like software development for manufacturing systems, team leader for MVS system programming, customer consulting, and Software Architect in knowledge-based systems, artificial intelligence, neural networks and intelligent text search before joining Workload Manager in 1999.

**Jim Becker** is an Advisory Software Engineer at the IBM Lab in Poughkeepsie, New York. He has planned and written user documentation for many IBM software products. Jim still has his "welcome to IBM" letter signed by Tom Watson Jr.

**Joe Bostian** is a Senior Software Engineer, and member of the Domino for z/OS development team. He has been with IBM since 1984, and has worked in z/OS Unix System Services before moving to the Lotus Domino team in 1996. His focus within the Domino team has been to maximize the performance and scale of the server.

**Bernice Casey** is a member of the z/OS Ease of Use team. She has over 20 years experience doing user interface and information design for IBM product.

**Mary Ellen Coleman** is an advisory software engineer in the eServer Information Solutions group, focusing on using multimedia to instruct and to enhance usability. She has been with IBM for sixteen years.

**Lesia Cox** has a background that includes marketing, technical support at IBM's International Technical Support Centers, as well as many years as a technical writer.

**Greg Daynes** has been an IMS and MVS systems programmer for almost 25 years. He is currently an IBM Senior Technical Staff Member responsible for the installation and deployment of z/OS and OS/390 software. This past fall, he led a team that created the new z/OS migration book. Greg regularly gives technical presentations at SHARE, EXPO and IBM Marketing calls on z/OS migrations, installation enhancements and software management.

**John Eells** 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, CBPDO, ESO, and SUF. His current assignments include ServerPac Design.

**Barbara Klein** is the Senior IMS Strategic Plan and Brand Manager. She has been involved with IMS, DB2, and CICS for more than 20 years and has held a variety of positions in Planning, Design, Development, Assurance, Systems Engineering, Industry Marketing, and Management.

**Frank Hildenbrand** is an EIM project manager.

**Peggy LaBelle** is an Advisory Software Engineer with IBM. She designs enhancements to the z/OS Security Server products. Her most recent project was the design of the z/OS implementation of Enterprise Identity Mapping. She is a graduate of Michigan Technological University and Syracuse University.

**Cheryl Loughlin** is an Advisory Software Engineer in the IBM eServer Information Solutions organization in Poughkeepsie, NY. She has been with the company for 15 years and currently enjoys working as a multimedia specialist in designing and developing conceptual information for IBM zSeries products using the latest multimedia technology.

**Ed Mekeel** has been with IBM for 33 years. His background includes 7 years involved with automated operations software. He is currently the content manager for the BCP as a member of the zSeries Software Portfolio Management department in Poughkeepsie.

**Roy Moebus** has over 20 years experience with IBM, and has most recently led the zSeries Parallel Sysplex Technical Support and Product Development Teams. His current assignment is as the zSeries Core Technology Strategy Manager. He has been very focused on systems availability in both of these roles.

**Mary Moore** is responsible for current and future market planning and marketing communications for z/OS and z/OS.e. She began her IBM career as a Systems Engineer, and has held numerous technical marketing positions in the S/390 division.

**Marsha O'Brien** spent 10 years as an SCLM developer on the ISPF development team before moving on to her current position in Marketing.

**Peter Pageroy** has worked for IBM for 24 years in a variety of assignments. He has worked on various hardware, software, and information projects with both his co-authors. He supported the RS/6000® SP for 6 years and most recently has been working on the usability of z/OS.

**Kelly Pushong** is an IBM Certified Senior Project Manager and as a member of the z/OS Product Development Team (PDT) is responsible for Technical Marketing Support for z/OS. Kelly has been with IBM for 14 years and has held a variety of technical and project management positions in the area of Parallel Sysplex design, development, and test.

**David Raften** joined IBM in 1982, working in the Performance Test organization. In 1993, David started looking at Parallel Sysplex and has been doing technical marketing support since its announcement in 1994. Recent focus areas also include availability analysis and z/OS product planning.

**Amir Rakha** is a business development manager with 18 years of experience in the design, development, implementation and marketing of large system software products. Amir is currently focused on new z/OS tools development for the Icing Software Group within IBM. Contact Amir at: arakha@us.ibm.com.

**Bob Rodgers** is a Senior Technical Staff Member on the z/OS design team and is a regular member of the zSeries Software Design Council. He is a frequent and popular speaker at technical conferences in the US and worldwide and occasionally teaches classes on the latest technologies to Software Vendors. He holds a number of software patents, including one for the storage management algorithm of Hiperbatch.

**Sandy Sherrill** is an advisory Information Developer at IBM's Silicon Valley Lab in San Jose, California.

**Don Spangler** is a Human Factors engineer at IBM, Poughkeepsie, working in z/OS User Interface Design and Development. He completed a Ph.D. in Human Factors Psychology in 1999, with backgrounds in Experimental-Cognitive Psychology and Computer Science.

**Shirley Swenson** is a senior software engineer with twenty-two years of experience writing, managing, and planning large systems software and hardware information. She is currently the team leader and planner for the IBM Poughkeepsie Softcopy Center, focusing on customer requirements, delivery, and feedback.

**Susan Van Parys** is a software engineer with over twenty years of experience writing, editing, and planning large systems software information. She is currently a softcopy collection owner for the IBM Poughkeepsie Softcopy Center.

**Jason Williams** is an Information Solutions representative for EIM.

**Jack Yuan** is a senior software engineer at IBM's Silicon Valley Lab in San Jose, California.



To: Readers  
From: Hot Topics Editors  
Subject: **New Hot Topics Newsletter email notification list!!!**

Email us now at [newsletr@us.ibm.com](mailto:newsletr@us.ibm.com) to get your name on the list of people notified when a new issue of Hot Topics is available. Then, every six months, we'll let you know right when the new issue hits the stands. All you have to do is nab a copy through the IBM Publications Center, and you'll have the honor and bragging rights of being the first person on your block to be in-the-know about z/OS's hottest topics!

© International Business Machines Corporation, 1999, 2003

Produced in the United States of America

2-2003

All Rights Reserved

The z/OS 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.

® AIX, AS/400, BatchPipes, BookManager, CICS, DB2, IBM, Infoprint, Language Environment, MQSeries, Multiprise, OS/390, OS/400, Parallel Sysplex, RACF, RS/6000, SystemPac, S/390, VTAM, and WebSphere are registered trademarks of International Business Machines Corporation.

™ DB2 Universal Database, DFSMSdfp, e-business, IMS, IP Printway, iSeries, MVS, Redbooks, z/Architecture, z/OS, z/OS.e, z/VM, and zSeries are trademarks of International Business Machines Corporation in the United States, other countries, or both.

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

Kerberos is a trademark of the Massachusetts Institute of Technology (MIT).

Linux is a registered trademark of Linus Torvalds.

Lotus, Lotus Notes, iNotes and Domino are trademarks of Lotus Development Corporation.

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

Tivoli and Netview are trademarks of International Business Machines Corporation or Tivoli Systems, Inc. in the United States, other countries, or both.

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

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

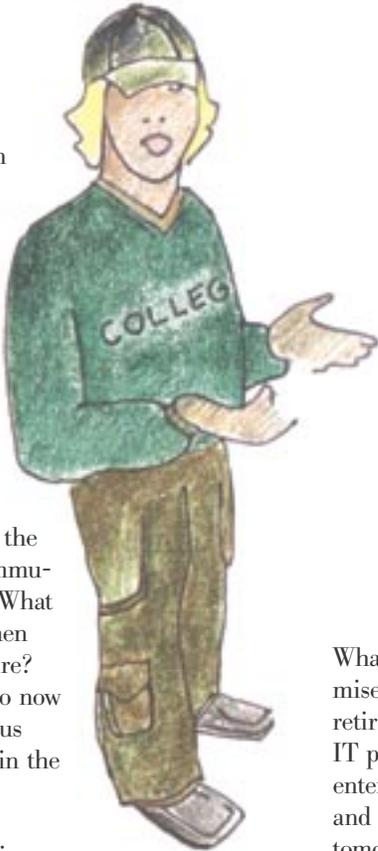
#### **OOPS! ISSUE 7**

In Issue 7 of the Hot Topics Newsletter (August, 2002), the article "TN3270 and Enterprise Extender: Your keys to SNA/IP integration" appeared with an incorrect byline. The correct authors were Sam Reynolds and Mac Devine. We deeply apologize to the authors and our readers for any confusion this may have caused.

# “Dude, where’s the mainframe?”

BY BERNICE CASEY, PETER PAGEREY, AND DON SPANGLER

By now, you’ve probably heard rumors of a “skills crisis” in the mainframe world. Recent polls show that a high percentage of today’s mainframe professionals have 20 or more years of experience. This has made the mainframe community wonder: “What will happen when the experts retire? What can we do now to avoid a serious skills shortage in the future?”



Over the past six months the z/OS Ease of Use team has taken an in-depth look at the skills crisis, its causes and possible solutions. In this article, we’ll share what we’ve learned and the actions that IBM is taking. In future issues, we’ll report to you on our progress.

## What’s behind the skills crisis?

In the first phase of our study, we surveyed 856 MVS, OS/390, and z/OS IT professionals. Most were working in North America, and almost half were system programmers. As Figure 1 shows, many of those we polled had extensive mainframe experience, while relatively few were at the beginning of their mainframe careers.

In addition to asking about years of experience and areas of expertise, we also asked how long they expected to continue working with mainframes. Figure 2 tells the story. We learned that over the next five years, for the population that we surveyed, there is a possible 36% decrease in the z/OS “skill reservoir,” and an even sharper decline by 2012, owing in part to retirements.

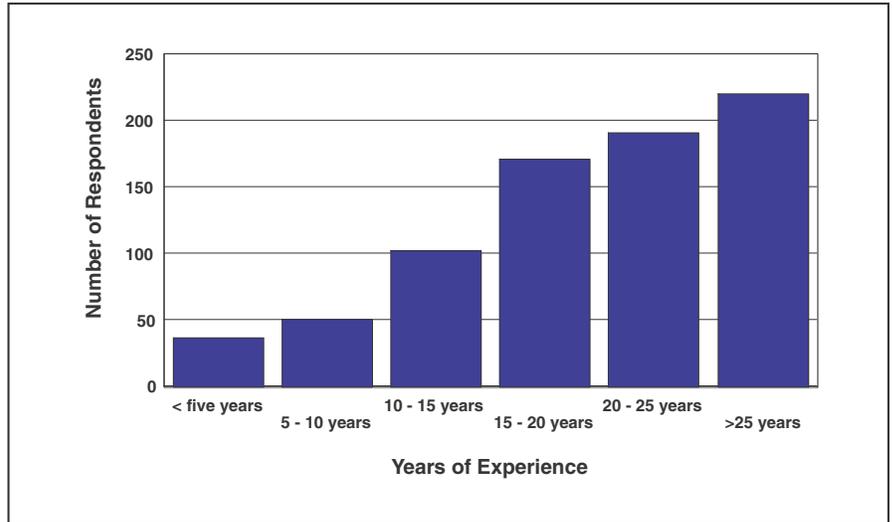


Figure 1 - Experience levels of survey respondents

What creates a sense of crisis, we surmised, is not simply that skilled people are retiring, but that there are so few newer IT professionals either in the pipeline or entering it, building up the skills reservoir, and training to become the z/OS gurus of tomorrow.

With our online survey complete, we launched phase 2 of our study, in which we contacted 13 customers and service providers to get their perspective on the skills crisis. In telephone interviews with IT executives and managers responsible for hiring and training, we heard a number of common themes:

- Customers feel that a skills crisis is looming and that action needs to be taken. Many said that it’s difficult to find skilled mainframe professionals even now.
- They prefer hiring experienced people, instead of hiring and training inexperienced people. Many said that it takes about two years to become a proficient z/OS system programmer, through a combination of hands-on experience and courses. “It’s a complex system with many interrelated components,” a few said. “It takes a while to get the Big Picture.”
- They seem to be looking to outside agencies, such as colleges and universities, to provide a ready supply of skilled IT professionals. However, they said, schools today don’t teach mainframe skills, and even if they did, students might not be attracted to careers working with mainframes because they don’t think of mainframes as exciting or “sexy.”

*continued on page 49*

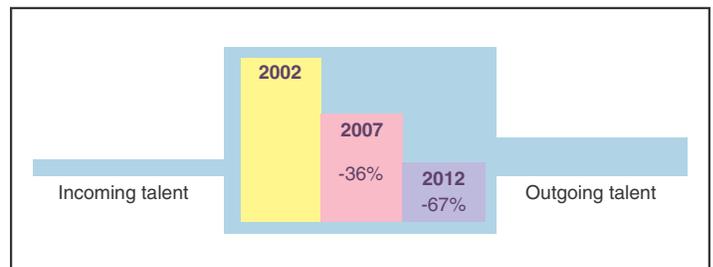


Figure 2 - Skill reservoir



GA22-7501-04