



## Royal Bank of Canada and IMS Connect - A Winning Combination

By Andrew Kulig, Dominic Tam, Dave Cameron, and Ralph Spadafora  
Royal Bank of Canada

*Chartered in 1869 as the Merchants' Bank of Halifax and later renamed, the Royal Bank of Canada (RBC) has become Canada's largest bank based on assets (over \$362 billion).*

*With over 57,000 employees and 12 million customers last year, the company offers over 4,500 ATMs (automated teller machines) and extensive telephone banking services to its customers. Innovation is second nature to this stalwart financial institution, and in January 1995, RBC became the first Canadian bank to offer information about its services through the World Wide Web. Then, in December 1996, it launched its Internet Banking service. This article is an account of RBC's experience implementing IBM® IMS Connect to achieve the company's online goals.*

RBC has been recognized as the first financial institution in Canada to integrate systems so that our clients can use a single logon ID to access their checking accounts, mortgage services, investment accounts, and lines of credit. RBC also introduced North America's first wireless banking service for handheld devices, offering secure end-to-end encryption of all information. By January 2002, we had over 1.9 million internet banking and trading customers and had introduced the "My View" New Account Aggregation feature within our internet banking service.

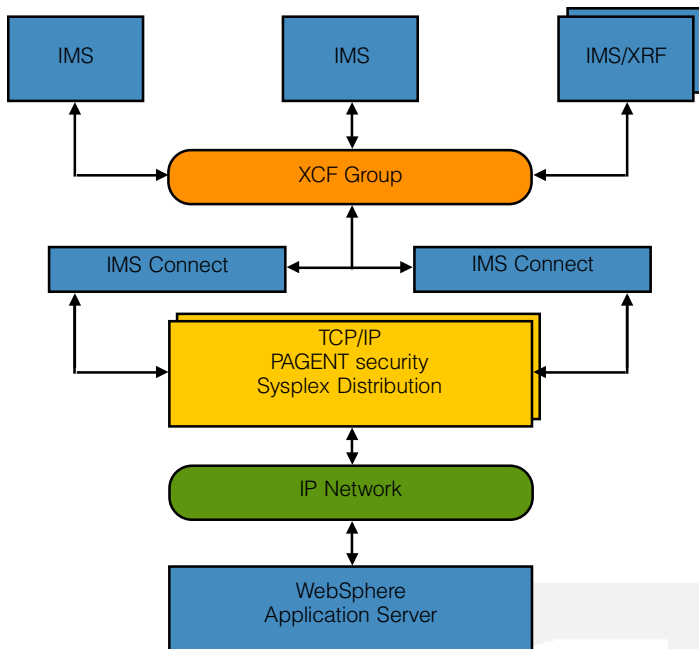
RBC has based its mainframe business on IBM software for many years: IMS®, DB2®, MQSeries®, and CICS®. The challenge was to transform these legacy systems to an online system for e-business. The platform chosen for this new development was IBM WebSphere® Application Server (WAS/390). The access to our existing IMS applications is enabled by IBM IMS Connector for Java, V1 (5655-E51) through IBM IMS Connect, V1 (5655-E51).

IMS Connect provides TCP/IP access to IMS using Open Transaction Manager Access (OTMA). We purchased IMS Connect and were ready to run our first trial of WAS services by the end of 2001. IMS Connect installation procedures use SMP/E, which is a benefit and was very straightforward. IMS Connector for Java is also packaged with IMS Connect and was required for our WAS development. After the initial install, IMS Connect was operational immediately. The code quality was good from the beginning, and we encountered no problems. The sample program proved to be a valuable tool that was easily modified with our installation customizations. Our developers used VisualAge® for Java® (VAJ) to build their application code, utilizing IMS Connector for Java to communicate with IMS Connect. They were able to execute IMS transactions from WAS in a very short period of time. We also have development underway using native TCP/IP socket calls to IMS Connect. This development has also been relatively problem free - the documentation was very clear, and our developers had little difficulty creating the code. From our initial experience, the largest hurdle for native IMS Connect clients was creating the IMS Connect message prefixes.



### In This Issue:

2002 Letter to Customers .....	3
What's New in IMS Tools .....	4
DM Tools Services .....	6
IMS Fast Path Online Tools .....	6
IMS Data Propagator .....	8
User Centered Design in DM Tools .....	10
Education .....	11
Upcoming Events .....	12
Redbooks .....	13
Hints & Tips .....	14
Hot Websites .....	15
Editorial .....	16
Trademarks .....	16



RBC currently has three production centers across Canada, each running multiple OS/390® or z/OS® images in a Sysplex environment. Each center runs three IMS systems, which are symmetrical across geographical centers. One of the IMS systems processes our high volume banking transactions and uses XRF (Extended Recovery Facility) for fail over. The other two IMS systems are clones (of each other) in a shared queues group and are capable of processing the same workload. Our peak transaction rate averages 450 IMS transactions per second for each center.

Currently, WAS/390 runs in one of the three sysplexes- eventually to run in all three sysplexes. WAS applications use IMS Connector for Java to access IMS in any of the three sysplexes. There are two instances of IMS Connect in each sysplex. The IBM Sysplex Distributor is used to balance the connections between them. Each IMS Connect has data store connections to all three IMS systems in the sysplex. This provides multiple redundancy for both the TCP/IP connection and OTMA to IMS. The entire solution runs securely with Policy Agent to limit TCP/IP addresses that can access the IMS Connect TCP/IP ports.

The first (pilot) WAS application delivered new functionality to our Web customers. The initial VAJ development was done using WAS on Windows NT. After WAS for z/OS was ready, the application was ported to the host. Using the current configuration, customers can now open new accounts on the Web using IMS Connect behind the scenes. By driving legacy IMS transactions using IMS Connect, customers can easily access their legacy data. Our internet banking application currently experiences up to 3 million page hits per day. Each page hit averages 2.6 IMS transactions. The internet access has put a new focus on true 24x7 IMS availability.

The use of IMS Connect is expected to grow. Branch office workstations that are currently connected through a Systems Network Architecture (SNA) network will eventually be moved to TCP/IP using IMS Connect. This network currently maintains a rate of approximately 380 transactions per second. In addition, we have an external company that has a requirement to access legacy IMS data from an AS/400™ platform. They have written an application, using native socket calls, that connects to RBC's IMS Connect and ultimately executes IMS transactions to retrieve their data. Several other in-house RBC initiatives are investigating and testing the use of IMS Connect to replace existing SNA IMS sessions with TCP/IP. One of RBC's fraud detection systems is in the process of packaging their software for re-sale. The application runs on AIX™ and has a requirement to send update transactions to an IMS application. Initially, the MQSeries IMS Bridge was used to deliver the transactions to IMS. After experimenting with a native TCP/IP sockets interface to IMS Connect, a decision was made to switch to IMS Connect for reasons of performance and cost (no additional software costs on AIX). Future design options include the use of SSL and J2EE for additional function and flexibility.

Overall, RBC believes that IMS Connect is a high-quality product that works as advertised. It has become part of our strategic direction for not only enabling new access to our existing IMS transactions, but also converting existing SNA sessions.

## DM Tools in 2002 and Beyond

I am pleased to have the opportunity to give you an update on our progress with Data Management tools. Tools have rapidly become an important and strategic part of our Data Management product line. IMS tools have been at the forefront of our development efforts.

### This Business Called Tools

When we started in this “tools game” in September of 2000, we announced 34 tools for IMS and DB2 for z/OS. Our objective was to deliver a comprehensive set of tools as quickly and effectively as possible. In our first delivery we covered all key management functions and delivered GA products within 3 months of our announcement. We also made clear a large, long-term investment in the tools business. We wanted to show you we are serious and committed to this business, that we have the strength to deliver a comprehensive portfolio from the outset, and that we have plans to expand our business quickly.

We added new function to many of these just-released products before the end of the year. We recognize that to meet your requirements we must respond quickly. We are not developing tools on a DB2 or IMS release schedule. We won't wait for new versions of DB2 or IMS to make enhancements to the tools. We're going to give you new and enhanced functions quickly — we're going to add new tools to our portfolio continuously. The tools business is an important business for IBM — because being able to acquire quality tools at affordable prices is important to the success of your business.

Since that original announcement and delivery in 2000 we have had new announcements and deliveries every six months. We are determined to supply best of breed products to support IMS and are investing now and in the future to deliver on that promise.

*Dan Wardman is the Director of IBM Data Management Tools for DB2 and IMS.*

*Located at IBM's Silicon Valley Lab in San Jose, California, Dan is responsible for the development, support, and service of tools for IMS and DB2 on all platforms. Dan and his team have developed and delivered more than 60 tools for IMS and DB2 in an 18-month time frame. Prior to leading the DM Tools group, Dan was the Director of Development for DB2 for OS/390. Previously he was the Director of Software Technical Operations for IBM's Software Group. Dan has held a variety of other responsibilities, primarily in the database area. He was the IMS product manager and the IBM Software Solutions division database strategy manager.*



### What is 'Day One' Availability?

With the IMS tools, we have availability on 'day one' of every new IMS release. So far, we've proved that we can do that, and with IMS V8 we're doing it again. When IMS V8 is made generally available later this year, IMS tools will have been executing in production environments side by side with IMS V8 for a number of months. Specifically:

- The IMS V8 QPP is well under way with 13 customers worldwide. These customers are using a total of 30 IMS tools with 'Run Under Support' for IMS V8.
- Of these 30 IMS tools, the top 4 tools are IMS Performance Analyzer V3 (5655-E15), IMS Batch Terminal Simulator V2 (5655-A14), IMS High Performance Pointer Checker V1 (5655-E09), and IMS Connect V1 (5655-E51).
- Of the 13 WW QPP customers, we picked the top 4 that rely on IMS tools the most and found out that they are all running IMS tools in an IMS production environment. Two of these top 4 WW QPP customers are also running IMS tools in their IMS test environments.

### Our Newest Tools Announcement

On March 6, IBM announced new and enhanced Data Management tools. This is the fourth in a series of tools announcements; it is key to our plan to provide a strategic and complementary set of tools for DB2 and IMS. We offer our tools to simplify administrative tasks, to optimize IMS and DB2 systems, and to control the costs of ownership for S/390® and z/OS systems. The 21 new and enhanced tools for DB2 and IMS that we announced in March will be available in the first half of this year and will support current versions of DB2 and IMS, including V7.

### Directions of IMS Tools

We're committed to expanding and enhancing our function, performance, coverage and support. Improving function and performance are key goals for our developers at Silicon Valley Lab, Boeblingen Lab in Germany, and Yamato Lab in Japan. We have built up our support teams at SVL, and we offer education, services, and field technical support to ensure that you are satisfied with our tools. We have delivered a comprehensive set of tools to cover the highest impact requirements. Our latest announcement extends the range of our existing products and adds 4 new IMS tools. We want you to be able to do everything that you need to do using IMS tools from IBM.

### Why choose IBM's IMS Tools?

This one's easy. First, customers like our tools. Many companies have chosen our tools, and more are choosing them over the competition every day. Our tools business has grown four-fold since we started in this business, and in 2001 alone we added over 1,000 new tools users.

*continued on page 8*

## What's New in IMS Tools

Just to prove that they haven't been resting from their past successes, in March of this year, the Data Management team had its fourth DM Tools announcement in eighteen months. The announcement included the new IBM IMS High Availability Large Database (HALDB) Conversion Aid for z/OS, V1 (5565-I01), enhancements to IBM IMS High Performance Prefix Resolution for z/OS, V2 (5565-I05), and new releases (service rollups) of IBM IMS Index Builder for z/OS, V2.2 (5565-E24), and IBM IMS Queue Control Facility for z/OS, V1.2 (5565-E99).

### IBM IMS High Availability Large Database (HALDB) Conversion Aid for z/OS, V1 (5565-I01)

As applications for e-business generate and require increasing volumes of data, database capacity has become a precious commodity. One way to optimize that capacity is to create a high availability large database (HALDB). Based on a partitioned database structure, HALDB enables you to store more data while making more of it readily available - without changing your applications. However, the task of converting IBM IMS full-function databases to HALDB format requires an experienced IMS DBA.

IBM offers a solution that guides DBAs through the transformation step-by-step: IBM IMS High Availability Large Database (HALDB) Conversion Aid for z/OS (5565-I01). IMS HALDB Conversion Aid provides the utilities and tools you need to analyze, model, and convert existing IMS Version 7 and 8 full-function database structures to a HALDB format. IMS HALDB Conversion Aid enables DBAs to perform the necessary conversions with minimal manual intervention. It requires no HALDB experience and very little IMS expertise. The tool's easy-to-use interactive system productivity facility (ISPF) interface guides the DBA through the entire conversion process. This interface enables DBAs to model database partitions, using either an ISPF dialog or standard setup rules. The tool:

- Performs index pointer healing
- Speeds the conversion process with partition modeling
- Creates database definitions (DBDs)
- Creates Database Recovery Control (DBRC) definitions

In addition, the tool also supports the migration of IMS Version 7 and 8 databases that have been partitioned by the IBM IMS Partition DB product (5697-D85) or the OEM Partitioning product.

### IBM IMS High Performance Prefix Resolution for z/OS, V2 (5565-I05)

This tool performs the same functions as the IMS Prefix Resolution utility (DFSURG10) but with considerably reduced elapsed times. Prefix resolution is required to resolve pointers for databases containing logical relationships and to create or recreate secondary indexes. Unlike V1, this version of the product does not require the BatchPipes™/MVS™ subsystem. Instead, a data transfer service called HPPRPIPE eliminates the intermediate Work File 2 (WF2) data set, which avoids much of the I/O, tape handling, and DASD requirements often associated with prefix resolution.

### IMS Index Builder for z/OS, V2.2 (5565-E24)

This tool can help you simplify index recovery and maintenance, reduce index maintenance time, and eliminate the need to image copy index databases. IMS Index Builder gives you an easy-to-use, one-step procedure to maintain primary and secondary indexes. This release of the product provides a complete service rollup, including SCAN performance fix. No new function is included with this release.

### IBM IMS Queue Control Facility for z/OS, V1.2 (5565-E99)

This tool delivers all the functions of IMS Message Requeuer and a lot more. QCF gives you an interactive ISPF interface, allowing you to Browse, Query, Unload, and Dequeue messages, set controls for the Queue Overflow Protection feature, and more. Message selection capabilities and other features quickly locate checkpoints. QCF supports IMS Standard Queues and Shared Queues. This release of IMS QCF contains no new function but provides a rollup of all service, including a user exit to free queue space, and dynamic QCF table manipulation for IMS QCF.

Need more proof that our commitment to meet customer requirements for affordable, functionally rich tools is being met? Visit our website: [www.ibm.com/software/data/db2imstools/](http://www.ibm.com/software/data/db2imstools/) for more details on all of the Data Management tools. Be sure to read our next IMS newsletter for What's New in Tools later in 2002.

*Margaret M Wilson*  
IMS Tools Market Manager

## Services Offering to Assist Clients Moving from IMS Partition Support Product to HALDB

On April 30, 2002, IBM program services for IMS Partition Support Product (5697-D85) was discontinued. It was replaced by Information Management System Version 7 (5655-B01), announced September 5, 2000, specifically through its High Availability Large Database (HALDB) function.

IBM can provide customized assistance to help you convert to HALDB. IBM can help you:

- Assess the work required
- Plan the conversion
- Obtain necessary education on HALDB or HALDB migration
- Implement the conversion plan

IBM will also provide on-site standby efforts during production cut-overs.

Also available to assist your conversion to HALDB is the newly announced IBM IMS High Availability Database (HALDB) Conversion Aid for z/OS, V1 (5565-I01). This tool provides a step-by-step guide for DBAs migrating IMS full-function databases. IMS HALDB Conversion Aid provides the utilities you need to analyze, model, and convert existing IMS V7 and V8 full-function database structures to a HALDB format.

We have already helped other customers with this migration and would be happy to assist you. For further information contact Donald D. Streicher with the IMS Product Affinity Services group at IBM Corporation, Silicon Valley Laboratory, 555 Bailey Avenue, San Jose, California, phone 1-408-463-2192 or e-mail: [ibmdds@us.ibm.com](mailto:ibmdds@us.ibm.com). To further explain the benefits of moving to HALDB, the tables below highlight some of the differences between IMS PDB and HALDB.

Number and Size of Partition	
PDB	HALDB
Up to 32 partitions for SHISAM and HISAM root-only databases, HDAM and HIDAM, VSAM KSDS, ESDS, or OSAM data sets with a size of 4MB for VSAM and 8MB for OSAM.	Up to 1001 4MB partitions per database for the new database types of PHIDAM, PHDAM and PSINDEX. Each partition can have 10 data set groups. HISAM must be converted into PHIDAM for HALDB.

Installation of the Product	
PDB	HALDB
SMP/E usermods are provided to add support for partitioned databases for current versions of IMS/ESA™. An additional macro 'PART' is added to the DBDGEN to define a partitioned database. This is a separately priced tool.	HALDB is part of IMS V7 and is installed and maintained using the same facilities as the rest of IMS. There is no additional cost for HALDB.

Command Support	
PDB	HALDB
No enhancements to database command support.	HALDB uses the same commands and keywords as non-HALDB databases. The use of a partition name instead of a database name in a command causes it to act on the partition.

Self Healing Pointers	
PDB	HALDB
No support	Reorganization of a partition does not require changes to secondary indexes or logically related databases. These pointers are resolved during normal processing after the partition is back online.

## IMS Fast Path Online Tools

IBM IMS Fast Path Online Tools for z/OS (5655-F78) extends IMS Fast Path Basic Tools (5655-E30) to an online IMS environment, enhancing system availability by allowing you to perform a variety of key functions without requiring IMS to be offline.

The product consists of three separate tools:

1. Online Pointer Checker (OPC)
2. Online Data Extractor (ODE)
3. Online Area Extender (OAE)

OPC runs as a Fast Path utility to collect pointer verification data from a DEDB area while that area is online and being updated by users. The collected data can then be passed to report generating programs, which are part of the DEDB Pointer Checker package.

OPC quickly and efficiently verifies the integrity of all IMS pointer values, free space element chains, VSAM control fields, and space utilization in IMS DEDBs. OPC uses an online DEDB database area as direct input. This means that the database area can remain fully accessible to users while pointer checking is being performed.

OPC performs functions similar to the program FABADA1 of the DEDB Pointer Checker, with the added benefit that the database can remain fully accessible to users. OPC generates files that can be used by the other functions of the DEDB Pointer Checker.

When OPC, DEDB Pointer Checker, and DEDB Tuning Aid are used together, they provide the comprehensive analysis information required to effectively manage the space utilization, performance characteristics, and physical attributes of IMS DEDBs. Together they are an invaluable aid for database administrators and technical support personnel in database management, maintenance, and performance tuning.

Other highlights of OPC include the following:

- Provides a special operating mode to detect pointer integrity errors quickly and efficiently
- Uses an online database as input; in other words, the database can be updated while OPC is running.
- Creates files (while the database is active) that can be used as input by certain functions of DEDB Pointer Checker.
- Produces a comprehensive set of reports, organized by DEDB area, to give multiple views of space utilization, record and segment profiles and placement, and physical I/O requirements.

- Produces the optional CI SNAP Dump, which enables you to analyze the CIs containing errors and to determine the best way to repair them without performing another run to obtain such information.
- Provides optional automated threshold reporting, which causes warning messages to be generated, and a special return code to be set, when conditions within a database do not meet user-specified criteria. This feature can be used to show when reorganization, expansion, or analysis is required, or when database performance is likely to fall below acceptable levels.
- Optionally generates a report identifying the UOWs that exceed a user-specified reorganization criteria. Identifying and reorganizing specific UOWs can significantly reduce the time and resources required for database maintenance.
- Optionally generates a history file containing the key space utilization and performance values for each area. This file lets you conduct statistical and trend analysis using packaged software or user-written programs.

### Online Data Extract

ODE extracts segment data from an online DEDB area and writes the output to a sequential file. The sequential file can be in one of the following formats:

- Standard format
- DBT Unload/Reload format
- DEDB High Performance Unload/Reload format

ODE can be used as an alternative to writing your own application to extract DEDB data. You can also use ODE for:

- Record reporting
- Extracting of records to populate a test database

ODE executes as an online IMS IFP utility, using existing IFP utility services to access DEDB data.

ODE operates on an area level. Each invocation of ODE specifies the area name as a parameter, so ODE will access only this area. However, multiple areas can be processed sequentially in a single run, similar to the DBFUMSCO and DBFUMDL0 utilities.

ODE selects data for specified segments only. The resulting extracted data does not include parents or children of the specified segment, unless they are also specified.

Other highlights of ODE include the following:

- Provides an easy-to-use, flexible tool to extract data from a DEDB area while the area is online.
- Lets you specify which segments are to be extracted by comparing data at an offset into a segment against a specified value.
- Allows more than one comparison per segment type.
- Gives you control over the amount of data you get back by letting you:

Skip a number of segment occurrences  
 Stop selection after a number of segment occurrences  
 Select every Nth occurrence

- Lets you extract SDEP segments for standard DBT format.
- Lets you write a subset of a segment's data to the output file.
- Provides an option to write output data in multiple formats compatible with the two unload/reload formats. However, SDEPs are not supported in unload/reload compatible output data.
- Provides an option to write records in a format suitable for sort.
- Lets you invoke the utility multiple times, for different areas, in a single run.
- Ensures a consistent view of the data by using appropriate locking protocols.
- Uses standard IMS data access techniques, ensuring that ODE functions correctly with all DEDB facilities, including data sharing, MADS, VSO, shared VSO, EQE, and IMS V7 EEQE support.

### Online Area Extender

Occasionally, it becomes necessary to increase the size of a DEDB area. OAE lets you dynamically extend the size of the DEDB area without taking the DEDB area offline. This dynamic capability is very helpful in environments where database availability is crucial. It also eliminates disruption of applications accessing the area.

OAE lets you dynamically extend either one or both portions of the DEDB area:

IOVF portion  
 SDEP portion

Other highlights of OAE include the following:

- Provides an easy-to-use, flexible tool to extend a DEDB area while the area remains online.
- Lets you specify which portion of the DEDB you want to expand: IOVF, SDEP, or both.
- Lets you specify which target area data sets you want to expand.
- Uses standard IMS data access techniques.
- Lets you restart the utility, if needed.

*Alan Smith  
 IMS Tools Development  
 IBM Silicon Valley Laboratory  
 San Jose, CA*

## *Letter to Customers continued from page 3*

Second, we are committed to pricing our tools according to their value. We recognize that the cost of third party tools has been an impediment to the expansion of use of the 390 platform. We have seen many customers switch from 390 to UNIX just because of third party tools costs when they want to upgrade processor capacity. We got into the tools business not only to offer a lower cost alternative, but also to make the 390 platform affordable. You can be assured that we will continue to stay focused on price and value because we see this as a platform affordability issue. Bottom line, we have fair prices.

Third, we are building lots of new tools quickly and therefore have a comprehensive portfolio. We have over 60 products for IMS and DB2. Although many are new, others are now coming out with their second and third releases. We have caught up with the competition in most if not all areas. And, as I mentioned before, we are increasing our development investment and our tools technology to provide superior tools and functions.

Finally, we're in the unique position of building both the tools and the databases. We are working closely with the DB2 and IMS development teams to make sure we have complementary products to the DBMSs. We are making tradeoffs about what function goes into the DBMS and what goes into the tool. We are trying to optimize the overall solution.

### **Tell Us What You Need**

We will be participating in a number of different IMS and user activities around the world this year, discussing and demonstrating IBM's data management tools for OS/390 and z/OS. You'll see us at some of the upcoming regional user group fairs, SHARE conferences, IDUG conferences, and at IBM Technical Conferences. On the web, you'll be able to hear replays of our series of webcasts, which include our answers to questions from customers worldwide. I invite you to visit our website at [www.ibm.com/software/data/db2imstools](http://www.ibm.com/software/data/db2imstools) and see all the new and enhanced tools we've introduced in response to your requests. This is just the beginning of our commitment to you, our valued IMS customers, to continue to bring you the tools you require to maintain successful business operations.

One last item. In case you haven't noticed, we've dedicated this Summer Edition of the IMS Newsletter to our IMS tools. Take a look around, read about our newest IMS tools in the 'What's New in IMS Tools' article, hear from one of our customers, Royal Bank of Canada, read the latest technical details from one of our developers, and take advantage of our education, hot websites, and Redbooks® section.

*Dan Wardman  
Director of IBM Data Management Tools for DB2 and IMS*

## **IMS Data Propagator™**

On October 5, 2001, IBM launched a new version of IMS DataPropagator (IMS DPRDP): IMS DataPropagator for z/OS, Version 3 (5655-E52). New in this version is the capability of propagating IMS database changes to DB2 Universal Database™ (UDB) for OS/390 and z/OS, using MQSeries to transport the captured IMS source database changes. IMS DPRDP Version 3 builds on the established capabilities of earlier versions of DPRDP and leverages the robust function of MQSeries.

IMS DataPropagator Version 3 provides, in particular, the ability to propagate IMS data to DB2 on OS/390 and z/OS asynchronously, but in near-real time, with a propagation latency between the source IMS database update and the target DB2 table update of just a few seconds. With IMS DPRDP's new MQSeries-based feature, point-in-time propagation on a user-determined schedule becomes possible. Additionally, significant consideration has been made to the operation and administration of the new MQSeries-based asynchronous feature, making the management of this feature simple and straightforward. IMS DPRDP Version 3 continues to provide high performance and reliability and is shipped with a revised and simplified packaging and price structure.

The new components of IMS DPRDP Version 3 are:

- A new IMS Change Data Capture (CDC) exit routine
- A DB2 Apply program
- An Event Marker facility

### **IMS DPRDP: Part of IBM's replication solution**

IMS DPRDP is part of IBM's larger replication solution, which has been repeatedly proven in customer environments worldwide for more than ten years. The solution includes DataRefresher™, IMS DPRDP, DB2 DataPropagator, and DB2 DataJoiner™. DataRefresher initially extracts and loads target DB2 tables from IMS database (and other data) sources. IMS DPRDP subsequently propagates changes made to IMS source databases to DB2 (on OS/390 and z/OS). Together with DataRefresher, IMS DPRDP brings IMS operational data into the relational environment.

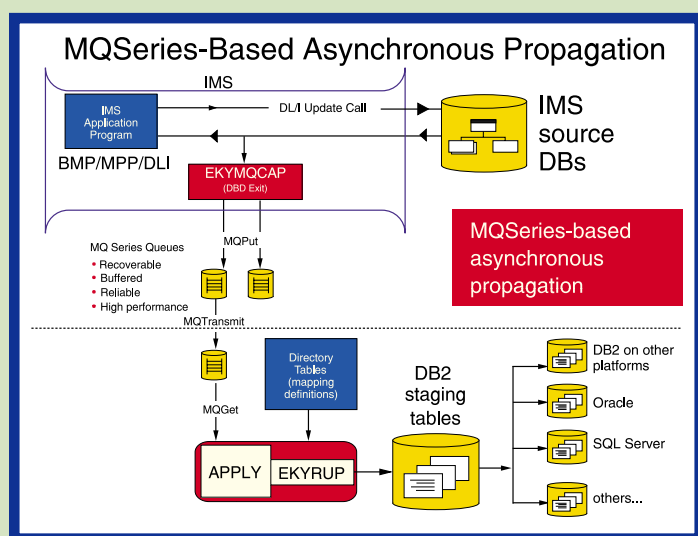
After being propagated to DB2, this data can be exploited for more comprehensive application integration with other relational data sources by using IBM's DB2 DataPropagator and DB2 DataJoiner. Data can be distributed to other relational targets as needed. By defining the DB2 target tables as consistent changed data tables (also known as CCDs or staging tables), full flexibility and integration between each of the different database managers is provided. This can reduce processing costs associated with maintaining a data warehouse and data marts.

With the addition of the MQSeries-based asynchronous replication feature of IMS DPRDP Version 3, continuous low-latency replication from IMS databases to DB2 databases is available. If needed, this low-latency capability of IMS DPRDP



can be used with the low-latency capability of DB2 DataPropagator, and replication of IMS source databases can be staged through DB2 to distributed relational targets. Thus, IMS database source data can be replicated to distributed relational targets continuously and with low latency.

IBM, through its set of replication products, provides a single vendor solution for a customer's replication needs. Data is readily captured and merged from relational and non-relational sources, and can then be replicated to distributed relational targets. Using the new MQSeries-based asynchronous replication feature, data warehouse and data mart environments can be updated in near real time.



### Extending IMS DataPropagator capabilities

With Version 3, IMS DPROG now offers three discrete methods of capturing and propagating IMS data:

- The new MQSeries-based asynchronous method
- The earlier log-based asynchronous method
- The original synchronous method

Users can choose one or a combination of these propagation methods. The choice of propagation method is largely dependent upon user needs regarding the use of any replicated data. Users should also consider the possible impact (although usually limited) a particular method might have on operational data systems.

The new MQSeries-based asynchronous feature provides the most versatile choice when considering which propagation method to use. With this method, changed IMS database data is captured to an MQSeries queue, which can then be made immediately available for replication to DB2 (OS/390 and z/OS). Replication, therefore, is accomplished in near real time. Alternatively, although the changed IMS data is still being captured as changes are made, these captured IMS changes need not be applied to their DB2 targets until some later point

in time chosen by the user, for example, at the end of a business day. Additional flexibility is also possible by utilizing the facilities of MQSeries. The queue that the captured IMS changed data is written to might be transmitted to a remote system MQSeries queue, where the DB2 propagation target tables might ultimately reside and the data changes applied at that remote system.

If your business needs require that DB2 target data be completely current and up-to-date with IMS source data in true real time, and the operational requirements of these methods are acceptable, then the synchronous feature can be used.

The asynchronous-log based feature can be used if propagation latency requirements can tolerate the need to wait for IMS online logs to be archived before initiating the replication process. Realistically, using the log-based method will result in a propagation latency of at least thirty minutes to one hour.

Whichever method is chosen, each assures data integrity and assures that only committed IMS data will ever be propagated through to the target system, with the delivery of captured data assured.

### For more information

Please contact your IBM marketing representative or IBM Business Partner, or call 1-800 IBM CALL within the U.S. Also, visit our Web site at [ibm.com/software/data/db2imstools](http://ibm.com/software/data/db2imstools).

*Anthony Bull*  
*IMS Replication Development*

## User-Centered Design in Data Management Tools

Competition and the constant desire to improve our products drive IBM. This is evident in our tireless efforts to provide the best solutions available to our customers. IBM's Data Management Tools team is no exception. In a market as competitive as Data Management Tools, it has never been more important to know who our customers are, the challenges they face, and the products they require to meet their data management needs. To build this knowledge, the Tools team follows a process called User-Centered Design (UCD), where feedback from customers drives the design of future versions of our products. As designers and developers for the tools, working to provide the most intuitive tools possible, we are constantly looking for opportunities to understand and hear from our users. Such an opportunity recently presented itself, and an example of the Data Management Tools team's commitment to UCD emerged.

### Understanding Our Users - A Case Study

In May a member of our team had an opportunity to visit one of our tools customers at their site. Often, customer visits provide the richest feedback on how our products are used and how the technology can be improved. This particular customer visit dealt with usability issues for the IMS Database Control Suite, an IMS tool that allows the user to manage and maintain their IMS databases from one central location.

### Surveys

To better prepare our team to address customer issues with this tool, a brief survey was sent to their DBA team a week before the scheduled visit. This survey was composed of three sections. The first section contained questions about the customer's environment. The DBA team was asked to specify the number of years they had been using IMS and the latest version they were using. Also included were questions addressing their experience with the Control Suite; how often they used it, what version they were using, and how many databases they maintained with it. Background information such as this is crucial to understanding the size of the customer's shop and their experience with and expectations for the product.

The second section of the survey dealt specifically with the functionality of the Control Suite. Each major function was listed, and the users were asked to indicate which functions they performed, along with their collective satisfaction when performing each function. This section ended with the users ranking the functions they performed by importance to their company. The responses here helped guide the visit and dictated to our team where we needed to concentrate our efforts for this particular customer. Functionality that was considered vital to the business or received poor satisfaction ratings was the first to be addressed when our team member arrived on site.


The final section asked the customer to specify other IBM Data Management tools they used as well as any competitor tools. With this information in the hands of designers, developers, and managers, the Control Suite team had a complete picture of the customer's tools environment. When the day of the visit arrived, everyone that was representing IBM was focused on the needs, expectations, and goals of this particular customer.

### Observation and Requirements Gathering

The first half of the visit centered around understanding the issues and concerns the customer had with the Control Suite version they were using. Four participants were present including two DBAs, a System Programmer, and a member of their Technical Support staff. During the first hour we discussed in detail the usability issues they were currently experiencing with the product.

With the Control Suite running in front of us, the participants generated the feedback we sought. One participant walked through the functions of the Control Suite while the others raised concerns and made suggestions for improvement. Issues were raised concerning everything from installation and set up to report printing and downloading documentation from the Data Management Tools Web site. IBM was receiving input that would be used to improve the total user experience. Every aspect of user interaction with our products is important when following a UCD approach. As a result of this activity, our team had a list of ways to enhance the Control Suite. Certain suggestions such as improved documentation and a more automated set up process were consistent with IBM's long standing goals for the tools; other suggestions, such as a logging function to help with audit responsibilities, were noted for the first time. All input was documented and will serve to direct our improvements.

Our team member was taken on a tour of one of the four buildings on the customer's site. During this time useful information was being gathered. It became clear the workers placed a premium on desk space. Perhaps



a quick reference card for the Control Suite would more easily fit into their environment than another large user's manual.

### Evaluating New Designs

A teleconference had been scheduled between the customer's team and the rest of the developers and managers at IBM's Silicon Valley Laboratory. During this time our developers demonstrated the version of the Control Suite currently under development.

Discussion between the participants and IBM's developers was clearly welcomed by both sides. Time was available during the presentation for questions and answers. It became evident that some of the customers' concerns with the Control Suite were already being addressed by development efforts. Issues raised during this time were being delivered to development and management directly from the customer. This type of feedback is at the heart of user-centered design. These participants provided IBM with valuable feedback early in the current development phase, ensuring that certain enhancements will be available in the next release of the Control Suite.

The day ended with a valuable IBM customer having its concerns heard and IBM obtaining the information required to improve our tools and provide the solutions our customers demand. In a market where innovation is synonymous with success, the benefits of listening to our customers is clear.

### You Can Be Involved

In all of our endeavors we need users on our team to ensure that we are designing the most useful and easy-to-use solutions possible. If you would like to participate in UCD activities for any of IBM's Data Management Tools, you can sign-up on the Web at [www.ibm.com/ucd](http://www.ibm.com/ucd).

*Joseph Sacco  
DM Tools User-Centered Design  
Kelly Spain  
DM Tools User-Centered Design*

## Data Management Education

An easy way to learn about the functions and features of IBM Data Management Tools is to go into the Data Management Tools web site at [www.ibm.com/software/data/db2imstools/](http://www.ibm.com/software/data/db2imstools/). From there, click on the recently added option [Education](#).

### Data Management Tools Education Web Page

The Education page is continuously updated. The page contains:

- Function and feature presentations of the different DM Tools
- DM Tools demonstrations
- Links to other learning-related pages where you can find DM Tools education; currently there are links to IBM Redbooks and IBM Learning Services.

### Function and Features Presentations and Demos

The DM Tools presentations and demos are listed in alphabetical order. They are easy to download and free of charge. They provide an overview of each tool and its functions and value.

If you want to learn more about a specific tool, but aren't quite ready for a formal class, these presentations and demos are the right choice.

The IMS Tools presentations that are available as of June 2002 are:

- IBM IMS Online Recovery Service for OS/390, V1 (5565-E50)
- IBM IMS Library Management Utilities for OS/390, V1 (5565-E04)

Soon to come are presentations for IBM IMS Connect, IBM IMS High Availability Large Database (HALDB) Conversion Aid for z/OS, V1 (5565-I01), IBM IMS High Performance Sysgen Tools for z/OS, V1 (5565-F43), and IBM IMS Performance Analyzer for OS/390, V3 (5565-E15).

### IBM Learning Services

IBM Learning Services continues to provide classroom and web-based classes. IBM Learning Services that are focused on Data Management education are listed, with links and references where applicable.

If you are looking for customer-tailored education or have specific education requirements, please contact your IBM Representative to get information about our chargeable DM Tools Education Services.

*Remember to keep a bookmark on our Education page since it is being continuously updated!*

*Brittinger Ottersgard  
DM Tools Customer Relations Manager*

## 2002 IMS Technical Conference

September 30 - October 3, 2002

St. Louis, Missouri

Adam's Mark Hotel

Have you read the latest business news? The economy is gradually heating up. Now is the time to strengthen the foundation for future growth, and there's no better event than the IMS Technical Conference to learn the latest insights about IMS. The Conference is "simply the best way to stay current in a fast-moving technological environment," an actual quote from a conference attendee last year.

### Technical Details

This year's event has three major focus items:

IMS Version 7 Product Features and Migration Information

IMS Tools Galore

IMS Version 8 Features and Early Experiences

Migrating to IMS Version 7 allows you to take advantage of:

- HALDB support - learn the details for implementation, operational features, and the utilities involved in migrating your full-function databases to HALDB.
- Shared Queues Enhancements - find out what makes them different from message queuing and hear about all the benefits of using shared queues.
- IMS Java support - learn the tools for not only developing new Java applications for IMS but also using JDBC to access IMS DB data from CICS Java applications, DB2 Java Stored Procedures and WebSphere EJBs.

IMS provides a wide selection of management, application development, and access tools. These tools provide your IMS system with enhanced:

- Connectivity - get all the details on how to provide your IMS systems high performance connections to the web and allow new business processing capability with WebSphere, Java, and XML.
- Recovery - find out how to increase performance in recovery situations, increase flexibility in recovery options, and increase simplicity in the recovery process.
- Business Intelligence - learn how to provide application and data consistency in updated IMS DB and DB2 data with asynchronous near real-time propagation.

IMS Version 8 is the new kid on the block. Hear from the product experts — many of them product developers — how this new release is focused on enablement for e-business as well as growth, availability, and systems management needed for an effective e-business environment. As with all other IMS releases, IMS Version 8 also focuses on continued performance capacity, open access, and supporting tools for the e-business environment.

Need more information to convince your manager that you need to attend the conference? Check the web site for detailed session titles, speakers, and detailed abstracts to justify your enrollment.

Coming in August, you'll find the detailed agenda posted on the web, along with a spreadsheet — ready to download — so you can determine your optimal schedule at the Conference.

### Early Bird Electives

Add an extra day's value to the conference by attending up to 4 elective sessions on Sunday! For no additional cost, supersize your conference experience by arriving on Saturday and attending elective sessions on Sunday!

Elective sessions held on Sunday include:

E31 Migrating to IMS Version 7 (double session)

E33 IMS Version 7 Installation and Generation (double session)

E37 IMS/OTMA Security Considerations

E38 IMS Connect Security Considerations

E13 IMS Java Application Development

E14 JDBC Access to IMS DB from DB2, CICS and WebSphere

### Hands-on Sessions

At the 2001 IMS Technical Conference, we debuted a hands-on workshop session. If you missed last year, come to St. Louis to find out what you missed, and choose from the popular session offered last year or the new hands-on sessions premiering this year!

Hands-on sessions include:

H01 IMS Connect Workshop: Creating Web Services Connecting to IMS (double session)

H03 Hands-on Workshop: Creating XML and Web Services Using IMS Source Files

H05 Hands-on: IMS HALDB Conversion Aid

H06 Hands-on: IMS Database Control Suite

H07 Hands-on: Tuning IMS with the IMS Performance Analyzer (double session)

## Product EXPO

Count on the Product EXPO for a firsthand look at the latest IMS technologies and solutions. Vendors provide demonstrations and answers to your questions about applications that can benefit your enterprise. Bring a list of your specific challenges and make side-by-side comparisons that can lead to solid purchasing decisions.

St. Louis, home of the graceful Gateway Arch, which towers above the mighty Mississippi River, has something for everyone — before and after conference hours, of course! There are sporting events to satisfy all fans. If you enjoy other types of exercise, you can find a wide variety of night life in the clubs including swing and salsa as well as the more traditional blues, jazz, and rock-n-roll music for dance floor enthusiasts. There are outlets for those who'd like to try and tempt "lady luck" on some exciting themed Riverboat casinos. If you plan to arrive early or stay late, you might want to make time to check out other daytime activities including tours of the famed Anheuser-Busch Brewery including the Budweiser Clydesdale stables, the world-class St. Louis Zoo, the Six Flags St. Louis amusement park, riverboat cruises, golfing on outstanding courses, State Park recreation, and so much more.

The Conference is "simply the best way to stay current in a fast-moving technological environment."

For more details and to enroll, visit the conference Web site at: [ibm.com/training/conf/us/ims/](http://ibm.com/training/conf/us/ims/)  
Your registration fee of \$1,695 is already paid if you have a US IBM Education Card.

*Amiet Goldman*  
*Senior Direct Marketing Strategist*

## IBM Redbooks

From the DM Tools website, you will see a link to our well-known IBM Redbooks, written by users of the products and very often containing usage scenarios and user experience with the tools. The list of DM Tools Redbooks is continuously updated. The IMS Tools Redbooks that are available as of June 2002 are:

*A DBA's View of IMS Online Recovery Service,*  
*SG24-6112-00,*

Redbook, published March 7, 2001

*IMS/ESA Database Tools Vol I: Database Manager Tools,*  
*SG24-5166-00*

Redbook, published November 2, 1999

*IMS/ESA Database Tools Volume II: System Extension and Other Tools, SG24-5242-00*

Redbook, published November 24, 1999

*IMS e-business Connectors: A Guide to IMS Connectivity,*  
*SG24-6514-00*

Redbook, published January 29, 2002

*IMS Version 7 Performance Monitoring and Tuning Update,*  
*SG24-6404-00*

Redbook, published January 13, 2002

*Ensuring IMS Database Integrity using IMS Tools,*  
*SG24-6533-00*

Redbook, published March 14, 2002

Also in plan for availability in 2002 are IMS Redbooks for DPROP Fast Path and Handling Data.



*Brittinger Ottersgard*  
*DM Tools Customer Relations Manager*

## Shared Virtual Storage Option Buffer Utilization Hints and Tips

Many customers are migrating to Shared Virtual Storage Option (Shared VSO) to replace main storage databases (MSDBs). In order for Shared VSO to provide a similar level of performance as MSDB, you need to minimize I/O, both from the coupling facility (CF) structure and from DASD. This article provides some tips to help you do just that.

To ensure that there is minimal DASD I/O, your CF structure must be large enough to hold all the root addressable and IOVF CIs from the Area. If the CF structure is too small, CIs are eliminated to reclaim space. A "least recently used" algorithm determines which data is removed from the CF.

Minimizing CF I/O is more complicated than it appears. Assuming the LKASID option is selected, the pool must be large enough to satisfy all concurrent requests for CIs from the Area. If LKASID is not specified, all requests read data from the CF structure. Dependent regions cannot share buffers in the pool. Therefore, determining the required number of buffers is not as simple as totalling the number of RAA and IOVF CIs in the Area. If most requests to the Area use a non-update PROCOPT (G, GO), many dependent regions can concurrently read the same CI. Each request requires its own buffer from the pool. For example, if the Area contains only one segment, and 100 dependent regions read (share) the segment concurrently, the required pool size is 100 buffers. Alternately, if 100 dependent regions access the segment with an update PROCOPT, only one dependent region owns the lock for the CI at a time, which means that only one buffer is required. The other 99 dependent regions will wait for the lock. When the lock is granted, the dependent regions reuse the same buffer serially without re-reading it. Note that segment-level locking is not supported for shared VSO.

Many MSDBs have a reference pattern that is heavily biased towards read rather than update access. To accommodate such applications, you can approximate the required private pool size. Multiply the number of dependent regions making non-update calls (calls resulting in SHR locks) by the number of CIs from the Area that each dependent region will access in each syncpoint interval. For example, if the application normally runs with 50 MPP regions, and each region will access three segments that have been carefully randomized to three different CIs, then 150 buffers are required - even though only three CIs actually contain data of interest. Use the DEDB keyword in member DFSVSMxx to define private buffer pools, including their primary allocation size, their secondary allocation size, and their maximum size. The primary size is allocated and page-fixed at IMS initialization. Secondary allocations are obtained up to the maximum size, as required. If the pool cannot be expanded, and an additional buffer is needed, the dependent region must wait for a buffer to be freed by another dependent region. The primary allocation size should be large enough to accommodate the normal workload. The pool will expand if needed, and secondary allocations are freed at checkpoint time if no longer required. By determining the proper primary allocation size,

you can prevent expansion and contraction overhead.

Another issue with shared VSO that can be confusing is "lookaside" within a dependent region. Normally, when an application reads a DEDB CI, a SHR lock is obtained and held until either the syncpoint is reached or the PCB position moves off the CI and buffer steal is invoked (NBA limit is reached). This means that if the application reads the same CI repeatedly, only the first read actually obtains the lock, the buffer, and data. Subsequent reads can find the buffer by using the lock, but they do not need to access the buffer pool, the CF, or DASD. However, all FLD calls, and those GET calls that specify VIEW=MSDB in the PCB, will behave differently than normal GET calls. In this case, a SHR lock is obtained, the buffer and data are obtained, and the SHR lock is released at the end of the DL/I call. The buffer is put on the pending-discard chain for the dependent region, which means that it cannot be used to satisfy the next DL/I call, even if it is for the same CI RBA. A new buffer must be obtained, and the data must be re-accessed. Re-accessing the data might require that it be:

- found in the LKASID private pool
- read from either the CF or DASD

The discarded buffers count towards the dependent regions' NBA limit, but they will be stolen when NBA is reached. The impact is that applications that use FLD calls or VIEW=MSDB processing against Shared VSO, and that repeatedly issue DL/I calls against the same CI, might need more NBA specification to perform well.

*Kevin Stewart*  
*IMS Advanced Technical Support*  
*IBM Silicon Valley Lab*

### Using GENJCL.USER to Allocate IMS HALDB Data Sets

HALDB databases often contain many data sets. Creating jobs to allocate these data sets could be a time-consuming task, for two reasons. First, there might be many data sets. Second, the allocation jobs must use the data set names that IMS assigned in the partition definition process. These challenges can be easily overcome by automating the creation of allocation jobs. The article at the URL provided below explains how you can easily allocate IMS V7 HALDB database data sets using DBRC GENJCL.USER commands. The article explains how to automate creation of allocation jobs and provides examples of jobs that use that automation. You can adapt these examples for use in your installation.

For hints and tips on migration to IMS Version 7, go to:  
[www.ibm.com/support/techdocs/atmastr.nsf/PubAllNum/TD100408](http://www.ibm.com/support/techdocs/atmastr.nsf/PubAllNum/TD100408)

*Rich Lewis*  
*IMS Advanced Technical Support*  
*Dallas Systems Center*

# Hot Websites

## IMS

[www.ibm.com/software/data/ims](http://www.ibm.com/software/data/ims)

## IBM Database Tools for OS/390 and z/OS

[www.ibm.com/software/data/db2imstools](http://www.ibm.com/software/data/db2imstools)

## VisualAge for Java

[www.ibm.com/software/ad/vajava/](http://www.ibm.com/software/ad/vajava/)

## WebSphere Application Server

[www.ibm.com/software/webservers/appserv/](http://www.ibm.com/software/webservers/appserv/)

## IBM Redbooks

[www.redbooks.ibm.com](http://www.redbooks.ibm.com)

## IBM User-Centered Design

[www.ibm.com/ucd](http://www.ibm.com/ucd)

## IBM Technical Conferences

[www.ibm.com/training/conf/us/ims/](http://www.ibm.com/training/conf/us/ims/)

## IMS Version 7 Migration Support

[www.ibm.com/support/techdocs/atasmast.nsf/PubAllNum/TD100408](http://www.ibm.com/support/techdocs/atasmast.nsf/PubAllNum/TD100408)

## IMS Newsletter

[www.ibm.com/software/data/ims/shelf/quarterly/](http://www.ibm.com/software/data/ims/shelf/quarterly/)



## Recent 2002 IMS News

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

**Editor-in-Chief: Glenn N. Galler**  
**Creative Director: Stacy A. Newman**  
**Managing Editor: Sandy Sherrill**  
**Graphic Artist: Robert Lee**  
**Senior Web Consultant: Tommy Peterson**

**Publisher: E&G Printing**

We welcome your feedback on any of the material contained in this issue.

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/shelf/quarterly/](http://www.ibm.com/software/data/ims/shelf/quarterly/).

or write to us in any of the following ways:

email: [imsmkt@us.ibm.com](mailto:imsmkt@us.ibm.com)

fax: (408) 463-4101

US Mail: IMS Newsletter

IBM Corporation

555 Bailey Avenue

DQY/A2

San Jose, CA 95141-1003

The graphic designer for this issue of the IMS newsletter is Robert Lee. Rob is a student from The State University of New York at Binghamton where he is a double major in psychology and studio art.

## They Must Be In Here Somewhere

The first couple of weeks following the mailing of the IMS newsletter are very enjoyable for me. This is the time when I receive your guesses on the locations of my children in the IMS newsletter. I also am able to read your feedback on the IMS newsletter. It is my great pleasure to read and share your comments with the entire IMS newsletter team.

### We want to hear from you.

We will send an IBM IMS pin to every person who can find the two pictures of my children correctly in this issue.

You can download a copy of the IMS newsletter in PDF format, or you can view the IMS newsletter on the IMS Web page at [www.ibm.com/software/data/ims/](http://www.ibm.com/software/data/ims/).

If you would like to receive an IBM IMS pin, please e-mail your external postal address and the picture locations in this issue of the IMS newsletter to: [imsmkt@us.ibm.com](mailto:imsmkt@us.ibm.com)

Or write to:

Editor-in-Chief, IMS newsletter  
IBM Corporation  
555 Bailey Avenue  
DQY/A2  
San Jose, CA 95141-1003  
U.S.A.



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

Glenn N. Galler  
Editor-in-Chief

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



© International Business Machines Corporation  
Printed in the U.S.A. 4/02  
All Rights Reserved

AIX, AS/400, BatchPipes, CICS, DB2, DB2 Universal Database, Data Joiner, DataPropagator, DataRefresher, IBM, IMS, IMS/ESA, MQSeries, MVS, OS/390, Redbooks, S/390, VisualAge, WebSphere, and z/OS are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

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


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

Adobe and Adobe Acrobat are trademarks of Adobe Systems Incorporated.

Other company, product, and 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.



GC26-9655-10

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