

## Assicurazioni Generali Makes the Move to HALDB

By Giampaolo Fiorese, IMS/DB2 Technical Support Team Leader, Assicurazioni Generali



Founded in 1831 and headquartered in the Adriatic port city of Trieste, Italy, Assicurazioni Generali ranks as Italy's largest insurance company and is the third largest insurance company in all of Europe. Assicurazioni Generali employs over 57,000 people worldwide with branches and offices in five continents. Through its network of subsidiaries, Assicurazioni Generali ranks among the 20 largest insurance companies in the world in terms of written premiums.

From our data center in Mogliano Veneto, Assicurazioni Generali controls approximately 1,000 agencies and subagencies throughout Italy, and for the past 30 years IMS has played a key role in our data center operations. Currently, we run six IMS V7 systems on OS/390 V2.10. Three are used for production operations, two are used for pre-production operations, and the remaining one is used for testing purposes. Our production systems typically process between 600,000 and 700,000 complex transactions per day, with each transaction generating, on average, more than 150 DL/I calls. All six IMS systems use ETO, and all are non-data sharing. Until recently, all of our IMS data (approximately 600 GB compressed at a 4:1 ratio) was stored in DL/I databases, which contained approximately 60 percent of our total customer data. DL/I databases served our business well over the years; however, as our volume of data grew along with our business, the limitations associated with the DL/I structure began to have an impact, and many of our DL/I databases began to exceed their 4 GB limit.

In 1997, we attempted to alleviate the restrictions imposed by DL/I's 4 GB limit by implementing the IMS Partition Database (PDB) product. PDB did allow us to bypass the 4 GB limit

and has been considerably stable since the initial implementation, but it had several functional limitations that were restricting, including:

### No primary index partitioning

All primary indexes for data partitions needed to reside in the same single data set. (Note that this limitation was eliminated in V2 of PDB.)

### No parallel recovery and reorganization

Recovery operations are performed at the partition (data set) level. This means that for databases that span more than a single partition, recovery is performed partition by partition. Reorganization is always performed at the database level.

### No partial database command

Commands can be issued against an entire database only, not against a partition.

### Low availability

Database availability is not granular; the entire database is either available or it is unavailable.

These limitations prompted us to look for a different solution, and we entered the Quality Partnership Program (QPP) for IMS/ESA V7 to explore the new High Availability Large Database (HALDB) functionality offered in this release.

... see Assicurazioni Generali on page 2



### In This Issue:

2002 Letter to Customers .....	3
eBusiness: CAM .....	4
Using FICON Channels .....	6
IMS UCD .....	8
IBM DM Tools .....	10
IBM IMS Support .....	11
IBM IMS Services .....	11
IMS Education .....	12
Upcoming Events .....	12
Hints & Tips .....	13
IMS Redbooks .....	14
Hot Websites .....	15
Editorial .....	16
Trademarks .....	16

*continued from page 1*

IMS HALDB enables you to partition a single database into a large number of partitions, each of which may be a smaller, more manageable element. HALDB addresses the following DL/I limitations:

### Capacity

HALDB significantly increases the amount of data that can be stored in a single database. Each HALDB database allows a maximum of 1,001 partitions, and each partition allows up to ten data set groups.

### Availability

Partition independence leaves more of the database available.

### Manageability

Smaller sections of the database are easier to manage.

On the surface, HALDB seemed to provide a comprehensive solution to these limitations, but the question of just what was involved in migrating our DL/I databases to the HALDB structure remained to be answered.

Before we could begin migrating to HALDB, we needed to understand the impact of several key requirements uncovered during the planning phase.

One involved secondary indexes — specifically, the requirement that HALDB indexes must be PSINDEX databases. The way in which secondary indexes are migrated to HALDB meant that, for all databases with a secondary index, we needed to migrate secondary index databases concurrently with their target databases before any HALDB in the database group could be used. Migrating to HALDB, though, provides significant benefits where secondary indexes are used because it allows secondary indexes to be managed separately from the target database. For example, reorganization of a non-HALDB database results in the secondary indexes for that database being recreated entirely from scratch. As a result, any user data contained within the secondary index is lost. Reorganization of a HALDB database, on the other hand, doesn't involve secondary indexes, which means that secondary indexes don't need to be created, and any user data contained therein is preserved.



We also needed to address the requirement that logically related databases must all be partitioned if any are partitioned. Logical relationships between HALDBs and non-HALDBs are not allowed, which means that we needed to migrate all of our logically related databases concurrently.

Another more basic issue that we needed to understand was the actual process of migrating from non-HALDB to HALDB. For most databases, migration to HALDB is straightforward: data is unloaded from the existing database, the database definition is changed to the new HALDB parameters, and the data is then reloaded into the new database. To help design our HALDB structure, we used the HALDB Migration Aid Utility (DFSMAID0) included in IMS V7. The Migration Aid Utility scans existing databases and provides you with statistics and recommendations that you can use to determine appropriate partition requirements for your new HALDB structures. You can specify any of the following three criteria as input to the Migration Aid Utility:

- Key range - the high key value desired for each HALDB partition
- Number of bytes per partition - the maximum number of bytes for each HALDB partition in the new database
- Number of partitions - the number of HALDB partitions desired in the new database

One limitation of the Migration Aid Utility worth noting is its lack of support for secondary indexes. Secondary indexes must be analyzed manually. The output of the Migration Aid Utility is a file of high keys per partition. This file can be used as input to the partition definition utility.

After we carefully designed our HALDB environment, we were ready to begin unloading our existing data. This step required us to select the most appropriate options considering the current structure of our data. The HD Reorganization Unload utility has been enhanced to include functions for migrating to HALDB structures. There are two migration options to choose from when unloading your data: MIGRATE and MIGRATX. The MIGRATE option is used primarily to migrate PDB and full function databases that do not have secondary indexes. The MIGRATX option, delivered by PTF UQ47416 (APAR PQ37015), was introduced to provide optimized performance when migrating databases with secondary indexes to HALDB.

*... see Assicurazioni Generali cont'd page 7*



I would like to take the opportunity to update you on the "health" of IMS. 2001 was a challenging year as we expanded our license base and our customer usage of IMS. As we continue to grow in 2002, I want to assure you of our direction and commitment toward continued IMS integration with your enterprise.

**IMS V7**

We began the year with the availability of our newest release, IMS™ V7. Some of the most notable new features are: High Availability Large Database (HALDB), Rapid Network Reconnect (RNR), enhanced connectivity, and IMS Java/XML support.

Now, after its first year in the field, how's IMS V7 holding up? I am very happy to report that IMS V7 is performing excellently. It's our best quality release in the 33-year history of IMS. Compared to the previous two releases, IMS V7 has received fewer customer defects (APARs) across the board. Customer migrations to IMS V7 began the year at a 25% faster rate than any previous release. More customers driving fewer defects is a great combination. Our overall customer satisfaction rating, measured through defect trailer surveys, also supports our improved quality. Through 2001, IMS achieved the highest customer satisfaction rating of any product in IBM® Data Management. Both the IMS team and I would like to thank you for your commitment and support of IMS!

Since IMS V7 General Availability, we have delivered additional product features through our service process. A few worth mentioning are:

IMS V7 Enhancements	Available in
OSAM/Log buffers above 2GB bar	PQ42127
Fast Path storage above the 2GB bar	PQ51419
250 OSAM Database Buffer Subpools	PQ41126
Providing Signon screen after LUSTAT x'082b'	PQ42530
Including the TCP/IP address in IMS messages	PQ41143
<b>IMS External Subsystem Attach for DB2 Groups</b>	PQ42180
Secondary index migration to HALDB	PQ37015
ILDS creation in HD Reload Utility	PQ36991
DFSMAID0	PQ37127
Recon Partition List Command support	PQ38822
Batch command initialization of HALDB and associated partitions	PQ35893 PQ44468
Unconditional HALDB Partition Initialization	PQ49638
<b>JDBC access to IMS DB from CICS/390 and DB2/390 environments</b>	PQ54548
System Log Data Set (SLDS) Read Support	PQ50657
Provide WebSphere access to IMS DB	PQ57320

ORS Enhancements	Available in
Stacked image copy support	PQ42443, PQ46522, PQ47471, PQ51623
Compressed Image copy support	PQ47216
Stacked log data set support	PQ45321

IMS Connect 1.1 Enhancements	Available in
Local Support	PQ45057, VAJ 3.5.3 + JR15813
Unicode Support	PQ47906
ACK/NAK required notification support	PQ46195
Output message structure change	PQ48182

**IMS e-business**

PQ53944 provides one of the most notable e-business enhancements to IMS. This APAR allows you to build CICS applications and DB2® stored procedures that can use JDBC to access IMS data using SQL statements. In addition, PQ57320 supports similar access using JDBC from a WebSphere for z/OS and OS/390® Enterprise Java Bean. These two APARs are extremely important to our long-term enterprise integration strategy, as one of our goals is to use industry-standard interfaces in building applications to access IMS. These two APARs support the same JDBC calls we introduced in IMS V7 for accessing IMS data from an IMS application and that are traditionally used to access relational data in systems like DB2.

Our job is to make the underlying database technology transparent to application developers. Our growth will be determined by our success in this technology. If you are interested in this technology, I encourage you to also examine PQ57328 and PQ57329, which provide the tooling support to enable you to rapidly build applications using JDBC. We will continue to bring our e-business application development function to the market at a rapid pace.

**In Closing**

The IMS team is positioned and committed to our key initiatives: 1) integration with the enterprise, 2) improved product quality, and 3) world-class customer support. These three initiatives become more important to us as we grow our e-business investments. Our customers have a tremendous investment in existing applications and data. Non-disruptive

expansion of IMS applications and access to IMS data is key to every installation.

The EMEA Technical Conference is in Vienna this June, and the US IMS Technical Conference is in St. Louis this September. I would like to thank you for your support of the IMS products and hope to see many of you at our events this year.



*Bob Gilliam*  
IMS Family Product Manager

## Common Application Metamodel (CAM) and IMS e-business

### Challenge for business application integration

The current trend for new applications is to embrace open Web standards that simplify construction and scalability. As new applications are built, it is crucial to integrate seamlessly with existing systems while introducing new business models and new business processes.

Today there are many IMS and CICS® application programs that run crucial business processes. It is important to leverage and reuse these existing assets to provide interoperability with applications from different systems. However, integrating legacy applications with new software is a difficult and expensive task due, in large part, to the necessity of customizing each connection that ties together two disparate applications. There is no single mechanism to describe how one application may allow itself to be invoked by another. Imagine that you need to pass an account number to an IMS transaction application program from a Web browser to transfer \$50.00. To do this, you will need to acquire new applications that can map to proprietary formats to integrate seamlessly with existing IMS applications, which could be written in COBOL, PL/I, or Assembler. Moreover, these existing applications are most likely based on 3270 screen formats.

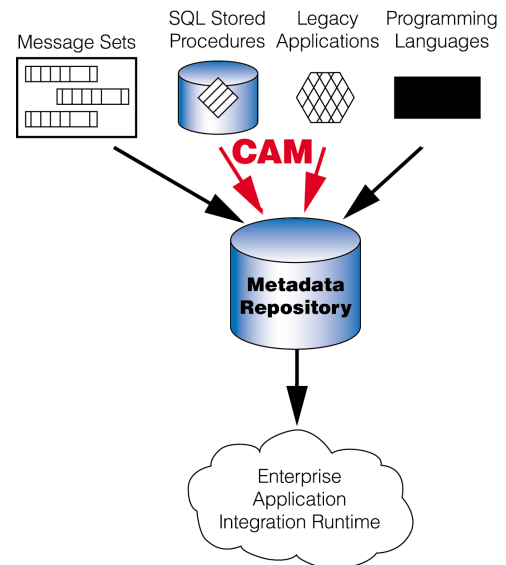
### CAM simplifies application integration

The Common Application Metamodel (CAM) makes IBM the first company in the industry to solve the problem of exchanging data between new software and legacy applications. CAM defines a standard for accessing enterprise applications defined in proprietary structures, including 3270 screen formats. This is called “metadata.” The goal is to simplify application integration by standardizing application metadata for invoking and translating application information. By doing so, CAM insulates businesses from further evolving technology and facilitates the reuse of existing assets in the new B2B environment without changing legacy applications. Through these metamodels, CAM provides a simple, standard, and predictable method of enabling data exchange between disparate systems, regardless of hardware or software platform.

CAM allows users to develop tools that can analyze existing enterprise applications and to describe these applications in a format that either can be exchanged among tools or that can be stored online in a repository for data transformation in a mixed-language environment. CAM opens the door to mapping between any industry-standard format and enterprise applications. Ultimately, CAM helps to modernize enterprise applications and to support execution of these applications from new industry-leading software using industry-standard mechanisms.

### CAM - A collection of interface metamodels

The IMS e-business team is leading the IBM cross-division CAM project, which includes members in five different locations: Silicon Valley, Toronto, New York, Raleigh, and Hursley. The team has defined a set of metamodels that consists of C, C++, Java, COBOL, PL/I, and High Level Assembler, with Type Descriptor and TDLang being language neutral metamodels used for data translation. In addition to language metamodels, CAM also includes metamodels for IMS transaction messages, IMS MFS (Message Format Service), and CICS BMS (Basic Mapping Support).



In a nutshell, CAM is needed for:

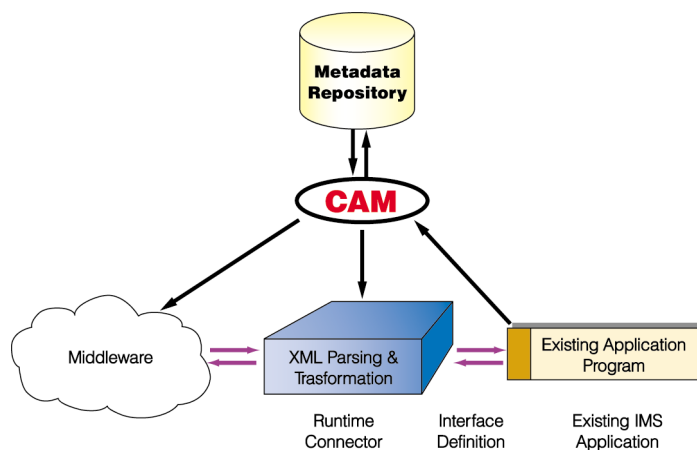
- Connector and/or connector-builder tools (development time)
- Data transformation in an enterprise application integration environment (execution time)
- Data type mapping between mixed languages
- Data translations from one language and platform domain into another
- Data-driven impact analysis for application productivity and quality assurance
- Viewing of programming language data declarations by developers

### Language Metamodels

The language metamodel (e.g., the COBOL metamodel) is used by enterprise application programs to define data structures that represent connector interfaces. An association between language metamodels (semantics) and the physical layout metamodel (syntactic) is necessary in order for the marshaller to correctly format the byte string. It is important to connector developers that connector tools show the source language, the target language, and the mapping between the two languages. The language metamodel is also used to support data-driven impact analysis for application productivity and quality assurance.

### Physical Representation Metamodels

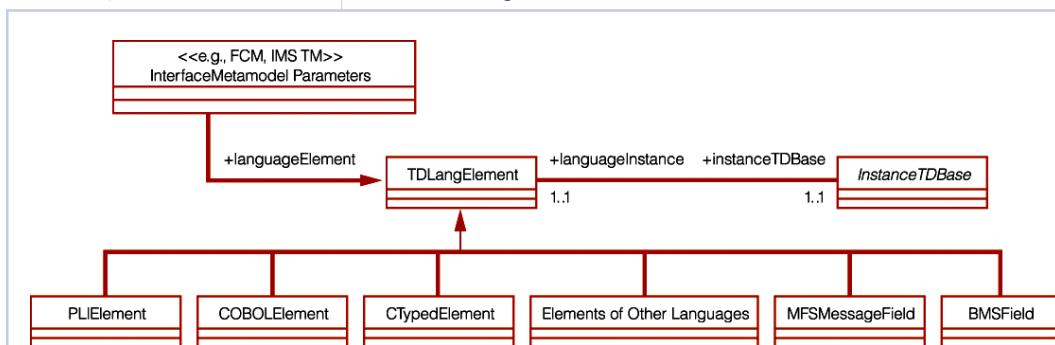
CAM has a Type Descriptor metamodel, which presents a language- and platform-independent way of describing implementation types, including arrays and structured types. This information is needed for marshalling and for connectors that have to transform data from one language and platform domain into another. Inspections of the type model for different languages can determine the conformance possibilities for the language types. For example, a long type in Java is often identical to a binary type (computational-5) in COBOL, and if so, the types may be inter-converted without side effect. On the other hand, an alphanumeric type in COBOL is fixed in size and, if mapped to a Java type, loses this property. When converted back from Java to COBOL, the COBOL truncation rules may not apply, resulting in computation anomalies. In addition, tools that mix languages in a server environment (e.g., Java and COBOL in CICS and IMS) should find it useful as a way to determine how faithfully one language can represent the types of another. Therefore, an instance of the Type Descriptor metamodel describes the physical representation of a specific data type for a particular platform and compiler.



### Enterprise Application Interface Metamodel

The application-domain interface metamodel describes signatures for input and output parameters and return types for enterprise application system domains. One good example of the enterprise application interface metamodel is the IMS Transaction Message (IMS TM) metamodel. Another good example is the Flow Composition Model (FCM), which is needed by IMS and CICS (pseudo) conversational transactions. FCM describes interactions and flows of information between applications. The payload of IMS TM and FCM metamodels needs CAM to define the signature (input and output) of applications to provide language-specific and wire formats for component inputs and outputs, and enables connectors and transformers to be defined and built so that data can flow properly between components.

In the following figure, TDLangElement is the parent class of all CAM language element classes. TDLangElement represents typed unit elements declared in a copybook or source code (that is, typed data elements without a subcomponent). Sample subclasses of TDLangElement are COBOLElement, PLIElement, and MFSMessageField given that the I/O messages for MFS-based IMS applications are language independent. TDLangElement also provides association with a physical representation model for each language element and MFS messages.



### Connectors Adapters and Web Services

The CAM team has also made major contributions to two OMG marketplace standards: EDOC (Enterprise Distributed Object Computing) and EAI (Enterprise Application Integration). Having the set of CAM models adopted by the OMG demonstrates IBM's leadership in the field of non-homogeneous data exchange to the IT industry.

In addition to CAM's acceptance by the standards committee, IBM has also started to rally around and adopt CAM internally. IBM's very first tooling solution, WebSphere® Studio Application Developer Integration Edition (WSADIE) 4.1 featuring CAM models for Importer, Adapter, and Web services, was Golden Mastered on March 1, 2002. With this new tool, customers are able to continue Web-enabling existing IMS COBOL and C applications and also to publish these applications on the Internet as Web services.

As part of the IBM's CAM efforts, IMS e-business is also building connector/adapter solutions for MFS-based applications. This support will allow customers to invoke MFS-based IMS applications via XML documents, and will also allow customers to publish MFS-based applications on the Internet as Web services. Businesses can then use evolving technologies to facilitate reusing transactions in the new B2B environment without changing the existing IMS applications.

*Shyh-Mei Ho  
IMS Senior Developer*



## Using FICON Channel to Improve IMS Performance and System Manageability

With the availability of the Native Fibre Connection (FICON™) support, IBM 2105 Enterprise Storage Server™ (ESS) can now offer significantly improved IMS performance and system manageability over ESS connected via Enterprise System Connection (ESCON®).

The FICON channel is a new interface for zSeries™ and S/390®. It uses Fibre Optic Technology and is based on the industry standard Fibre Channel architecture. The new FICON support is a successor to ESCON support. It relieves ESCON architecture limitations, while supplementing and coexisting with ESCON support.

Some FICON Benefits over ESCON include:

- Higher data transfer rate
  - FICON: 100 MB/sec in full duplex mode
  - ESCON: 18 MB/sec in half duplex mode
- Higher effective throughput
  - FICON: 60-65 MB/sec
  - ESCON: 10-13 MB/sec
- Channel Command Word (CCW) pipelining and frame multiplexing
  - This feature supports asynchronous command execution, thus concurrent I/O.
- Longer distance attachment support
  - FICON supports up to 20 km direct attach.
- Channel consolidation-reducing cabling infrastructure
  - A typical 4:1 ESCON-to-FICON channel folding in IMS environment is possible.
- Increased number of unit addresses per channel
  - FICON: 64K addresses
  - ESCON: 4K addresses
- In an intra-sysplex IMS MSC transaction routing VTAM® CTC environment, we have observed 10% throughput improvement in IMS CTC and 5% in VTAM CTC when using FICON over ESCON.

### Measurement Environments

The following online logging, BMP, and utility measurements were taken using an environment that consisted of IMS V7 running on OS/390 V2.8, a 2064-116 processor, and ESS-F20. A few selected database tools were also used. The following Channel Consolidation measurements using Data System Workload (DSW) were taken using an environment that consisted of IMS Version 6 running on OS/390 V2.8, a 9672-RX6 processor, and ESS-F20.

### Online IMS Logging Evaluation

Over the years, the online IMS logging bandwidth (logging capacity) requirement has continued to grow due to the higher online transaction rate and the higher degree of concurrent BMP processing. Connected via FICON, ESS is able to sustain 40% higher logging of 26 KB blocks via BSAM while utilizing much less channel path capacity.

Figure 1 compares the bandwidth measured of IMS Online Log Data Sets (OLDS) using ESCON and FICON channels. The workload used was a Fast Path application with a large volume of logging data such that the capacity of the IMS log was the primary inhibitor to higher transaction throughput. Each transaction generates approximately 26 KB of log data.

	ESS-ESCON	ESS-FICON	Improvement
Log data / sec	14.5 MB	20.5 MB	41%
Transaction rate / sec	557.7	788.5	41%
Single channel path busy	92.6%	14.8%	84%

Figure 1. IBM DASD measurements from a log-constrained IMS environment

### BMP and Utility Evaluation

The measurements shown in Figure 2 compare various I/O-intensive processes against a Fast Path Data Entry Data Base (DEDDB) while using ESS connected via ESCON and FICON channels. The area involved contains 24 KB control intervals and consists of 362,405 root segments. The three processes observed are: an account consolidation BMP, a database tools (DBT) DEDDB unload (FABCUR1+SORT), and a DBT reload (FABCUR3).

	ESS-ESCON	ESS-FICON	Improvement
BMP elapsed (sec)	320	241	25%
Unload elapsed	75	54	28%
Reload elapsed	172	107	38%
Total	567	402	29%

Figure 2. IBM DASD measurements for I/O-intensive processes

### Channel Consolidation Evaluation

The following IMS DSW/VSAM measurements have validated a frequently recommended rule-of-thumb: 4 ESCON-to-1 FICON when ESCON is about 50% average in utilization. Figure 3 shows that 4 FICON has outperformed 8 and 16 ESCON measurement scenarios in IO/sec, mean response time, transaction rate, and mean transaction elapsed time.

DB and Logs	IO/sec	Mean Resp Tm (ms)	Tran Rate	Tran Mean Elapsed Tm (ms)
4 FICON	5,969	6.3	541	75
8 ESCON	5,616	8.4	532.4	111
16 ESCON	5,692	5.7	538.6	78

Figure 3. IBM DASD measurements for channel consolidation

There are other I/O-intensive processes in IMS that can benefit from faster and more effective data transfer from FICON. For example, I/O-intensive database data sets, Fast Path VSO data hardening, and Shared Message Queue structure checkpoints. Any process that involves heavy I/O, even if it is for a short burst of time, is a candidate process for using ESS with FICON. For technical details about FICON, see the IBM Redbook titled *Introduction to IBM S/390 FICON*, SG24-5176.

*Cedric Chen, Hsiung Tang, and Bovorit Pibulsonggram  
IMS Performance Specialists*

*continued from page 2*

We also had to decide how to define our HALDB partitions. When IMS V7 first became generally available, the only option available was the Partition Definition Utility (PDU). The PDU is an ISPF-based tool used to create and manage partitions. In addition to defining partitions, the PDU also can be used to delete partitions or to modify partitions either before or after they're initially loaded (required, for example, if your initial settings proved to be too small).

Soon after GA, however, DBRC batch commands grammar was extended to provide an alternative method for defining HALDB partitions. Specifically, the INIT.DB command was enhanced to include the TYPHALDB parameter, which allows you to register a HALDB to DBRC. The ability to actually create a HALDB partition structure was provided by a new DBRC command, INIT.PART. Although HALDB batch initialization is not as versatile as the PDU (for example, modifying and deleting partitions is not supported by HALDB batch initialization, and some INIT.PART parameters apply to all the partition DBDSs created as a result of this command, which is different from the PDU where these parameters can be specified separately for each partition DBDS being created), HALDB batch initialization proved to be the appropriate approach for us because it is run as a standard batch job.

Another decision we had to make was whether to use the same database names for our HALDBs as we used for our non-HALDB databases, or assign new names. This choice has implications for the process of unloading and reloading data. For example, if you choose to use the same names, you need to save your RECON information for fallback purposes before deleting the records for your databases. If you choose different names, RECON information can be retained indefinitely if desired. Also, if you choose to use new database names, all involved PSBs will need to be changed. Note that if you place a DBD in the DBDLIB with the same name as the old database as a logical database with all segments set up as "source=the new HALDB," then ACBGEN will be required, but PSBs need not be changed. We found that using the existing names was the best choice for us because this method did not require us to modify the PSBs.

When it came to reloading our data, we had the choice of using either of the two HD Reorganization Reload utility input statements: NOILDS or ILDSMULTI. The NOILDS option was introduced in V7 as a way to improve performance during the migration reload by disabling ILDS processing. By deferring the rebuilding of ILDSs until after the data reload, less CPU and less DASD resource are consumed. Of course, immediately after the reload process is finished, you must rebuild the ILDSs. This is accomplished by using the HALDB Index/ILDS

Rebuild utility (DFSPREC0). The design of DFSPREC0 is based on the belief that if ILDS recovery is required for multiple HALDB partitions, time is critical. Therefore, it is faster to submit several jobs to recover multiple HALDB partitions concurrently than it is to submit a single job to recover them serially. DFSPREC0 can recover the ILDS (as well as the prime index) for a HALDB partition during a single execution of the utility.

If system resources are not an issue though, consider using the ILDSMULTI input statement. When you use ILDSMULTI, the data reload and ILDS insert processing happen in parallel. This process uses more CPU and DASD resource, but honors free space. It also has a much shorter total elapsed time than using NOILDS and then subsequently rebuilding the ILDS by using DFSPREC0.

Overall, our migration of our test and pre-production systems from non-HALDB to HALDB has gone smoothly. Our largest HALDB, so far, is over 10 GB with software compression on 9 partitions. By mid-spring 2002, we plan to have completely migrated our production systems as well.



Assicurazioni Generali's leverage of newer IMS technologies does not stop with HALDB, though. ETO has been in full production since 1991, and the MQSeries IMS Bridge has been in production since its GA. We participated in a beta program for Online Recovery Service in 2001 and found that it was easy to use and that it had good performance. As such, we will use it as a replacement for batch recovery processes in our development environment. The most exciting new projects, though, are Web-based IMS applications. One is a recently deployed IMS application that uses IMS Connect and WebSphere Application Server running on OS/390 to pull data from a DB2 database. Plans are in the works to develop more applications like this. Some will be entirely new, and some will use existing code. And our interest in the latest IMS enhancements continues. We will be participating in the IMS V8 QPP program to evaluate various APPC enhancements and enhancements to Java and XML, including: Java-dependent regions, Java standards enhancements, and XML support.

*Author's Note: I would like to thank the following people for their outstanding efforts in making the transition to HALDB. I present this case study on their behalf:*

*Marina Canestrelli  
Massimo Pettenò*

*IBM Global Services - ITS  
Assicurazioni Generali IMS/DB2  
System Programmer*

*Federica Polese*

*Assicurazioni Generali IMS/DB2  
System Programmer*

*Fabrizio Piaggi*

*Assicurazioni Generali IMS/DB2  
Senior Application Developer*

*Francesco Zanotel*

*Assicurazioni Generali IMS/DB2  
System Programmer*

## User-Centered Design in IMS

Have you ever worked with IMS and thought, I wish they designed this feature differently; it makes more sense to do it this way? Or, If only I had this option, my job would be so much easier? Or, There's got to be an easier way of doing this? If you answered yes to any of these questions, then we need you as part of the IMS design team. Users are essential to ensuring that new features in IMS both solve important problems and are easy to use. That's why the IMS team uses a process called User-Centered Design (UCD). UCD is based on four basic principles:

- Understand users
- Design the total user experience
- Evaluate designs iteratively
- Assess competitiveness

### Understand Users

Even a long-established product like IMS has a constantly changing set of users. Experienced users are learning and trying to integrate IMS with new technologies. New users come to IMS with a very different set of skills, experience, and expectations than those of a decade ago. Organizations change the way their work is organized and who is expected to do what.

Following the UCD methodology means that we talk with users and compare our assumptions with their reality. The IMS design team cannot design in a vacuum. Prior to design we need to understand the user's roles and tasks. For example, in recent years we have expanded our definition of system programmer, refined our knowledge expectations for operators, and added specialist roles for automation and performance. These changes influence how we write our documentation, what APIs we make available, and what new tools we develop. To understand how users interact with the product, we typically use one or both of these methods:

**Customer Visits** In a customer visit, we come to your shop and interview you about how you do your job. But the interview is not just talk — we need to watch you as you work. What are the important tasks? How often do you perform them? Whom do you work with, and what information passes between you and other team members? Where do the breakdowns occur? When we better understand your work and problems, we can use our understanding of the technology to design solutions that even you may not think of.


**Group Task Analysis** In a group task analysis, you come to us. Groups can include one- to three-dozen users gathered together to describe how they organize their work. In a sense, everyone can talk at once because we use group software on networked computers. Each user types in his or her work, and the software consolidates it into categories. We display the categories back to the users, who then rank the importance and their satisfaction with how they do it today. Like the customer visit, this helps us to identify the biggest problems.

### Design the Total User Experience

Ease-of-use is often equated with a graphical user interface (GUI), but UCD takes a broader view of usability. UCD looks at ease-of-use in everything the user sees and touches. This includes: how you learn about new releases and how you obtain and install them; the commands, APIs, messages, and documentation for using the product; technical support, product fixes, and education.

As you can imagine, this means that the whole IMS team is involved in the process: marketing, technical architects, programmers, visual designers, testers, writers, support, and human-computer interaction specialists. Each brings a unique perspective and deliverables that are important to the user. That is why you, the user, are the central and essential member of the team — it is your experience with the product that we want to improve.

**Scenarios** One technique we can use even before a new product or feature is designed is the end-to-end scenario, which describes how a user obtains the product or feature, installs it, uses it, resolves problems with it, and so on. The scenario typically includes all parts of the product and identifies whom the user interacts with. We sit down with users and step through the scenario with them. They can identify what they like, what they don't like, and maybe what's missing.






## Evaluate Designs Iteratively

When we talk about evaluating designs, we mean observing users' reactions to the designs. We see what problems they encounter and listen to their feedback. Iteration means we make changes based on user feedback and evaluate the changes again and again to ensure that the problems have been fixed. For example, IMS V8, which is entering QPP this spring, has been through four rounds of design walkthrough, as well as hands-on testing and the following remote studies:

**Design Walkthrough** This group activity allows users to explore a design more thoroughly than is possible with just a scenario. The UCD team describes the design in detail, including the problem we are trying to solve and the externals of the solution. Externals may include new command options, syntax, or whole new APIs. When the design involves an ISPF or GUI interface, we include paper mock-ups of the screen and ask users how they would expect to interact with them.



**Hands-on Testing** When a design has progressed to the point where we have a working prototype, we ask users to complete tasks with it. We don't give too much instruction — just a goal and whatever materials the user might normally expect, like documentation. Then we shut up and watch. It is often humbling to see how much trouble users have with what we thought was a good design. But this is how we improve the design and ensure that when users get the final design in their office, it is easy to use.

**API Studies** Hands-on testing is not restricted to ISPF or GUIs. For example, when designing the new IMS Java API, we asked Java programmers to create applications. We gave them the class libraries and the documentation, and watched as they wrote the Java code using a text editor. The users could compile the code and look for errors. These observations helped us improve the documentation and make the API more usable.

**Remote Studies** Sometimes users come to the Silicon Valley Lab to do a hands-on test with our designs, but more and more often we are conducting studies remotely. Using remote control and online meeting software, we can have a prototype running on a computer in Silicon Valley, and the user can be interacting with it from another state, another country, or anywhere in the world. This saves time and money and allows for quick follow-up with new designs.

In evaluating designs, we look at two measures of success: task performance and user satisfaction. We want users to be able to complete tasks quickly and with few errors. If the tasks take a long time or the user has many errors, we don't blame the user — the failure is in our design and we need to change it. But even more importantly, we want users to be satisfied that the design is useful and easy to use. So at the end of each study, users rate their satisfaction with the design. This is our ultimate measure of success.

## Assess Competitiveness

Ease-of-use is often relative to the state-of-the-art, as well as what is available in competitor products. Though it is not easy for many customers to replace IMS with another product, it is possible to migrate to different solutions or to develop new applications without IMS. UCD includes competitor products in studies to ensure that we know what users like and dislike in the alternative solutions.

**Side-by-side Comparisons** With any of the user studies described above, we can include both the IMS solution and a competitor's solution. Users may prefer some parts of one and some parts of another. By identifying their preferences, we are able to make changes that result in a better overall solution. We also evaluate a feature often left out of ease-of-use studies — price. When two products are mostly equal, price can be the deciding factor. Even when one product is significantly better, an inferior but adequate product might be preferred if its price is less.

## You Can Be Involved

In the past few years, IMS has conducted over 4,000 hours of UCD studies. In the coming years, we will be developing many new administration tools, enabling application access from new technologies, and simplifying the base product function. In all of these projects, we need users on the team to ensure that we are designing the most useful and easy-to-use solutions possible.

Participating is easy. Some studies take only an hour or two and you can participate from your own home or office. If you would like to participate in one of the UCD activities described above for IMS or any other IBM products, you can sign-up on the Web at [www.ibm.com/ucd](http://www.ibm.com/ucd). If you would like a more personal contact or more information about the process, contact Dana Marsh at [dkmarsh@us.ibm.com](mailto:dkmarsh@us.ibm.com).

*Kevin McBride  
IMS User-Centered Design  
Tracy Hutcheson  
IMS Human Factors Engineer*

## What's New in Tools

The Data Management team continues to announce new and updated tools for IMS and DB2. Recent arrivals for IMS included IMS Network Compression Facility, IMS High Performance Change Accumulation Utility (HPCA), and new versions or releases of several other tools.

As a potential or current Data Management tools user, you probably have some questions about books and CD-ROMs for IMS tools:

- What information arrives with my order?
- Where are the User's Guides?
- Are there any BookManager® books available?
- Can I order a CD-ROM with IMS tools books?
- How can I effectively use the information provided?

### Information that Arrives with an Order

When an IMS tool ships to a customer, the box contains the software product on tape or CD-ROM and a hardcopy of the Program Directory. The Program Directory is intended for the system programmer responsible for installation and maintenance. The contents include a description of the basic and optional programming materials and documentation, the IBM support available, program and service level information, installation requirements and considerations, and installation instructions.

### Locating User's Guides

To customize and use the product, you will need to refer to the User's Guide. Most IMS tools have a single User's Guide, though a few of the tools have more than one User's Guide. In order to provide frequent updates to documentation, we do not print the guides in hardcopy. Instead, the guides are delivered electronically from the Web and provided on CD-ROMs.

The Program Directory tells you where to locate the publications on the Web. You will find the publications both on the library and on the product pages at this Web address:

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

You can download the files to your workstation, LAN, or Intranet server at no fee.

There are two formats available on the Web:

- Portable Document Format (PDF) is ideal for printing pages or an entire book.
- BookManager books (served as HTML on the Web) are excellent for searching and locating information within a large book or library of books.



### Books on CD-ROM

As an IMS user, you are probably familiar with using publications on CD-ROMs, especially in BookManager format, which have been available for a number of years for z/OS™ products. The Data Management team supplies both BookManager and PDF files on the following collections:

- z/OS Software Products Collection (SK3T-4270), BookManager and PDF
- OS/390 Software Products Collection (SK2T-6700), BookManager format

The two CD-ROM collections also contain the IMS publications, so your company might already receive a copy of these collections. The collections are manufactured quarterly.

When IBM delivers a new or enhanced tool, the publications are immediately available on the Web. The publication appears on the next quarterly release of the CD-ROM collection.

### How to Order CD-ROMs

To order your CD-ROM, use our e-business site, IBM Publication Center, at [www.elink.ibm.com/public/applications/publications/cgi-bin/pbi.cgi](http://www.elink.ibm.com/public/applications/publications/cgi-bin/pbi.cgi).

Select the country of your choice, and click **Go**. Then click **Search for Publications** to display the search page. In the Publication number field, enter the form number for the CD-ROM (SK3T-4270 or SK2T-6700), and click **Go**. Add the CD-ROM to your shopping cart, and check out. The site has hundreds of IBM books, some of which can be ordered in hardcopy for a fee, and many of which are in PDF and BookManager formats for download at no fee.

### Keeping Current with Documentation

With the flexibility of electronic delivery, the Data Management team can more easily refresh publications. Thus, the Web edition of the book could be more current than the CD-ROM edition. You can easily find out if you have the most recent edition of the book. Check the date of the edition notice on the inside of the title page of the PDF book.

### For More Information

For more information about how to effectively use the electronic documentation visit [www.ibm.com/software/data/db2imstools/support.html](http://www.ibm.com/software/data/db2imstools/support.html). Select Hints and Tips and search for "Books and CD-ROMs."

Tell us how we can improve Data Management Tools information by contacting us at [comments@vnet.ibm.com](mailto:comments@vnet.ibm.com).

*Mary Paquet  
Data Management Tools  
User Technology Team Leader*

## IMS and IMS Tools Web Support

On July 10, 2001, IBM launched a new software support strategy, in part, to improve, simplify, and enhance existing support. As a result of this strategy, you can now access all IBM support for IMS and IMS Tools through their respective support Web pages, regardless of your location or time zone.

### IMS Support Page

Your IMS Support hub is now the IMS Support page: [www.ibm.com/software/data/ims/support.html](http://www.ibm.com/software/data/ims/support.html).

From the IMS Support page, you can now do the following:

[Submit and track problems](#)

[Conduct an advanced search across all platforms](#)

[Access over 140 frequently asked questions \(FAQs\) and technotes that are updated weekly](#)

[Find 1,230 Hints & Tips \(HATs\) that are no longer available in information authorized program analysis reports \(APARs\).](#)

[Track and download program temporary fixes \(PTFs\)](#)

[Search all open and closed APARs](#)

[Access the preventative service planning \(PSP\) buckets](#)

[View the IMS Technical manuals Web page, which now includes the following two licensed manuals:](#)

*[IMS Version 7 Diagnosis Guide and Reference](#)*

*[IMS Version 7 Failure Analysis Structure Tables for Dump Analysis](#)*

(Note that to access these books, you must supply your IMS customer number.)

[Locate IMS Redbooks](#)

[Link to IMS Presentations and Whitepapers](#)

[Look up error message text](#)

You can do all of this in seven different languages through dynamic translation with the help of WebSphere.

### IMS Tools Support Page Now Available

Data Management (DM) Tools also now offers an external support page to provide 24 X 7 service and improve customer satisfaction. The DM Tools support page is available at [www.ibm.com/software/data/db2imstools/support.html](http://www.ibm.com/software/data/db2imstools/support.html).

The DM Tools support page lets you search FAQs, HATs, closed APARs, and Redbooks. You can also find the latest product enhancements made available through PTFs by going to the "Support Flashes" section. Additionally, a direct link is provided to search PSP buckets, order PTFs electronically, and view a complete PTF listing for all DM Tools. From the left-hand navigation panel, you can also link to the latest product information, services, events, news, and downloads. Be sure to check out the DM Tools support page regularly because a variety of support documentation and useful links will be added on an ongoing basis.

*Sara Guyer*

*IBM Data Management Tools Specialist*

*Tommy Peterson*

*IMS User Technology*

## IMS Services for Tools

IBM offers a number of services that you can take advantage of when purchasing IMS tools:

**Assessment:** IBM will combine remote calls and notes with on-site visits to help you plan and size the amount of work required to migrate from an OEM tools vendor environment to a Data Management (DM) Tools environment. This work includes understanding the number of items to be changed, the functional usage mix of the tools being replaced and installed, and the amount of effort you wish to contribute to the migration. This work normally will also uncover unusual usage issues specific to your environment.

**Education:** IMS can provide both on-site and remote classes on various tools. The remote classes currently use a combination of teleconferencing and pre-sent materials. In the near future we hope to provide Webcast education for the most commonly taught tools.

**Installation:** IMS can provide on-site assistance to install the product(s) and to apply all applicable maintenance. This can be combined with education to allow you to gain basic familiarity with the tools.

**Migration:** IMS can provide a wide range of assistance for the actual tools migration. The range extends from remote phone and note assistance only, through short on-site visits to consult and assist during the early phase of the conversion, to having a full IBM team come in and do the conversion for you. The full conversion uses a combined IMS/IGS team that will both find existing tools jobs, and also provide new ones for the purchased DM Tools.

**Testing:** IMS can provide a wide range of assistance for tools testing. The range extends from remote note/phone support to sending people on-site to provide you with extensive assistance.

**Performance and Tuning:** IMS can provide personnel to assist you with performance tuning of the new tools. This includes both parameter analysis and traditional performance tuning activities.

**Cutover Support:** IMS can provide skilled personnel to be either on-call or on-site during key migration milestones.

For further details call Don Streicher at 408-463-2192 or email to [ibmdds@us.ibm.com](mailto:ibmdds@us.ibm.com).

*Don Streicher*

*Manager, IMS Planning and Field Support*



## Upcoming IMS Education

### Affordable IMS On-Site Training for Small Groups

For a limited time, we're offering a new on-site pricing option for 1-7 students. This offer applies to both our standard courses and to custom offerings. Here are some examples of on-site pricing for small groups:

CM010: IMS Fundamentals	
3 days for 1-7 students	\$8,280 (regularly \$10,350)
U3717: IMS DB Application Programming	
3 days for 1-7 students	\$10,044 (regularly \$12,555)
U3718: IMS DC Application Programming	
2 days for 1-7 students	\$7,148 (regularly \$8,935)
<i>Instructor travel expense is an additional charge.</i>	

And you can save another 10% by using the IBM Education Pack program. See below for details. To learn more about the IMS curriculum, visit [www.ibm.com/training/spotlight/ims/](http://www.ibm.com/training/spotlight/ims/).

### Get \$5000 Worth of IBM Education for Less

Get \$5,000 of IBM education for only \$4,500 by using the IBM Education Pack. The IBM Education Pack includes \$5,000 worth of training coupons in \$500, \$100, and \$50 denominations for IBM Learning Services public and private courses and IBM technical conferences in the USA. For more about this offer, visit [www.ibm.com/training/us/savings/edpack/](http://www.ibm.com/training/us/savings/edpack/).

### New Courses

You owe it to yourself to check out IBM Learning Services' IMS curriculum. We have updated and added many courses:

IMS Fundamentals (CM010) covers the basic IMS facilities, how they work together, and how application programs interact with them in today's enterprise systems.

IMS V7 Product Enhancements (CM701) covers the new features delivered in IMS V7, including High Availability Large Databases (HALDB) and Online Recovery Services (ORS).

IMS Database Performance & Tuning (U3720) explains how to tune IMS databases in IMS/Batch, IMS/DC, CICS-Local-DL/I, and DBCTL environments.

IMS Logical Relationships: Implementation & Performance (U3724) shows you how to implement IMS logical relationships and to improve their performance.

IMS Connectivity & Web Enablement Workshop (CM651) covers how to implement IMS in an e-business environment, how to set up and administer IMS communications, and how to enable existing IMS transactions and IMS databases.

Call 1-800-IBM-TEACH to enroll in any IBM course, or for more information, contact Linda Stanislawski (LSTAN@us.ibm.com).

*Amiet Goldman  
Senior Direct Marketing Strategist  
Joe Kreitz  
IMS Worldwide Marketing Manager*

## 2002 IMS Technical Conference

**Sept 30 - Oct 3, 2002 St. Louis, Missouri**

St. Louis is the ONE STOP you need to make in 2002 to learn about the new and enhanced IMS V7 and V8 technologies for transforming your enterprise into an e-business superhighway. Join us for a most special learning experience at the IMS Technical Conference, where you will:

- Dive into the depths of IMS Connect for e-business legacy integration. See how to easily use it with IMS Connector for Java to develop new solutions that can readily access existing IMS applications and data.
- Learn how to expand the capabilities of your enterprise with IMS V8 in areas such as Sysplex terminal management, Java and XML, APPC/OTMA, MSC, and recovery.
- Get the details on significant enhancements that IMS V8 provides for managing your IMS systems — Sysplex-Wide Resource Manager, Single-Image Operations Manager, Syntax Checker, and Transaction Trace.
- Learn how to deploy Java applications, servlets, and Enterprise JavaBeans in a WebSphere Application Server environment to access IMS transactions.
- Gain insights on how to increase your e-business database capacity to over 40 terabytes with HALDB, including valuable hints and tips on how to migrate an existing IMS database to HALDB.
- Explore how you can improve systems and data management capabilities to get the best possible performance ever with IBM's robust portfolio of new and enhanced IMS tools.
- Get a jumpstart on migration considerations that will help you plan for a smooth and orderly transition to the latest IMS releases.
- View demonstrations of what's new and next in IMS e-business solutions at the Product EXPO, where you will also have an opportunity to address specific issues with key exhibitors.
- Hear the hottest tips, tricks, and techniques through dozens of eye-opening sessions from product developers and technical experts that will have you developing better and faster the day you get back to the office.

Please bookmark our conference Web site, as we will be continually updating it with the latest details on this event: [www.ibm.com/training/conf/us/ims/](http://www.ibm.com/training/conf/us/ims/).

You won't want to miss it!

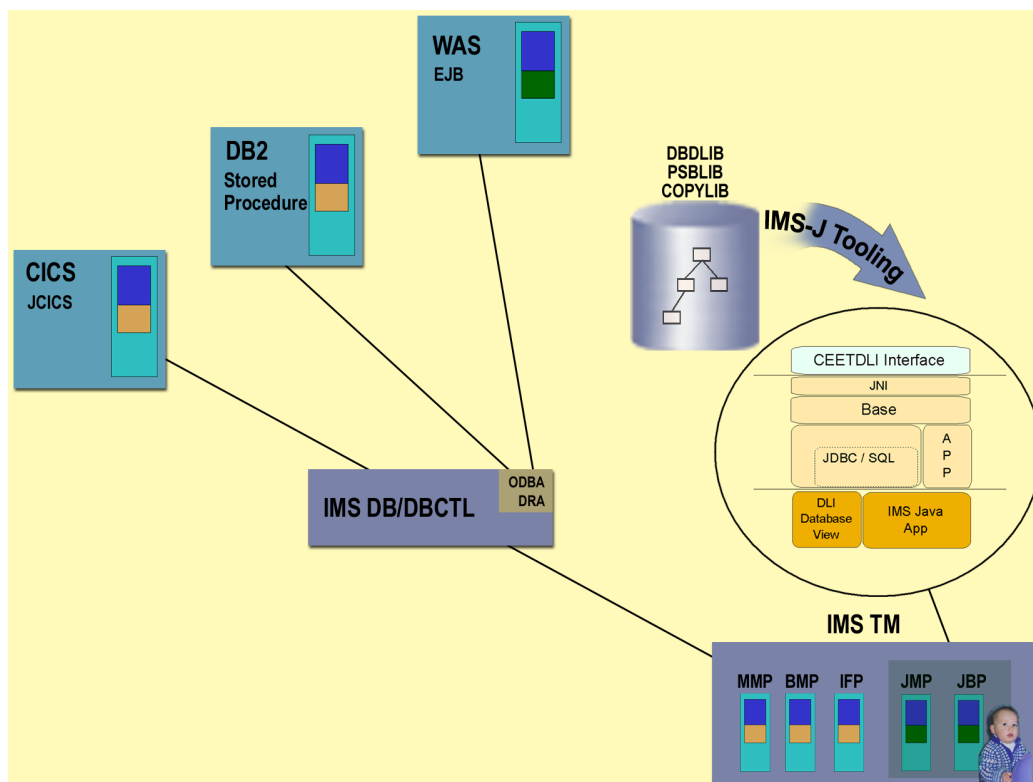
*Amiet Goldman  
Senior Direct Marketing Strategist*

## JVM vs. HPJ

IMS is pleased to announce several new enhancements to our support for Java application development. To briefly summarize, our offering is a layered set of class libraries that allow you to write Java programs to access and manipulate IMS databases. The highest layer is an implementation of the JDBC application programming interface. JDBC is the common programming interface for database connectivity for Java. Using this layer, experienced Java programmers can quickly and easily write applications to access IMS data without detailed knowledge of IMS itself.

We introduced support for Java in IMS V7. The initial offering allowed you to use the IBM High Performance Java (HPJ) compiler to build Language Environment (LE) applications to run in IMS dependent regions, including MPP, BMP, and IFP regions. To IMS, these applications appeared no different than any other LE load module. IBM has since made the decision to stabilize the HPJ technology at the JDK 1.1.8 and JDBC 1.0 level of Java. This limits our ability to take advantage of emerging technology, so we will not support HPJ-compiled applications in IMS V8.

The good news is that we now have full support for running IMS applications in a Java Virtual Machine (JVM), allowing you to take advantage of all the features of Java, including the ability to dynamically load classes (something you could not do with HPJ). To support the use of a JVM, we introduced two new dependent regions designed specifically for running Java applications. The new region names are JMP (analogous to an MPP) and JBP (analogous to a non-message driven BMP). Adding this feature has allowed us to update our Java support to JDK 1.3 and JDBC 2.0. This support will appear in IMS V8 and is already available in IMS V7 via APAR PQ53944 (the dependent region support) and APAR PQ54039 (the IMS Java support for JDK 1.3).



We have also introduced Java support for several additional running environments. With APAR PQ54548, you can now build applications that can access IMS databases from CICS TS 2.1 and DB2 V7 stored procedures using the same JDBC interface you use in an IMS application. Soon, we will introduce similar access from WebSphere 4.0.1 via an Enterprise Java Bean (EJB).

Up until now, we required you to manually build a set of Java classes that describe the metadata of your IMS databases. In our JDBC driver, we use this metadata to automatically map the SQL queries you build to the necessary Segment Search Arguments (SSAs) and DL/I calls to access the data. This metadata is also used to automatically convert from one data type to another. With APARs PQ57328 and PQ57329, we have added a tool to automate this process for you. This tool parses DBDs and PSBs and creates the Java metadata classes necessary for us to process your SQL queries and create the SSAs. This tool also produces a user-friendly report for Java programmers called the Java Report, which provides the information necessary to code SQL queries (segment names, field names, field types, etc.).

We look forward to hearing from you about our latest product offering.

*Kyle Charlet  
IMS Developer*

**The following new IMS Redbooks will be available during spring 2002:**

**Ensuring Data Integrity Using IMS Tools (SG24-6533)** explains and documents a collection of IBM Data Management Tools for IMS that provides functions for controlling, validating, recovering, and repairing IMS data. This book provides the descriptions and usage examples for the following IMS Tools: IMS Database Control Suite, IMS Database Repair Facility, IMS High Performance Pointer Checker, IMS Image Copy Extensions, and IMS Index Builder.

**IMS V7 Java Update (SG24-6536)** introduces you to enhancements in the IMS Java support for IMS V7. This book helps you install, tailor, and configure the IMS TM, CICS, and DB2 environments that support Java applications that use JDBC to access IMS databases. The new persistent and reusable Java Virtual Machine (JVM) environment that IMS provides is described with the new Java dependent region types, Java Message Processing (JMP) region, and Java Batch Processing (JBP) region. The book provides a sample auto dealership application that may be helpful when designing your own Java applications that access IMS resources.

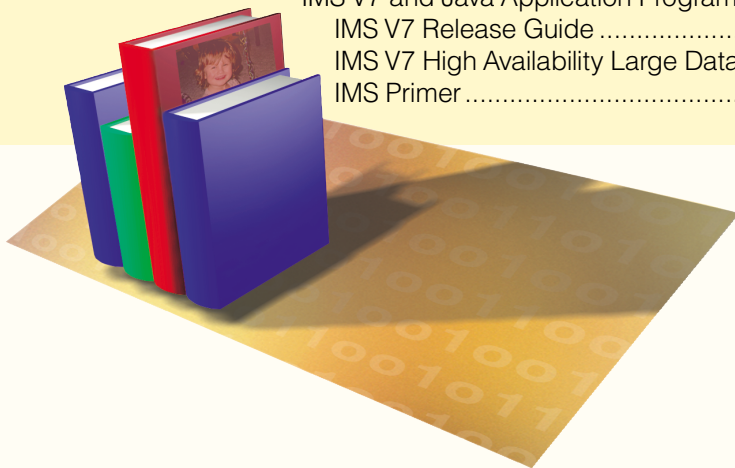
**IMS e-business Connectors: A Guide to IMS Connectivity (SG24-6514)** describes the changes in the IMS e-business connectors that have been implemented over the last two years. This book documents the changes to IMS Connect, documents the Java™ support added to VisualAge® for Java, provides an ODBA example, and documents changes to IMS support in WebSphere® Application Server. The book is full of examples and detailed descriptions of the processes needed to access IMS from the World Wide Web and other e-business clients.

For more information about IBM Redbooks™, joining a Redbook residency project, or to download any IBM Redbook, visit the IBM Redbooks Web site at [www.redbooks.ibm.com](http://www.redbooks.ibm.com).


*Jouko Jääntti  
ITSO IMS Specialist*

The most recent Redbooks are available at the ITSO Web page ([www.redbooks.ibm.com](http://www.redbooks.ibm.com)) and include the following titles:

TITLE	Order Number	Publication Date
IMS V7 Performance Monitoring and Tuning Update .....	SG24-6404 .....	January 2002
IBM Web-to-Host Integration Solutions .....	SG24-5237-03 .....	January 2002
Using XML on z/OS and OS/390 for Application Integration .....	SG24-6285 .....	November 2001
A DBA's View of IMS Online Recovery Service .....	SG24-6112 .....	March 2001
IMS V7 and Java Application Programming .....	SG24-6123 .....	February 2001
IMS V7 Release Guide .....	SG24-5753 .....	September 2000
IMS V7 High Availability Large Database Guide .....	SG24-5751 .....	September 2000
IMS Primer .....	SG24-5352 .....	January 2000





**IMS**

<http://www.ibm.com/ims/>

**IMS Support**

<http://www.ibm.com/software/data/ims/support.html>

**DB2 and IMS Tools**

<http://www.ibm.com/software/data/db2imstools/>

**IMS Tools Support**

<http://www.ibm.com/software/data/db2imstools/support.html>

**IBM Publication Center**

<http://www.elink.ibm.com/public/applications/publications/cgi-bin/pbi.cgi>

**IMS Learning Services**

<http://www.ibm.com/training/spotlight/ims>  
<http://www.ibm.com/training/us/savings/edpack/>

**IBM Technical Conferences**

<http://www.ibm.com/training/conf/us/ims/>

**ShopzSeries**

<http://www.ibm.com/eserver/shopzseries/go/>

**IBM Redbooks**

<http://www.redbooks.ibm.com>



**The IMS Newsletter was awarded the level of Merit in the 2001 Touchstone Competition of Northern California this past March. This competition is sponsored by the Society for Technical Communication, which is an international professional association that advances the arts and sciences of technical communication.**

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

**Publisher: Corporate Printers**

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  
 DQY/A2  
 San Jose, CA 95161-9023

**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

It is such a great pleasure to receive your notes with your guesses on the locations of my children in the IMS newsletter. The number of winners of the IBM IMS pin continues to grow with each issue, and your feedback is very much appreciated by the IMS 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: [imsmtk@us.ibm.com](mailto:imsmtk@us.ibm.com)

Or write to:

Editor-in-Chief, IMS newsletter  
IBM Corporation  
DQY/A2  
San Jose, CA 95161  
U.S.A.



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. 4/02  
All Rights Reserved

BookManager, CICS, DB2, Enterprise Storage Server, ESCON, FICON, IBM, IMS, MQSeries, OS/390, Parallel Sysplex, Redbooks, S/390, VisualAge, VTAM, WebSphere, z/OS, and zSeries are 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.



GC26-9655-09

IMS newsletter  
IBM Corporation  
DQY/A2  
San Jose, CA 95161-9023

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