

Extending the Power of IMS Connections for z/OS

Ben Johnson & Jack Yuan

IBM® offers a new tool that allows users to increase the performance, availability, and security of the connections between e-business applications and IMS™ databases. IMS Connect Extensions for z/OS, as its name implies, is a companion tool for IMS Connect for z/OS®, IBM's high-performance TCP/IP server that connects TCP/IP clients and local z/OS applications to IMS.

IMS Connect Extensions provides a new level of user control, enabling users to easily monitor, manage, tune, and secure IMS Connect for z/OS as it opens and closes TCP/IP sockets and routes transactions. Figure 1 shows an IMS Connect for z/OS environment that uses IMS Connect Extensions.




Figure 1: An IMS Connect for z/OS environment that uses IMS Connect Extensions

This article describes the functions and services that IMS Connect Extensions provides to IMS Connect for z/OS users, including:

- Event collection and reporting
- Transaction routing
- Workload balancing
- Transaction pacing
- Enhanced security management
- Online command interface
- Status Monitor

Collecting and reporting event data

IMS Connect Extensions monitors and records four different categories of events that are related to the management of TCP/IP connections, the processing of transactions, and IMS Connect for z/OS trace functions. For each event of a given type, IMS Connect Extensions creates an event record and stores that record in the Active Journal data set. Each event record contains an event number and details about the event.

IMS Connect Extensions records the following four types of events:

- Events related to the status of the IMS Connect for z/OS system
- Events related to the processing of messages
- Events related to traces initiated by the IMS Connect for z/OS RECORDER command
- Events related to an IMS Connect for z/OS trace

Event records related to the status of the IMS Connect for z/OS system identify changes in the status of the IMS Connect for z/OS system, such as the initialization and termination of the z/OS region in which IMS Connect for z/OS runs, the beginning and ending of accept socket actions, changes in the availability of the IMS datastores, and the opening and pre-closing of recorder trace data control blocks (DCBs) on the operating system.

Event records related to the processing of messages identify events that are related to the processing of incoming message requests, such as the beginning and ending of close socket actions, the beginning and ending of z/OS System Authorization Facility (SAF) requests, the sending and receiving of messages to IMS Open Transaction Manager Access (OTMA), the entry and return from user message exits, the reading from and writing to sockets, and the timing out of OTMA.

Both recorder trace event records and IMS Connect for z/OS trace event records identify and capture selected events recorded by traces that are initiated from IMS Connect for z/OS.

The Active Journal data set and reporting events

When the event collection function is active, IMS Connect Extensions continuously records the events at the specified collection level in its Active Journal data set on DASD. You can then create either summarized or detailed reports from the formatted records in the Active Journal data set using the IMS Connect Extensions Print utility.

IMS NEWSLETTER

**Spring 2006
In This Issue**

IMS On Demand
page 4

IMS User Technology
page 6

IMS Connect
page 10

Upcoming Events
page 12

IMS Profiles
page 13

Redbooks
page 14

Hot Websites
page 15

Editorial
page 16

For more in-depth analysis and reporting of the records that are stored in the Active Journal data set, you can use the IBM IMS Performance Analyzer Version 3.3 (Program Number 5655-E15). The IMS Performance Analyzer reports provide performance, response time, and throughput information. For example:

- The Transit Log report provides performance details about every transaction that is processed by IMS Connect for z/OS
- The Transit Analysis report provides a summary of IMS Connect for z/OS transaction performance
- The Trace report provides a detailed list of all IMS Connect for z/OS events for every transaction that is processed by IMS Connect for z/OS
- The Transaction Exception report provides the details of transactions that exceed a certain performance threshold.

When Active Journal data sets become full, or when a specified time or activity interval is reached, IMS Connect Extensions archives the records in the Active Journal data set to the Archive Journal data set on tape or DASD.

Event collection levels

You can specify one of five event collection levels for IMS Connect Extensions. Collection levels determine the events and event categories that IMS Connect Extensions records. Collection level 0 indicates that you want to record the minimum number of events and collection level 4 indicates that you want to record all events. Table 1 presents a summary of the five event collection levels.

Collection Level	Descriptions
Level 0 - Minimum level	Collects start up and shut down events along with some error events.
Level 1 - Accounting level	Collects Return and Exit events, OTMA timeout, and session error events. Level 1 provides accounting information in terms of the number of messages by Transaction, User Exit, and so on.
Level 2 - Transit time reporting level	Collects the minimum number of records to run simple transit time reports.
Level 3 - Comprehensive performance analysis level	Collects all TCP/IP read and write events, which provides for analysis of TCP/IP activity.
Level 4 - maximum Level	Collects all event records.

Table 1. Event collection levels for IMS Connect Extensions

Table 2 represents a sample message flow for an incoming IMS Connect for z/OS transaction message that has a synchronization level of NONE. For each event, the event number and the collection level that is required to record that event are shown.

To record all of the events of the message flow shown in Table 2, you would specify a collection level of 4.

The Transaction Routing function

The Transaction Routing function of IMS Connect Extensions allows you to re-route incoming transactions to different IMS datastores than those specified by the client applications that are submitting the transactions. Transaction Routing provides you with more control and flexibility over which IMS datastores process your incoming transactions without requiring changes to the client applications.

Transaction Routing processing is defined by rules and logical groupings of datastores. This function provides for a redistribution of system resources that is transparent to the client. For example, without modifying the client, you can route all transactions from a single IMS Connect to a preferred, or primary, datastore. If the datastore is taken offline for maintenance, the transactions can be routed to other available datastores also without modifying the client.

IMS Connect Extensions does not re-route all incoming message requests. For example, a Resume TPIPE message cannot be re-routed and, for conversational messages, all subsequent messages of a conversational transaction are routed to the datastore that processed the first message in the conversation.

If your incoming transactions must run on the destination datastore that is specified by the submitting client application, you do not need to use the Transaction Routing function.

The Workload Balancing function

The Workload Balancing function of IMS Connect Extensions allows you to redirect transactions based on the capabilities of individual datastores rather than on the destination of the transactions. For example, in a group of three datastores, if one of the datastores can process transactions at twice the speed and volume of the other datastores, it should receive twice the number of transactions to process as the others in order to balance the workload and use resources optimally.

If Workload Balancing is not active, transactions are routed based on an algorithm that distributes transactions equally among all available datastores.

The Transaction Pacing function

The Transaction Pacing function of IMS Connect Extensions allows IMS Connect Extensions to protect datastores from surges in the number of transactions to be processed by detecting the surge and automatically rejecting the incoming message requests. Transaction Pacing allows IMS Connect Extensions to issue warning messages or to reject transactions if specified incoming transaction threshold values are exceeded.

You can define Transaction Pacing during system definition or dynamically using online commands.



Event number	Event Description	Required collection level
X'3C'	Prepare Read Socket - the beginning of this message request	2
X'49'	Read Socket	3
X'3D'	Message Exit called for READ	2
X'3E'	Message Exit return for READ	1
X'41'	Message sent to IMS	2
X'42'	Message received from IMS	2
X'3D'	Message Exit called for XMIT	2
X'3E'	Message Exit return for XMIT	1
X'4A'	Write Socket	3
X'0C'	Begin Close Socket	4
X'0D'	End Close Socket	3
X'48'	Trigger event CLO - the end of this message request	2

Table 2. Collection levels for events in a sample IMS Connect for z/OS message flow

Enhanced security management

You can improve performance by using IMS Connect Extensions to check security for User IDs and passwords, and for the authority to access IMS Connect for z/OS.

To validate User IDs and passwords, IMS Connect Extensions can cache the access control environment element (ACEE) control block structures for each User ID to reduce the system overhead that is incurred by each security call. You can also specify how long each ACEE structure remains cached for inactive User IDs.

To validate authority to access IMS Connect for z/OS, IMS Connect Extensions verifies that users are authorized to use the IMS Connect for z/OS system by checking a pre-defined security resource class. To improve performance, the security resource class is preloaded in storage if this security function is active.

Online command interface

You can manage IMS Connect Extensions and your IMS Connect for z/OS user exit routines dynamically using the online command interface. By using the online command interface, you can perform the following tasks:

- Modify your user exit routines without interrupting IMS Connect for z/OS execution.
- Refresh SAF class rules and delete cached ACEE control blocks.
- Rebuild the IMS Connect Extensions system definitions.
- Temporarily override IMS Connect Extensions system definitions.
- Rotate the Active Journal data set that is currently in use.

The Status Monitor function

The Status Monitor function of IMS Connect Extensions displays the current activity and processing rates over a selected period of time for either an IMS Connect for z/OS system or for the ports in the IMS Connect for z/OS system.

Using the System View of the Status Monitor, you can view information about the IMS Connect for z/OS system, datastore groups, datastores, and user exit routines. The System Overview panel of the Status Monitor provides a global view of activity on all four of these elements. In the System Detail panel, you can view detailed statistics for each of the elements individually.

Using the Port View of the Status Monitor, you can view the message activity on the ports in the IMS Connect for z/OS system. The Port Overview panel displays activity on all ports. The Port Detail panel displays information for a single port. You can tailor the Status Monitor so that it presents only the information you need.

Summary

IMS Connect Extensions is a powerful tool that allows users to monitor, manage, tune, and secure IMS Connect for z/OS as it processes TCP/IP transaction requests. By using the Event Collection and the Status Monitor functions of IMS Connect Extensions, you can analyze message flows to identify problems and improve performance. By using the Transaction Routing and Workload Balancing functions of IMS Connect Extensions, you can ensure the availability and responsiveness of your datastores for your TCP/IP clients. Finally, by using the security features of IMS Connect Extensions, you can reduce the processing overhead that is associated with security checking, thereby increasing performance.

Ben Johnson
IMS Information Development

Jack Yuan
IMS Development



<IMS>Modern Business World</IMS>

XML is rapidly gaining industry-wide acceptance as the de facto means for data representation for two primary reasons. First, XML is independent of platform, operating system, and programming language; this independence makes XML ideal for smoothing out the differences between disparate systems. Second, XML uses a standard tagging system that greatly eases programmatic parsing and handling of both its contained data and structure. These reasons remain at the center of nearly all industry-wide discussions on how to simplify and navigate the increasingly distributed business environment and have placed XML at the center of the emerging Service Oriented Architecture (SOA).

While the promises of SOA and XML technologies offer tremendous potential for adapting, growing, and managing the increasingly complex modern business world, the obvious obstacle for established businesses today is that little (if any) of their existing data is natively stored or managed as XML. This new problem of how to fit existing non-XML data into the growing XML-oriented business world is receiving a tremendous amount of attention and effort within the database community today. The reason for this effort is that XML is a natively hierarchical markup language. As a result, it is both conceptually difficult and computationally intensive to store, index, query, and retrieve hierarchical XML documents from collections of referentially-joined relational tables.

While the relational community struggles with complex relational transformations or the migration of data to a new breed of unproven hierarchical native XML databases, IMS stands nearly alone in the database community in its ability to store, index, query, and retrieve hierarchical data, and stands completely alone in doing so with the reliability, stability, and performance of over 37 years of experience.

In IMS Version 9, IMS extends its peerless management of hierarchical data by providing the facilities to store and retrieve XML documents from both existing and new IMS databases. By requiring no data migration and allowing concurrent access to IMS data as both native IMS records and as XML documents, IMS XML database support enables users to fully retain existing core business investments in IMS while opening up IMS data to participate in the emerging SOA and XML-oriented modern business world.

How does IMS XML database support work?

IMS XML database support relies entirely on the ability to align the hierarchical structure of an XML document with the hierarchical structure that is defined by an IMS database. Structural alignment is enforced by limiting support to only XML documents that are valid (or conform) to an IMS database's defining XML Schema.

Tooling is provided to assist the user in generating a defining XML Schema that is structurally aligned with a particular IMS database. This aligned XML Schema can be used in two

important ways. First, the XML Schema defines the complete set of XML documents that can be retrieved from or stored into a particular IMS database and is used to filter out all XML documents that do not structurally align with that database. Second, the XML Schema describes the XML layout that includes tag names, element types, and internal hierarchical structures of every XML document already stored in that IMS database.

After an aligned XML Schema has been generated for a desired IMS database, XML documents can be stored into that database by:

1. Verifying that the XML document is valid to the aligned XML Schema.
2. Removing and discarding all XML tags.
3. Converting the character-encoded element and attribute values from original XML types into IMS defined types.
4. Grouping all the one-to-one related values into individual segments.
5. Storing these segments into an IMS database in hierarchical order as an IMS record.

This process, referred to as decomposing or shredding an XML document, discards all information that is not required to reconstruct the original XML document. Tag names, original XML type information, and the hierarchical structure of individual fields within grouped segments are all already described in the defining XML Schema and can therefore be safely discarded.

The remaining information includes the number of occurrences along a particular hierarchical path, the actual data values of elements and attributes in each occurrence, and the overall document ordering of these occurrences can all be represented as a standard IMS record. Removing the XML structural information that is already contained in the defining XML Schema and converting the remaining data into a traditional IMS record allows IMS to then efficiently store, index, and query XML documents with no changes to the underlying database engine.

To retrieve a decomposed XML document from an IMS record, the same process is followed in the reverse order. This time, however, all the shared information that was previously discarded during decomposition is restored using the corresponding descriptions in the database's defining XML Schema.

By not storing any XML specific information into the IMS database, XML documents that have been decomposed and stored as traditional IMS records are indistinguishable from existing native IMS records. As a result, IMS XML database support works equally well with databases that have been created with the intent of managing collections of XML documents as it does with databases that were created with no intention of using XML. The only distinguishing feature between a traditional IMS database and an IMS XML enabled database is the existence of a defining XML Schema. As a result, IMS XML



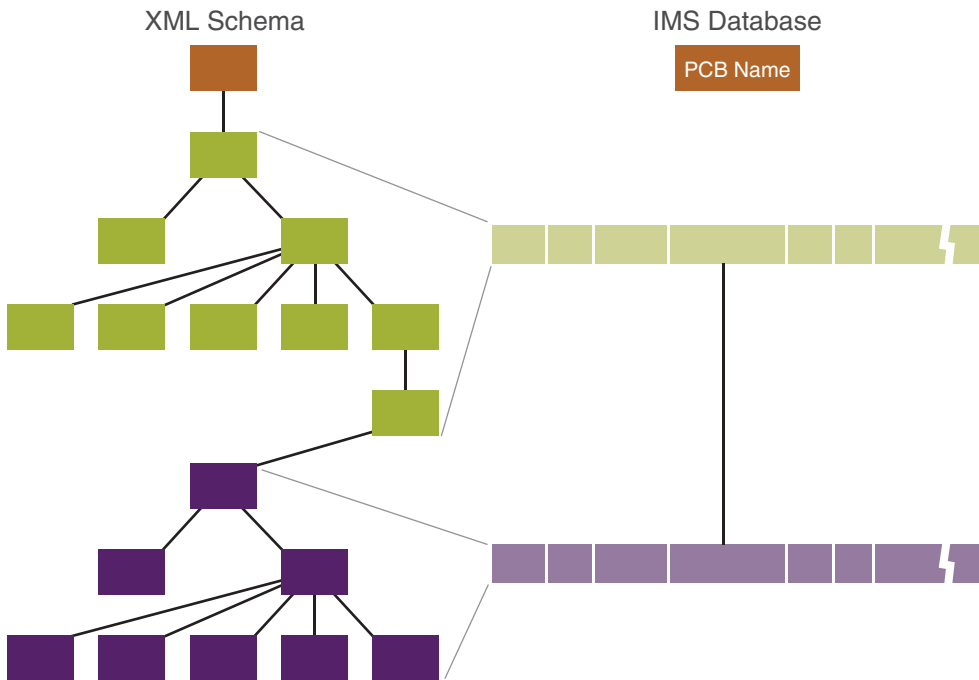


Figure 1. How an XML Schema hierarchy might align with an IMS database

enabled databases are created, monitored, tuned, recovered, replicated, and administered in exactly the same way IMS are today.

How do I use XML database support?

IMS Version 9 includes extensions to both the DLIModel utility and the IMS JDBC resource adaptor that enable IMS XML database support in Java™ applications.

The DLIModel utility was introduced in IMS Version 7 to assist in generating a metadata description of an IMS database. This metadata enables the IMS JDBC driver to interpret and run high-level SQL queries against an IMS database. The utility constructs this metadata by parsing available DBD, PSB, and COBOL copybook descriptions. In IMS Version 9, the DLIModel utility is extended to support the generation of a hierarchically-aligned XML Schema. The XML Schema generation includes options for introducing namespaces, renaming of fields, and choosing the representation of an IMS field to be an attribute, an element, or to be completely left out of the resulting XML document. This schema generation gives the end application flexibility in how XML documents can be constructed.

IMS XML database extensions to the IMS JDBC driver include the introduction of two new SQL User-Defined Functions (UDFs) to handle the storage and retrieval of XML documents. As of IMS Version 9, users can store XML documents into an IMS database by issuing an SQL INSERT call with the storeXML() UDF. Similarly, users may retrieve XML documents by issuing an SQL SELECT call with the retrieveXML(segment) UDF. In both cases, the UDF works on the SQL Character Large Object (CLOB) type. The CLOB type is a standard SQL type introduced

to handle generic large text-based values (in our case, these are XML tagged large text-based document).

The segment argument, passed into the retrieveXML(segment) UDF, indicates the root segment from which an XML document is to be constructed. For example, to retrieve customer information, including all account information, as an XML document for every customer whose account contains more than \$30,000, you would issue the following SQL call:

```
SELECT retrieveXML(Customer)
FROM PCB1.Account
WHERE Account.balance > 30000
```

The FROM and WHERE clauses operate, unchanged from a standard IMS JDBC application, to position the PCB on a particular segment in an IMS record.

The retrieveXML(Customer) UDF indicates to the JDBC engine to move the PCB cursor up the hierarchical path to the Customer segment and begin building the resulting XML document following the description in the defining XML Schema.

What's next: The Virtual XML Garden

The IMS Version 9 UDF extensions to the IMS JDBC resource adaptor enable only the storage and retrieval of full XML documents and take advantage of only a straightforward implementation of the IMS XML database concepts. Our group recently teamed up with IBM Research in investigating how to extend these IMS XML Database concepts to enable the evaluation of full XQuery expressions against an IMS database.

The result of this teamwork is an evolving Java runtime library called the Virtual XML Garden. A freely available, downloadable version of the Virtual XML Garden, complete with IMS samples and installation instructions, is available from the IBM alphaWorks Web site at www.alphaworks.ibm.com.

Conclusion

With the ability to natively fit IMS databases into the emerging SOA and XML oriented business world without sacrificing any of its 37 years of industry proven stability, reliability, and performance, IMS, beginning in IMS Version 9 with IMS XML database support, is enabling IMS customers to truly embrace the modern business world using XML.

Chris Holtz
IMS Developer for e-business



The Evolution of IMS Information Part II

In the Winter 2004 edition of the IMS Newsletter, we introduced the changes we're making to the IMS product information. In describing the overall evolution of the IMS information, we capped that description by introducing the Information Management Software for z/OS Solutions Information Center (publib.boulder.ibm.com/infocenter/dzichelp/index.jsp), which delivers all IMS Version 8 and IMS Version 9 information from the same Web-based tool. We also highlighted the new IMS Roadmap (www.ibm.com/software/data/ims/roadmap/), which consolidates links to a variety of IMS resources, including Redbooks, white papers, presentations, educational materials, and more. When we wrote that article, we didn't realize that we would need to provide a Part 2 to it; however, the very nature of evolution implies continual change, so we are taking this opportunity to share with you more important changes we're making to the IMS library.

Mirror the product

The last time we checked, IMS was approaching its 38th year. Over the course of those 38 years, our intrepid team of IMS developers have used a number of programming languages: Assembler, PL/I, COBOL, C, C++, Java, and more. As a technical user of IMS, you can no doubt see an evolution within that list of languages. As the product has evolved and been rearchitected into smaller, more reusable components, the information that supports it has likewise evolved. And it's not just IMS following this trend, of course. The nature of most businesses is to be innovative, to find ways of doing things better. Our strategy for the IMS information is driven by the concept of rearchitecting into reusable components.

Driven by IBM strategy

The User Technology group at IBM is home to the people who deliver IBM product information. We count technical writers, visual designers, editors, human-factors engineers and programmers among our talented crew. We are historically an innovative and creative group. However, we also look to our corporate User Technology strategists to drive innovation on our behalf. Over the last few years, these strategists have devised new ways for IBM user technologists to author and deliver information. The Information Management Software for z/OS Solutions Information Center is a perfect example. What's next? Something called DITA: Darwin Information Typing Architecture. Before you worry that this sounds like a technical anthropology class, let us explain.

What does Darwin have to do with it?

With all this talk about evolution, it's inevitable that Darwin's name would come up, right? Let's see if we can simplify this. DITA is a standard that supports the writing of topics. Not books, not PDF files, not BookManager files, not help panels

– just topics. Sounds like an oversimplification? The beauty of distilling information down to topics is what you can do with them. We can combine and recombine topics. We can better organize them into cohesive sets of task information and supporting reference information. We can share them with other product groups for reuse in multiple information centers. We can assign metadata values to topics, such as expert or novice, z/OS or AIX® or AS/400®, application programmer or system programmer or database administrator. We can write a topic that is meant specifically for an expert system programmer on the z/OS platform. It's all about granular levels of classification, about writing IBM product information that works for you and your job, and organizing that information in such a way that you can find it more quickly and easily. So what does that mean for the IMS User Technology team? It means that we have to rearchitect the IMS library, all 12,000 pages (give or take a few) into DITA topics. And we're pretty excited about that. Really, we are.







Rearchitect and reorganize

As you might imagine, such an undertaking will require some time. The IMS User Technology team has devised a plan to rearchitect our information into topics over a minimum of two releases of IMS. A preliminary step to that rearchitecture involves reorganizing large parts of the IMS library to eliminate duplication and to group information into more clearly delineated categories. The single-topic books that emerged in the last few releases will be consolidated into higher-level books so that the same information is no longer duplicated in multiple books. Single-topic books contain information about installing, administering, operating, customizing, and troubleshooting several key IMS functions: Base Primitive Environment, Common Queue Server, Database Recovery Control, Common Service Layer, and Open Transaction Manager Access. We'll move information from the respective books for those IMS functions into the "core" library, which is undergoing some changes of its own. An added benefit of the reorganization work is some literary housecleaning. By going through the existing 12,000 pages of information with a fine-tooth comb, we are confident that we can eliminate redundant and obsolete information, which ultimately will make it easier for you to locate the information you need. Ongoing enhancements to the open source Eclipse framework, on which our information center is built, will provide additional muscle for search, support, and service – giving you a true, one-stop shop for all IMS product information.

What does the reorganized library look like?

Figure 1 shows the breakdown of the reorganized IMS product library by primary tasks. Some of the titles (and content) are unchanged. Other titles have changed to better reflect the



 <p>Installation & Related</p>	<p><i>Installation Verification Guide</i> <i>System Definition Guide</i> <i>System Definition Reference</i></p>
 <p>Administration & Related</p>	<p><i>Communications and Connections Guide</i> <i>Database Administration Guide</i> <i>IMSplex Administration Guide</i> <i>Operations and Automation Guide</i> <i>System Administration Guide</i></p>
 <p>Troubleshooting</p>	<p><i>Diagnosis Guide</i> <i>Diagnosis Reference</i></p> <p><i>Messages and Codes Reference, Volume 1: DFS Messages</i> <i>Messages and Codes Reference, Volume 2: Non-DFS Messages</i> <i>Messages and Codes Reference, Volume 3: IMS Abend Codes</i> <i>Messages and Codes Reference, Volume 4: IMS Component Codes</i></p>
 <p>Programming</p>	<p><i>Release Planning Guide</i></p>
 <p>Planning</p>	<p><i>Application Programming Planning Guide</i> <i>Application Programming Guide</i> <i>Application Programming API Reference</i> <i>System Programming API Reference</i></p>
 <p>Other Information</p>	<p><i>Command Reference, Volume 1 and 2</i> <i>Database Utilities Reference</i> <i>Exit Routine Reference</i> <i>System Utilities Reference</i> <i>Master Index and Glossary</i></p>

newly reorganized content; for example, the Customization Guide becomes Exit Routine Reference. Failure Analysis Structure Tables for Dump Analysis becomes Abend Reference. The titles highlighted in bold will be rearchitected to DITA topics for the next release of IMS; the others will follow.

Tell us what you think

In October, we presented IBM's overall information strategy, and IMS User Technology's specific plan, at the IMS Technical Conference in San Jose, California. Early feedback was very positive. We'd like to hear more! If you have questions or com-

ments about the content of this article, please send an e-mail to imspubs@us.ibm.com. If you are interested in partnering with the IMS User Technology team to help us validate the IMS user tasks that you perform, which will contribute to the optimal organization of our information, send an e-mail to that address as well.

Sandy Sherrill
IMS User Technology



IBM Information >>> On Demand

2006

REGISTER TODAY!

**REGISTER TODAY FOR THE IBM INFORMATION MANAGEMENT
FIRST EVER GLOBAL CONFERENCE!**

IBM Information On Demand 2006 global conference is a premier information management event for Business and IT executives, managers, professionals, DBAs, architects, and developers.

- Participate in the PREMIER discussion on the future of Information Management
- Learn how the transformation to information as a Service will help you unlock business value and drive competitive advantage
- Hear how peers are realizing ROI
- Understand the roadmap to long term strategic advantage
- Learn best practices in your industry
- Receive the best in technical education and free certification
- Extensive opportunities for networking with both your peers and industry experts
- Most of all, by combining 5 individual technical conferences into one global event, you gain access and exposure to the full breadth of IBM Information Management technologies and experts

IBM.COM/EVENTS/INFORMATIONONDEMAND



TAKE BACK
CONTROL

The theme for our information On Demand 2006 global event is 'Take Back Control' - take back control of your information, your market, your risk, your bottom line and your future.

REGISTER TODAY AND SAVE!

Register before August 15, 2006 and save \$200!

Early Registration: \$1,695 - After August 15, 2006: \$1,895

OCTOBER 15 - 20, 2006
IBM INFORMATION ON DEMAND 2006
ANAHEIM, CALIFORNIA



Take my IMS Connect asynchronous output . . . please!

In this age of workload balancing and failover protection provided by multiple instances of the same product, IMS Connect has had an Achilles heel: how it handles asynchronous output—but not anymore. Recent APARS to IMS Connect introduce the new IMS OTMA super member solution for IMS Connect. The IMS OTMA super member manages all asynchronous output for multiple instances of IMS Connect and allows any instances of IMS Connect to retrieve the asynchronous output that is generated by any other instances of IMS Connect. Users can now enable true workload balancing and protection from a single point of failure with multiple instances of IMS Connect without risking the temporary loss of asynchronous output.




Figure 1. Two instances of IMS Connect with z/OS Sysplex Distributor

None of these network solutions allow you to specify a specific instance of IMS Connect and, consequently, do not allow users to retrieve asynchronous output messages for a TCP/IP client. Any message retrieval requests or any Resume Tpipe requests might be routed by the network solution to a different IMS Connect than the one to which the asynchronous output messages belong. There was no way to know which IMS Connect would get the message retrieval request for the asynchronous output messages.

See Figure 1 for an example of a configuration using z/OS Sysplex Distributor with two instances of IMS Connect without the IMS OTMA super member.

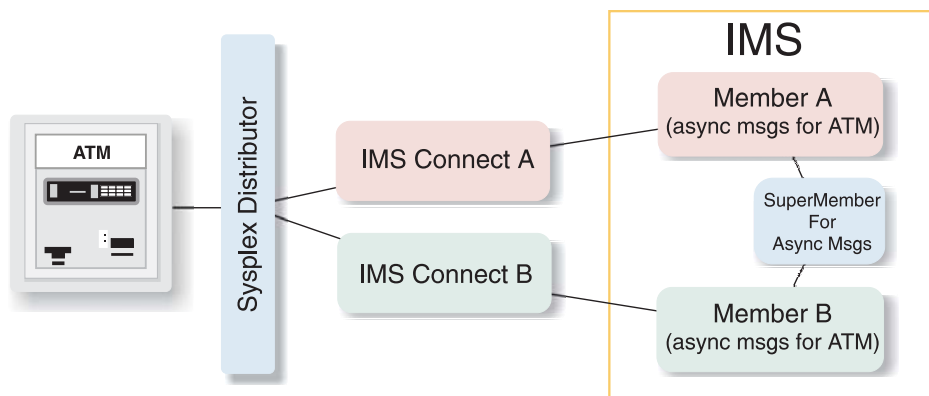


Figure 2. Two instances of IMS Connect with super member function

The IMS Connect asynchronous output problem

Prior to the recent APARS, all asynchronous output messages for IMS Connect had exclusive affinity to the specific instance of IMS Connect for which they were generated. When asynchronous output messages destined for an instance of IMS Connect were queued to an IMS output queue, the output messages could not be retrieved by any other instance of IMS Connect. To retrieve these asynchronous output messages, the TCP/IP application program had to issue a message retrieval request to the same instance of IMS Connect to which it had submitted its original input message. If that instance of IMS Connect was unavailable, the asynchronous output messages could not be retrieved until that instance of IMS Connect became available again.

This problem was further complicated for users of IMS Connect who tried to balance their TCP/IP workload and achieve greater availability by using multiple instances of IMS Connect with a network solution, such as:

- WebSphere® Edge Components Load Balancer
- z/OS DNS/WLM Connection Optimization function
- z/OS Sysplex Distributor

```

HWSC001I  HWS ID=HWSA  RACF=N
HWSC001I  MAXSOC=50  TIMEOUT=5000
HWSC001I  RRS=N  STATUS=REGISTERED
HWSC001I  VERSION=220  IP-ADDRESS=009.030.123.014
HWSC001I  SUPER MEMBER NAME=SM01
HWSC001I  DATASTORE=IMS1  STATUS=ACTIVE
HWSC001I  GROUP=XCFGRP1  MEMBER=HWSA1

HWSC001I  TARGET MEMBER=IMS1
HWSC001I  DEFAULT REROUTE NAME=HWS$DEF
HWSC001I  RACF APPL NAME=APPLID1
HWSC001I  DATASTORE=IMSA  STATUS=ACTIVE

HWSC001I  GROUP=XCFGRP1  MEMBER=HWSAA
HWSC001I  TARGET MEMBER=IMSA
HWSC001I  DEFAULT REROUTE NAME=HWS$DEF
HWSC001I  RACF APPL NAME=APPLID2
    
```

Figure 3. IMS Connect command output for ViewHWS command

Configuring and using the IMS OTMA super member solution

Activate the super member function by specifying a 1- to 4-character super member name in the new SMEMBER parameter of the HWS configuration statement in the IMS Connect configuration member (HWSCFGxx). Super member names must be unique and cannot be the same as existing OTMA member names.

After a super member is activated by multiple instances of IMS Connect, you can display the super member name by issuing the IMS Connect command VIEWHWS. The output for VIEWHWS is shown in Figure 3.

You can issue the IMS commands /DISPLAY OTMA and /DISPLAY TMEMBER TPIPE to display information about the super member. The

output for each these commands is shown in Figure 4

When you issue the OTMA /START TMEMBER TPIPE command, the /STOP TMEMBER TPIPE command, or the /TRACE TMEMBER TPIPE command with a regular member specified, OTMA expands the scope of the command to include any related

```

IEE600I REPLY TO 47 IS;/DIS OTMA
DFS00I  GROUP/MEMBER  XCF-STATUS  USER-STATUS  SECURITY  TIB INPUT  SMEM
DFS00I  XCFGRP1      IMS1
DFS000I  -IMS1      ACTIVE  SERVER  NONE
DFS000I  -IMS1      ACTIVE  ACCEPT  TRAFFIC
DFS000I  -IMS1      ACTIVE  ACCEPT  TRAFFIC  NONE  0  0  SM01
DFS000I  -HWSA1    ACTIVE  ACCEPT  TRAFFIC  NONE  0  0  SM01
DFS000I  -SM01      SUPER  MEMBER  SM01
DFS000I  *05229/144352*  IMS1
    
```

```

0090 * DFS996I*IMS READY* IMS1
0290 R 32 /DIS TMEMBER SM01 TPIPE ALL
0090 DFS00I MEMBER/TPIPE ENQCT DEQCT QCT STATUS SMEM
0090 DFS00I SM01 SM01
0090 DFS00I -TPIPE1 1 0 1
0090 DFS00I *05261/153158* IMS1
0090 *33 DFS996I *IMS READY* IMS1
    
```

Figure 4. Super member information displayed using IMS commands

The IMS OTMA super member solution

To solve the “instance dependency” of IMS Connect asynchronous output, IMS OTMA introduces the super member. A super member is a special OTMA member that can be shared by multiple instances of IMS Connect to manage asynchronous output messages from IMS. All of the asynchronous output messages are stored in the super member instead of in the regular OTMA members that are only available to a single, specific instance of IMS Connect. To retrieve the output messages from the super member, any instance of IMS Connect can issue a message retrieval request. The super member solution is transparent to the end user, who has no idea which instance of IMS Connect handles the message retrieval request.

In the configuration shown in Figure 2, the asynchronous output messages are not queued to the regular member of each instance IMS Connect, but to a new, shared super member. To retrieve the asynchronous message, either instance of IMS Connect can access the super member and any of the messages stored in its message queues.

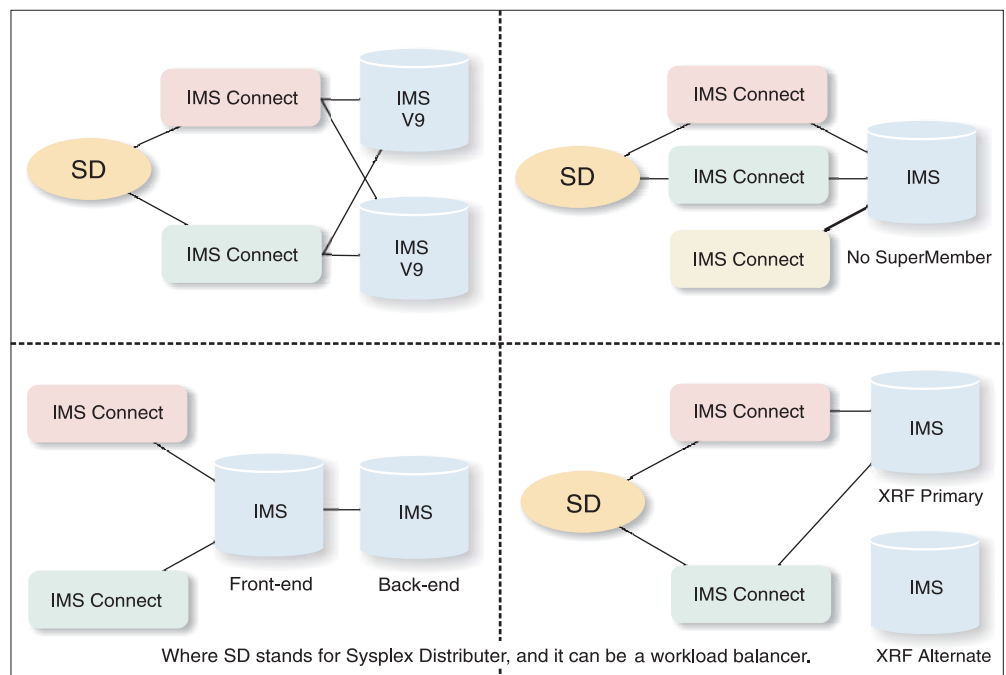


Figure 5. Example configurations running with super member

The IBM Information On Demand Global Conference 2006

October 15-20, 2006
Anaheim Convention Center
Anaheim, California, USA

Register now for the premier information management event for business and IT professionals. IBM's new 5-day event combines the following conferences into one:

IBM DB2® Information Management Technical Conference
IBM Content Management Technical Conference
Business Intelligence Customer Solutions Summit
IMS Technical Conference; and Information Integration Live!

Because so many conferences are combined into one global event, you gain access and exposure to the full breadth of IBM Information Management technologies and experts. More IBM technical and business solutions content will be available in one place than ever before! Select from over 800 sessions: a 2-1/2 day business leadership track with 180 sessions and a 5-day technical track with 650 sessions.

Register before August 15th and save \$200 with our early bird rate!

For more details visit:
www.ibm.com/events/InformationOnDemand

Amiet Goldman
Marketing Manager
IBM Global Services, supporting IBM IT Education Services

The poster features the IBM logo at the top left, followed by 'IBM Information >>> On Demand 2006'. Below this is the text 'REGISTER TODAY!' and 'REGISTER TODAY FOR THE IBM INFORMATION MANAGEMENT FIRST EVER GLOBAL CONFERENCE!'. A list of benefits includes participating in a PREDMET discussion, learning how to transform information into a service, and understanding the roadmap to long-term strategic advantage. The URL 'IBM.COM/EVENTS/INFORMATIONONDEMAND' is provided. A row of five diverse people's faces is shown. At the bottom, it says 'TAKE BACK CONTROL' and 'REGISTER TODAY AND SAVE!' with registration details for August 15, 2006. The dates 'OCTOBER 15 - 20, 2006' and location 'ANAHEIM, CALIFORNIA' are listed, along with the IBM logo.

continued from page 11

super member. You can also issue OTMA commands directly to an existing super member. For example, you can dequeue asynchronous messages from a super member by issuing the command /DEQUEUE TMEMBER supermember_name TPIPE tpipe_name PURGE.

Asynchronous output messages that are created by ALT-PCB processing are stored in the super member directly. If the input messages are submitted through an IMS Connect that supports the super member, the input parameter lists for the DFSYPRX0 and DFSYDRU0 user exits will contain a new flag indicating that the message is from a super member supported client. Also, the OTMA state data prefix pointed to by the input parameter list will contain the super member name, if any.

If multiple IMS systems are associated with a super member, then those IMS systems must also have IMS shared queues implemented. If only one IMS system is connected to by multiple instances of IMS Connect, then shared queues are not required. Figure 5 shows examples of various configurations in which the super member function can be used.

The OTMA super member solution includes the following additional enhancements that benefit users that use shared queues:

- The /DISPLAY OTMA command now shows the OTMA members on back-end IMS systems in a shared queues environment.
- The restriction against cleaning up OTMA tpipes in the shared queues environment has been removed.
- IMS Connect can retrieve the ALT-PCB output messages created by a back-end IMS system from a front-end IMS system.

APARs for the IMS OTMA super member solution

The following IMS Connect and IMS APARs provide the IMS OTMA super member function:

- **IMS Connect V2R2 users:** Apply APAR PK10910 to IMS Connect and APAR PK09944 to IMS V8 and APAR PK09946 to IMS V9.
- **IMS V9 Integrated IMS Connect users:** Apply APARs PK10911 and PK09946 to IMS V9.

Acknowledgements

Special thanks to the IMS people and teams who provided guidance for the super member solution.

Jack Yuan
IMS Development



Robin Hoyt began her career at IBM at the tender age of 17 by working as a part-time Customer Engineer (CE). Her father, who boasts a 35-year career with IBM, opened the door for her employment and inspired her desire to work with mainframe computers. She remembers arriving at customer sites in the 1980s and being asked “So where’s the CE?” Even so,

it didn’t take long for her to win over the doubtful with her skill, knowledge, and drive. She recalls, “It wasn’t unusual to see the CEs, down on their hands and knees, pulling on the hoses and cables under the raised floors.” Within a few years, Robin moved from the hardware area to the software area and eventually landed the level 2 support engineer position that she still holds today. Reflecting on her career, Robin believes that it is the people of IMS that have made all the difference: “IMSerS, both IBM employees and our customers, are just good people!” When Robin’s not putting in extra hours to solve customer problems, you’ll find her working out at the local gym or taking care of her three beautiful children.



Angelique “Angie” Greenhaw began her career at IBM in 2000 after she graduated Summa Cum Laude from Arizona State University with a degree in Computer Information Systems. She chose IBM and IMS even though many of her fellow graduates were, as she put it, “jumping on the bandwagon” by joining the numerous internet

start-up companies in the Silicon Valley. Angie recognized that IBM and the IMS organization provided not only a stable environment, but a great place to begin and grow her career. Her instincts paid off, and she quickly moved from the IMS test group to the IMS Systems 1 development group, specializing in the Online Change function. Angie says: “I’m not sure that I would have had the same kind of opportunity at a smaller company – not to mention the fact that most of those smaller companies are no longer in existence.” Currently, her time is split between her work as an IMS developer and fixing APARs on the IMS change team. In addition, she enjoys being the IMS development representative for SHARE, a technical conference for IBM product users, which is held in various locations biannually. When Angie isn’t solving complex IMS problems or working on future product releases, you’ll find her kickboxing at the local gym or expressing her artistic side through water-color painting and pottery.



The Latest in IMS Redbooks from the ITSO

The following IBM Redbook was updated in October 2005: *IMS Version 9 Implementation Guide: A Technical Overview*, SG24-6398-00.

This Redbook provides the essential information that you need in order to migrate your IMS systems to IMS Version 9 and provides an overview of the new functions and enhancements in IMS Version 9. IMS Version 9 contains over 50 enhancements, which address availability, scalability, capacity, usability, manageability, operations, on demand business, and application development requirements.

The following new Redbook was published December 2005: *Publishing IMS and DB2 Data using WebSphere Information Integrator: Configuration and Monitoring Guide*, SG24-7132-00.

This Redbook describes how to exploit WebSphere Information Integrator's IMS data and DB2 data event publishing capabilities on the z/OS platform. The Redbook uses a customer information facility (CIF) scenario for the financial industry to describe configuration and monitoring considerations.

The following new Redbook was published February 2006: *IMS Connectivity in the On Demand Environment: A Practical Guide to IMS Connectivity*, SG24-6794-00. This book builds upon and replaces the Redbook *IMS e-business Connectivity: A Guide to IMS Connectors*, SG24-6514-00.


This Redbook provides an overview of the IMS Open Transaction Manager Access (OTMA) function and IMS SOAP Gateway. The Redbook also provides extensive information about IMS Connect and its usage, including the IMS Connect Extensions product and how to enhance the IMS Connect operating environment by using IMS Connect Extensions. The book also provides examples of how to develop Java connectors, and how to use Open Database Access (ODBA) with stored procedures and with IMS Remote Database Services.

To join a Redbook residency project, attend a workshop, or download a Redbook, Draft, or Tip, go to the IBM Redbooks Web site at www.ibm.com/redbooks.

Paolo Bruni
ITSO DB2 Information Management Project Leader



Hot Websites



DB2 Information Management Software Information Center for z/OS Solutions
publib.boulder.ibm.com/infocenter/dzichelp/index.jsp

IBM Redbooks
ibm.com/redbooks/

IBM Software for zSeries®
ibm.com/software/zseries/

IMS
ibm.com/ims/

IMS Education
ibm.com/software/data/ims/education.html

alphaWorks: Virtual XML Garden
www.alphaworks.ibm.com/tech/virtualxml

IMS Newsletter
ibm.com/software/data/ims/shelf/quarterly/

IBM DB2 and IMS Tools
www.ibm.com/software/data/db2imstools/

IMS Technical Conference
ibm.com/training/us/conf/ims/

New book: An Introduction to IMS
publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.imsintro.doc.intro/intro.htm

The IMS Newsletter

Spring 2006

Editor-in-chief Glenn N. Galler
Creative Director Robert Lee
Graphic Designer Clark Gussin
Managing Editor Lara Cone
Consulting Editor Moira McFadden Lanyi
Senior Web Consultant Bob Turek

To receive a free subscription to the IMS Newsletter, please complete the form on the IMS Web page at www.ibm.com/software/data/ims/newsletter/

Or contact us in any of the following ways:

Email: imsmtk@us.ibm.com

Fax: (408) 463-4101

US Mail: IMS Newsletter
 IBM Corporation
 555 Bailey Avenue
 DQY/A2
 San Jose, CA 95141-1003



Do you want to receive the latest IMS announcements in your inbox? Go to the IMS website. Click "Receive Recent IMS news" (under **Communities** on the right side of the page). Fill out the form that displays. Click **Submit**. You will be added to an email distribution list that is growing every day!





They Must Be In Here Somewhere

Once again you can find my children's pictures hidden somewhere in this issue. We will send an IBM IMS pin to every person who can correctly find the two pictures of my children in this issue of the IMS Newsletter. If you would like to receive an IBM IMS pin, please e-mail your external postal address and the picture locations of my children to: msmkt@us.ibm.com

While you're at it, tell us what you think of the IMS Newsletter!

Glenn N. Galler
Editor-in-Chief

© International Business Machines Corporation
Printed in the U.S.A. 4/06
All Rights Reserved

AIX, AS/400, DB2, IBM, IMS, OS/390, Redbooks, WebSphere, z/OS, and zSeries are trademarks or registered 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.

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

The information in this document concerning non-IBM products was obtained from the suppliers of those products or from their published announcements. IBM has not tested those products and cannot confirm the accuracy of the performance, compatibility, or any other claims related to non-IBM products.

Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

GC26-9655-18



IMS Newsletter
IBM Corporation
555 Bailey Avenue
DQY/A2
San Jose, CA 95141-1003

PRESORT
FIRST CLASS
US POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1