

Avnet and IMS MFS Web Services

Paul Oh, Principal Engineer, Avnet Global Information Solutions



Avnet, Inc. is one of the world's largest Business-to-Business (B2B) distributors of semiconductors, interconnect, passive, and electromechanical components, enterprise network and computer equipment, and embedded subsystems from leading manufacturers. Headquartered in Phoenix, Arizona, Avnet serves customers in 68 countries. In 2005, Avnet was listed as number 217 on the Fortune 500 list of largest U.S. companies by revenue, and was also ranked third in InformationWeek's 2004 annual list of the best users of information technology. Avnet reported revenues of \$11.07 billion for fiscal 2005, ending July 2, 2005.



from IMS databases. Both of Avnet's U.S. telesales locations maintained separate copies of the PowerBuilder-based application and accessed locally-stored Oracle data. When a salesperson needed to enter a quote or sales data into the IMS system, they had to open a 3270 session in the background in order to do so.

Figures 1 and 2 show the high-level design of Avnet's original IMS system.

continued on page 2

Pass 1: Retrieve customer profile

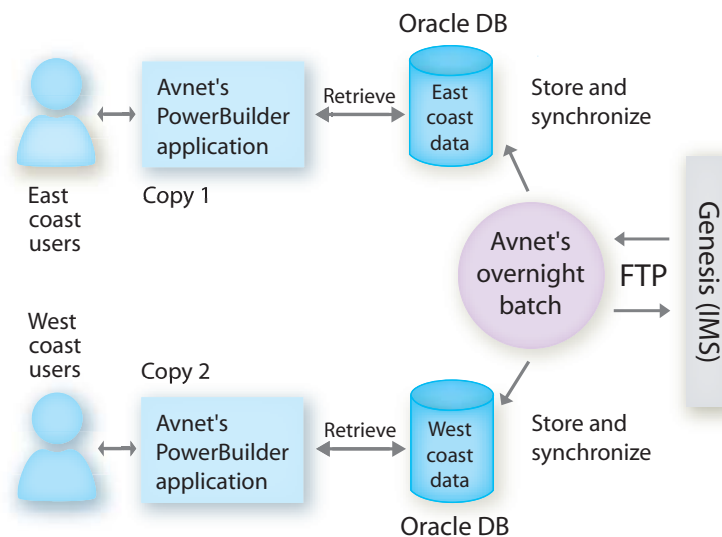


Figure 1

Pass 2: Enter a quote or sales data

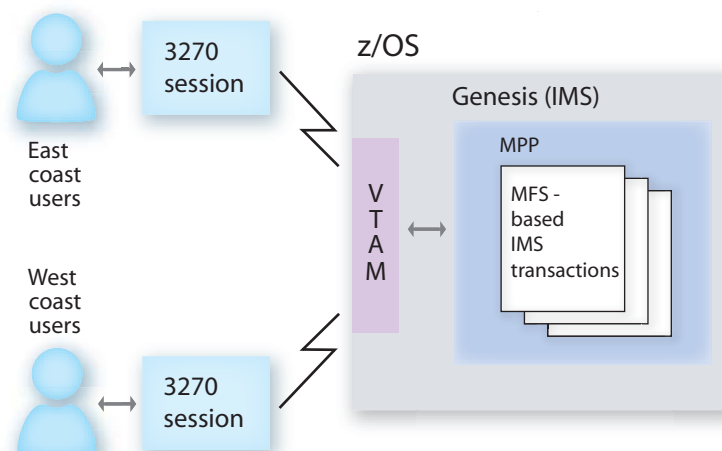


Figure 2

For 10 years, Avnet used a Sybase® PowerBuilder-based Microsoft® Windows NT® application as a front end to interface with its IMS™ system in its telephony system. In this system, incoming customer calls were intercepted by a phone switch and passed to the PowerBuilder-based application, which triggered a screen pop-up. The screen pop-up contained a pre-populated customer profile, including sales, invoice, and call history data. This information appeared before each call was answered by a salesperson, providing them with vital customer service information.

The customer data was retrieved from Oracle® tables that were refreshed nightly through downloads

IMS NEWSLETTER

Fall 2005 In This Issue

Letter to Customers
page 3

IMS Security
page 4

IMS Profiles
page 5

IMS Version 9
page 6

IMS MFS Web Support
page 8

IMS Technical Conference
page 10

Data Management Tools
page 12

Hot Websites
page 15

Editorial
page 16

Avnet and IMS MFS Web Services

continued from cover page

In order to improve the speed and efficiency of the overall process, Avnet wanted a more flexible architecture. They needed to transition from a Windows NT-based platform to a Web services-based application environment in order to meet this goal. This transition required a complete redesign of their system. It also was crucial that Avnet maintain the same basic functionality in their new application as in their original system.

Avnet needed to develop a new Web services-based application that could invoke both their Oracle and IMS data. Avnet chose to use the IBM IMS MFS Web Services solution to access IMS because of the solution's flexibility, total cost of ownership, open architecture, and its ease of integration into other applications. With the redesigned system, the IMS MFS Web Services would be generated from existing MFS source files using IBM WebSphere® Studio Application Developer Integration Edition Version 5, running on IBM WebSphere Application Server. Figure 3 shows the high-level design of Avnet's new system.

The new Avnet system, named Express Customer Profile, is a complete redesign of the original system. Express Customer Profile uses a Java®-based Web application, and connects with the new customer interaction infrastructure using Avaya Interaction Center software. The system accesses the centralized application and database servers and communicates with the back-end IMS system.

Existing IMS transactions are accessed from the Web application without any required programming changes. This is accomplished by running the IMS Resource Adapter (IMS Connector for Java) on WebSphere Application Server, with IMS Connect running at the host. IMS transactions are seamlessly integrated into the new Web application with sub second response times. IMS transaction-level security is maintained using the RACF® userid and password (specified as IMSConnectionSpec properties).

By using IMS MFS Web Services, Avnet has been able to rearchitect how it drives its existing MFS-based IMS transactions. Avnet can use its existing business logic in the evolving Business-to-Business environment without using VTAM® and a 3270 emulator. This helps reduce the total cost of ownership while also revitalizing existing 3270-based transactions.

As an IMS customer for 18 years, total cost of ownership was a major factor in Avnet's decision to use IMS MFS Web Services. Avnet's Vice President of Business Applications, Bob Pischke, describes IMS MFS Web Services as "...a better application at a lower cost — and certainly more sustainable."

The IMS MFS Web Services solution is available as part of the IMS Integration Solutions Suite at ibm.com/ims.

Paul Oh
Principal Engineer
Avnet Global Information Solutions

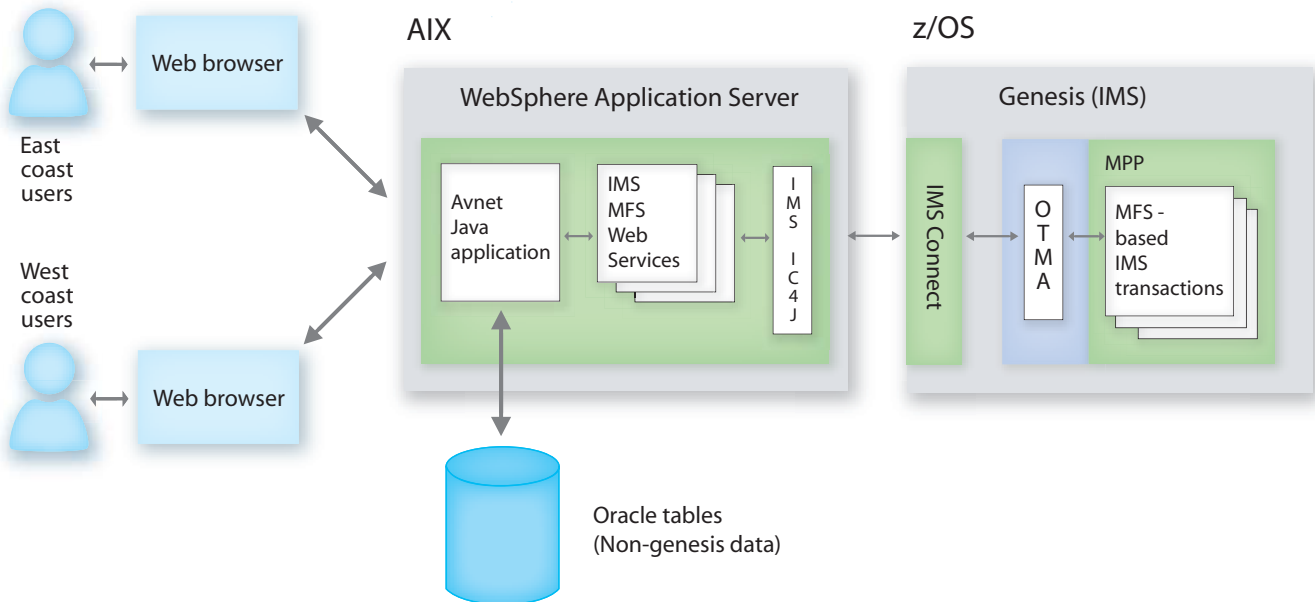


Figure 3

Bob Gilliam bids farewell...



After five years as the IMS Product Director, I've accepted a new position within the IBM™ Tivoli® organization. Working with IMS clients and the IBM development and support teams has been one of the highlights of my career. It was an extremely difficult decision to leave one of IBM's best and most respected products. I will miss the many IMS clients and friends that I have met around the

world. IMSers are a close-knit community comprised of IBM employees, business partners, and clients who are proud of their technical knowledge and business contributions. IMSers understand that they are the key to running the core of the world's top business applications. I feel proud to be associated with such a group and thank you for supporting me through five terrific years.

My new position allows me the opportunity to visit with many of you. I am the Director of IBM Tivoli's Availability and Business Service Management team. This area has some of the finest systems management products that aid in managing your operations, understanding your business processes and systems, tracking your service levels and advising you of the corrective actions. I hope to visit future IMS Technical Conferences with some of my peer team members from the NetView (host and distributed), the Candle acquisition (including OMEGAMON® for IMS), and Tivoli Monitoring (including the Tivoli Enterprise Portal).

Now for the good news: Greg Lotko has been named as my replacement! Greg has already been associated with IMS in one of his previous jobs as the IMS Customer Service Manager. I feel very confident that Greg will continue to move the IMS product forward while ensuring that service to our clients remains the top priority.

Greetings from Greg Lotko

I'm sure that many of you are thinking "deja vu" as you read this. As Bob pointed out, I'm not new to IMS. I decided to return to IMS because of the wonderful people who are associated with IMS both within IBM and outside in our clients' shops. And so much business runs on IMS. IMS is where the action is — it runs the world!

Throughout my 17-year career at IBM, I've held a variety of positions including IMS and DB2® financial application programmer, business analyst, project manager, insurance application delivery executive, and several senior management positions within IBM Software Group product development. Because of my experience working on IBM internal and out-sourced applications, I feel more able to put myself in the shoes of our

clients. I can understand the problems and opportunities they face each day. With this knowledge, I will strive to continue a good working relationship with our clients and focus on what makes their jobs easier.

Prior to taking my current position, I was the Director of Development for WebSphere Information Integration and was responsible for implementing product features that provided access to IMS DB and IMS TM. While in that position, I met with a number of clients and became keenly aware that preserving their existing investment and leveraging it with new technology was better than migrating away from IMS.

As IMS Product Director, I plan to keep the team's focus on our clients. Bob did an excellent job of establishing good relationships with clients and I hope to expand upon the strong foundation that he created. We see our clients as our partners. This means that if their business isn't successful, ours certainly can't be. Our goal is to continue this partnership and to make IMS an integral part of their business by going forward to help them achieve your goals. With this in mind, our folks know that when our clients have a problem, it's our number one priority to just do the right thing and help wherever we can.

Before we look into the future, I would like to share some of our IMS accomplishments over the past year. IMS Version 9 General Availability was in October 2004 with a number of our early Quality Partnership Program clients already in production. With this and a number of new and improved tools and service stream enhancements, IMS clients have been provided significantly improved integration and continuous access to critical business information with unlimited database capacity for on demand service-oriented architectures. IMS clients have also been provided rapid response, rapid deployment of new Web-based applications at the lowest possible cost, and solutions for managing their increasingly complex environments.

As we look to the future, key initiatives from the past are still at the forefront for IMS. These initiatives include increasing product quality and customer responsiveness and service, and delivering product technological advancements that address our clients' continually-evolving needs.

Two exciting technical conferences are scheduled for 2005. In October, we have the IMS Technical Conference in San Jose, California, and in November, we return to Koenigswinter, Germany for the IMS Technical Symposium. I hope you can join us at one or both of these conferences. I look forward to seeing many new and familiar faces at these upcoming events.



Encrypt your IMS and DB2 data on z/OS

Today's complex IT environments make the need for both data privacy and compliance with data security regulations even more critical. IT professionals are looking for solid, efficient solutions to meet these needs. Data encryption technology is a key component of any such solution.

We often mistakenly assume that only financial and insurance institutions face concerns about data security. However, because of recent world events and the rapid increase in data piracy crimes, data security has become a paramount issue for all businesses, including local, state, and federal governments, manufacturing companies, retail companies, and airlines.

IBM offers a unique data encryption product for the z/OS® platform: IBM Data Encryption for IMS and DB2 Databases. This product is the only data encryption offering for both IMS and DB2 on z/OS.

The problem

Many IT shops must conform with data privacy regulations, but the available resources for meeting those regulations are scarce. Therefore, encryption solutions must be both efficient and easy to implement.

Although many IT shops write their own data encryption software, such software can be difficult to maintain and support for long periods of time. Common challenges with maintaining and supporting any data encryption software include:

- Performance overhead
- Application changes
- Key management
- Software costs
- Security of diagnostics (dump analysis)

The IMS and DB2 z/OS data encryption solution

The IBM Data Encryption for IMS and DB2 Databases tool provides you with the capability to protect all of your sensitive z/OS data, in IMS at the segment level and in DB2 at the table level, in a single product. You can implement the tool simply by using standard IMS and DB2 exit routines. The IBM Data Encryption for IMS and DB2 Databases tool exploits the features that are already built into the S/390® and zSeries® Integrated Cryptographic Services Facility hardware, which means lower overhead encryption and decryption processes.

The IBM Data Encryption for IMS and DB2 Databases tool uses the Triple DES version of the ANSI Data Encryption Algorithm (DEA), also known as the U.S. National Institute of Science and Technology (NIST) Data Encryption Standard (DES) algorithm. The DES algorithm is available in three flavors: Single DES, Double DES, and Triple DES. The Double DES algorithm encrypts the data twice, and the Triple DES algorithm encrypts

the data three times. It takes significantly more time decoding software that is encrypted with the Triple DES algorithm than software encoded with the Single DES or Double DES algorithms. The Triple DES algorithm is trillions of times stronger than the Single DES algorithm. It can take decades, or even centuries, to break the coding, even with the current rate of advancement in computer processing speeds. However, the hardware performance cost of using the Triple DES algorithm is triple the cost of using the Single DES algorithm.

The IBM Data Encryption for IMS and DB2 Databases tool is easily customizable and easy to implement because you implement it in your environment and it relies on your environment. Few application changes are required in order to use the tool, and passwords need not be stored in applications. Passwords are passed at an exit routine level, and all passwords are managed by the tool.

Encrypting DB2 z/OS data

For DB2, the Data Encryption for DB2 and IMS Databases tool exploits the DB2 EDITPROC procedure and is compatible with the DB2 Load and Unload utilities. The performance of the tool is similar to that for row-level data compression.

The Data Encryption for DB2 and IMS Databases tool provides the following support for DB2:

- Row-level encryption
- Support for all versions of DB2
- Support for both secure key and clear key encryption
- Requires no changes to applications, except to assign an EDITPROC procedure on the CREATE TABLE operation
- Applications do not need to track keys

To implement data encryption for DB2, you must define a key, build an encryption routine that uses the tool to exploit the DB2 EDITPROC procedure, unload the table, and then reload the table using the EDITPROC procedure to encrypt the data.

Encrypting IMS data

For IMS, the Data Encryption for DB2 and IMS Databases tool exit routines are implemented at the segment level. The IMS exit routines can be different for each segment. The tool uses the IMS Segment Edit/Compression exit routine.

The Data Encryption for DB2 and IMS Databases tool provides the following support for IMS:

- Support for all versions of IMS
- Support for both secure key and clear key encryption
- Requires no changes to applications, except to change a DBD to define the exit routine



To implement data encryption for IMS, you must define a key, build an encryption routine that uses the tool to exploit the IMS Segment Edit/Compression exit routine, unload the database, and then reload the database using the encryption exit routine to encrypt the data.

Summary

Several IBM customers have already implemented the Data Encryption for IMS and DB2 Databases tool, with favorable results. One of the strongest selling points for the IBM Data Encryption for IMS and DB2 Databases tool is the documentation, which includes comprehensive information about the overall process of preparing for and implementing data encryption. The information includes “cookbook scenarios” that help you to easily run the tool. For example, it might take merely one day to prepare for the encryption and only 10 minutes to run the IBM Data Encryption for IMS and DB2 Databases tool. The end-to-end process and the concepts and tasks that are associated with implementing the tool are documented thoroughly. You can find the documentation for the tool at <http://publib.boulder.ibm.com/epubs/html/decuga12/decuga12tfm.htm>

Tom Vogel

DB2 and IMS Tools User Technology

Geoff Jackson


Data Management Tools Development



Brian Kitzberger

Brian Kitzberger started working for IBM and IMS five years ago developing and maintaining IMS internal tools for test and development. He supports two IMS internal tools, Test Case Builder and Code Coverage Analysis Tool (CCAT). Brian states, “I enjoy making it easier for testers and developers to get their jobs done and in doing so, I’m improving the quality of our product.” Known for his responsiveness, Brian has added several new features to both tools. These improvements include

enhanced statistical analysis and test case logic debugging which can be accessed through the application web based interface. When Brian isn’t working on IMS tools, he enjoys restoring classic cars and lists his favorite as a 1964 Chevy Chevelle Malibu SS convertible. Born and raised in the sunny Silicon Valley, he also enjoys weekends working on home projects with his new wife.



Pedro Vera

Pedro Vera began his IBM career twenty-four years ago as an intern. He recalls a story where several employees were standing in front of a terminal, stumped by an MVS™ catalog problem. The baffled onlookers parted, and in walked an individual who solved the problem with a few key strokes. In that moment, Pedro decided that he wanted a career at IBM where he would be the “go to guy”. Being the “go to guy” certainly describes his early career as a System Programmer supporting

the local MVS system at the Silicon Valley Lab. Today, Pedro is leading a small team of programmers responsible for the TSO user interface for IMS. He gains satisfaction knowing that what he does, “makes life easier for IMS customers.” Pedro is particularly proud of the role he played at a new IMS customer site in China, meeting their requirements on time and with high quality. When Pedro isn’t working on IMS, he enjoys being a Boy Scout leader with his two sons, who share his passion for the great outdoors. He’s also been known to kick up his heels as a ballroom dancer, his favorite dances being the Cha Cha Cha and the Salsa.



HALDB Online Reorganization

Since HALDB was introduced in IMS Version 7, many IMS customers have experienced the benefits of migrating their databases to HALDB. These benefits include increased storage capacity, improved manageability of databases, and data availability. If you've been considering migrating some of your databases, or you have databases that are good candidates for migration, there's another compelling reason to do so — HALDB Online Reorganization (OLR).

OLR is a major enhancement to HALDB for IMS Version 9. OLR is designed to extend the improvements made to database availability in IMS Version 7 by allowing concurrent update activity while a reorganization process is running, keeping the databases available to all applications throughout the reorganization process. OLR eliminates reorganization as the major reason for planned database outages and is available to all IMS Version 9 database users at no additional cost.

Overview of OLR processing

OLR is an online process, which runs as an IMS task under the DLI separate address space in an IMS online system. Each OLR task reorganizes one database partition. Multiple OLR tasks can be started to concurrently reorganize different HALDB partitions. OLR applies only to HALDB PHIDAM and PHDAM databases. OLR cannot reorganize PSINDEX databases, but databases with secondary indexes or logical relationships can be reorganized.

OLR reorganizes a partition by reading database records from the database data sets and writes the database records to new data sets. HALDB uses the characters A-thru-J in the DDNAMEs and data set names. This naming convention has been expanded for IMS Version 9 by implementing the characters M-thru-V (and Y if PHIDAM) for an alternate (or paired) set of data sets. For example, the corresponding data set for data set A is data set M and the corresponding data set for B is N, and so on. For a PHIDAM database, the corresponding data set for X is data set Y. There is no corresponding pair for L and all the updates to the Indirect List Data Set (ILDS) are made to this data set for the partition. In addition, no corresponding pair exists for PSINDEX, because OLR does not currently support reorganization of PSINDEX databases. Figure 1 illustrates the naming convention changes for OLR.

Before OLR begins, a single set of data sets exist for the partition: either the A-J set (and X if PHIDAM), or the M-V set (and Y if PHIDAM). These data sets are called active data sets. After OLR has been successfully started for the partition, ownership for the OLR is established. Only one IMS owns the OLR for a partition at any time. When OLR begins, the active data sets are considered the input data sets, and the paired data sets are the output data sets. The successful initiation of OLR for the partition is recorded in the RECON with cursor-active status, meaning that the partition is now

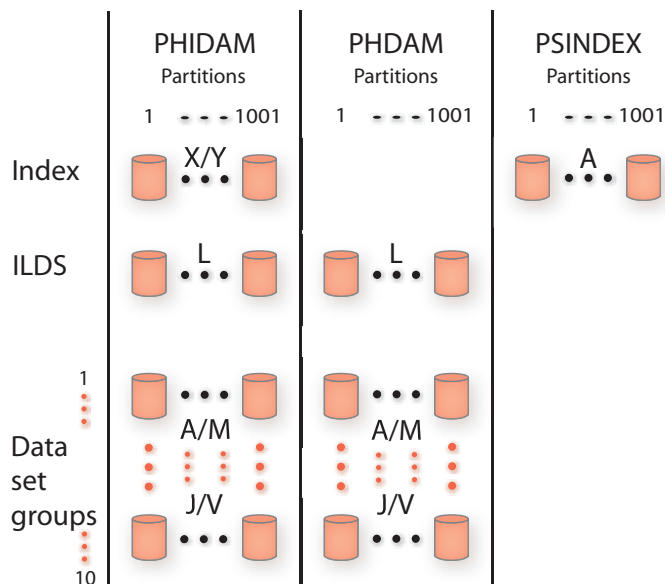


Figure 1

comprised of both the A-J data sets (and X if PHIDAM) as well as the M-V data sets (and Y if PHIDAM). After successful initiation, either OLR must run to completion, or an offline reorganization must be run to reset the OLR settings in the RECON data set. Until the reorganization completes, both sets of data sets can be accessed by applications.

Figure 2 illustrates how OLR reorganizes a partition by reading segments from the input data sets and inserting them to the output data sets. The ILDS is updated during reorganization if the inserted segment is a target of a secondary index or has logical relationships. Because application programs can concurrently read or update the partition, OLR uses locking protocols to provide data integrity. To provide data recoverability, OLR also logs all inserts to the output data sets during reorganization processing. After reorganization completes, the output data sets become the active data sets, and the old input data sets become the inactive data sets because they are now obsolete. When the next reorganization of the partition is started, the active data sets become the input data sets, and

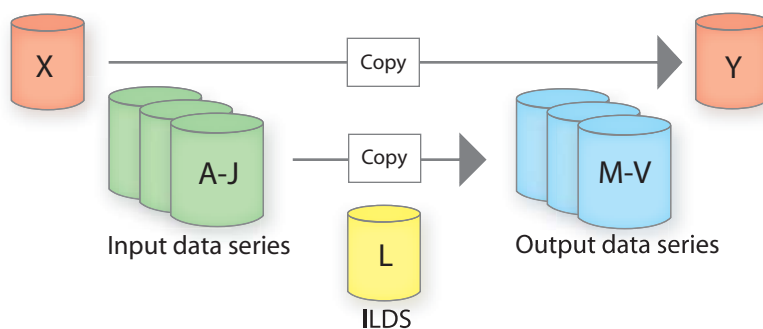


Figure 2

the inactive data sets (if they were not deleted) become the output data sets. You do not have to know which data sets are active. IMS determines this from information saved in the RECON.

Cursor-based reorganization and application availability

Figure 3 illustrates how OLR copies data. OLR is a cursor-based reorganization where the cursor identifies the last record in the partition that was copied to the output data sets. OLR reorganizes database records in groups called a unit of reorganization. As the copying proceeds, the cursor moves forward through the partition.

All the database records in a unit of reorganization are locked before they are copied to the output data sets. After copying is complete, all of the updates in the unit of reorganization are committed, the cursor is advanced, and the locks are released. The cursor is also used when IMS processes DL/I calls. IMS determines which data sets (input or output) a database record is retrieved from, and into which data sets a database record is inserted. IMS makes this determination based on the position of the record and whether it is before or after the cursor. This process is completely transparent to an application program accessing the partition.

One of the primary design goals of OLR is to provide a non-disruptive reorganization with no impact to database availability. The performance characteristics of OLR are similar to that of an online BMP application that receives a large number of updates. Just like a heavy-update BMP application, OLR needs to read and write every segment in a partition and therefore uses more resources than an offline reorganization. Even so, it is designed to be a well-behaved application. OLR uses a dynamic algorithm that autonomously adjusts the size of the unit of reorganization and limits the number of locks held during each unit of reorganization. Therefore, even though OLR creates some lock contention, the impact is minimal. In addition, when OLR creates a deadlock, it is always chosen as the victim. The other participants in the deadlock are not affected. The current unit of reorganization is dynamically backed out, and OLR is automatically restarted from the cursor. To avoid buffer and logging contention, and to isolate the logs written by OLR, you can dedicate an IMS to run OLR.

Getting started with OLR

Implementing and starting OLR is simple. Before using OLR, you need to register the HALDB with DBRC or make it OLR capable. To make a HALDB OLR capable, use either the INIT.DB or CHANGE.DB command with the OLRCAP keyword, which sets the online reorganization capable flag for all the partitions. To register the HALDB with DBRC, simply use the Partition Definition Utility. By default, all new HALDB PHIDAM and PHDAM databases are online reorganization capable when

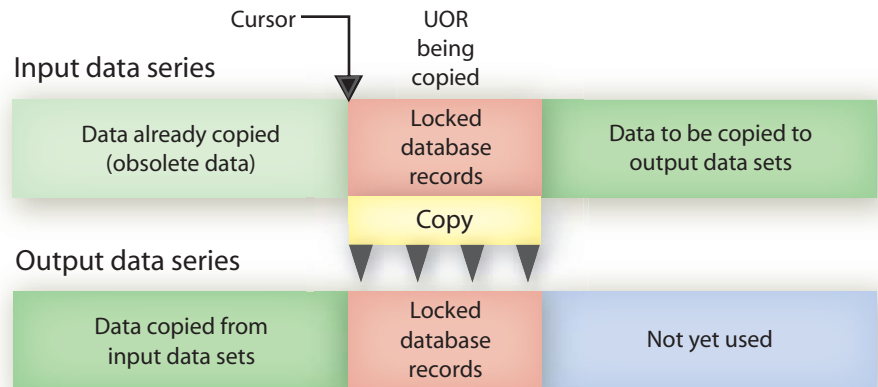


Figure 3

they are registered with DBRC. After registration, simply start OLR with either a type-1 /INIT command or a type-2 INIT command. Multiple partition names can be specified on the INIT command. Only one partition name can be specified on the /INIT command. When multiple partition names are specified on an INIT command, multiple OLR tasks are started and run in parallel, one for each partition.

DBRC and database utilities support for OLR

DBRC stores and maintains information about OLR in the RECON data sets. Each partition database record indicates the status of OLR. DBDS records for the M-V data sets are created in the RECON data sets when ownership of the OLR is established, if they do not already exist. OLR creates an online reorganization record for each DBDS. A-J and M-V data sets are treated as one logical set of DBDSs. In addition, OLR is supported by backup and recovery database utilities.

Summary

OLR is a major enhancement to HALDB. OLR is designed to address the requirement of eliminating reorganization as the key reason for scheduled database unavailability. In addition, it reorganizes data efficiently while being a good citizen in the online environment by using a dynamic algorithm to autonomously adjust the size of the units of reorganization. It's also easy to set up and use because after you register a HALDB with DBRC, you can start OLR with one command. OLR also provides full data integrity and recoverability. Finally, you can take advantage of OLR at no additional cost because it is included with IMS Version 9.

For more information on HALDB, see *IMS Version 9 Administration Guide: Database Manager* and the IMS Redbook *The Complete IMS HALDB Guide: All You Need to Know to Manage HALDBs*.

Claudia Ho
IMS Development

IMS MFS Web Support

IMS MFS Web support

As the need for IT communication grows beyond traditional computing and into the new Business-to-Business (B2B) and Business-to-Client (B2C) environments over the Web, ease of data exchange is a greater and greater concern for businesses. Existing 3270 data stream and terminal emulation facilities are limited, complex, and expensive, and no longer meet the increasingly complex needs of users. In order to meet this challenge, MFS-based IMS applications require a standard interface between these dissimilar systems. XML is rapidly becoming the universal data format that can be input to and output from any application. IBM IMS MFS Web support offers industry-standard techniques to modernize existing MFS-based IMS applications via XML processing into Web services (B2B) and Web enablement (B2C) applications.

The IMS MFS Web support solution facilitates total cost reduction by eliminating the dependency on traditional 3270 emulators and VTAM. The solution includes two product offerings:

- IMS MFS Web Services, which enable business applications to interact with MFS-based IMS transactions.
- IMS MFS Web Enablement, which enables users to seamlessly navigate through dynamically-generated MFS-like Web pages in a browser.

Both of the IMS MFS Web support solution product offerings include three components:

- MFS Metamodel, which reserves existing IMS applications by modeling MFS definitions.
- MFS Importer, which parses and converts MFS source files into XML Metadata Interchange (XMI) files.
- MFS Adapter, which translates XMI metadata to and from an equivalent MFS online processing format.

MFS Metamodel

The MFS Metamodel is a nonproprietary industry-wide UML-specified representation of the MFS definitions, and is part of the Common Application Metamodels of the OMG Enterprise Application Integration (EAI) marketplace standards. The MFS Metamodel captures the meaning of the MFS definitions, precisely representing the elements and the relationship between the device format and the message descriptor.

MFS Importer

The MFS Importer reads and parses existing MFS source files in order to generate instance XMI metadata that conforms to the MFS Metamodel. The MFS Importer is one of many importers developed by IBM to parse and generate XMI metadata based on the Common Application Metamodels. The XMI metadata represents all of the application interface information that is encapsulated by the MFS source, including the input and output messages, display information, MFS flow control, device

characteristics, and operation semantics. The four types of MFS control blocks: Message Output Descriptor (MOD), Message Input Descriptor (MID), Device Output Format (DOF), and Device Input Format (DIF), are grouped as MOD/DOF pairs and MID/DIF pairs. Each pair is stored into an XMI metadata file.

The collection of XMI metadata can reside in a repository on any system or platform, thereby replacing the traditional 3270 MFS format library on z/OS. The XMI metadata is retrieved for runtime processing by the MFS Adapter.

MFS Adapter

The MFS Adapter loads XMI metadata from the XMI repository to perform data transformation between MFS-based IMS application programs and a Web-based application such as a servlet or a Java bean serializer/deserializer. The MFS Adapter can reside on any system in a network. By sending and receiving data to and from IMS via IMS Connect and OTMA, the MFS Adapter bypasses the 3270 emulator, VTAM, and the MFS online processor.

With the new architecture, the MFS Adapter replaces the MFS online processor (inside the IMS control region), which processes MFS format and data requests, retrieves the MFS MID, MOD, DIF, and DOF control blocks that reside in the MFS format library, and sends message data to the IMS message queue.

Figure 1 shows how the MFS Adapter, on any system in a network, processes the MFS format and data requests, retrieves the XMI metadata from a repository that can reside on any platform, and sends message data to the IMS message queue.

IMS MFS Web Services

The IMS MFS Web Services solution enables business applications to access the existing MFS-based IMS business logic as Web services. The IMS MFS Web Services tooling component, the MFS Importer Wizard, is integrated as part of IBM WebSphere Studio Application Developer Integration Edition

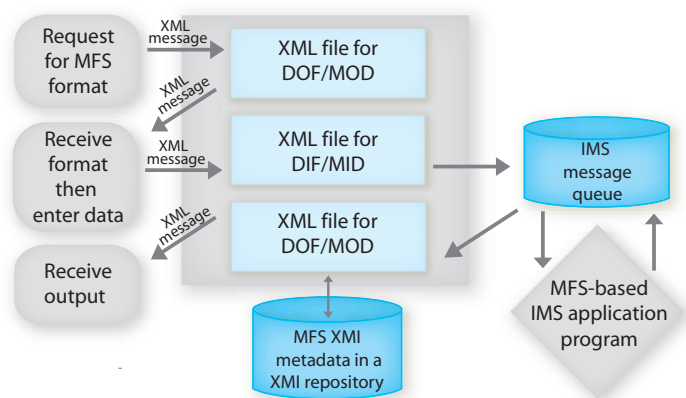


Figure 1

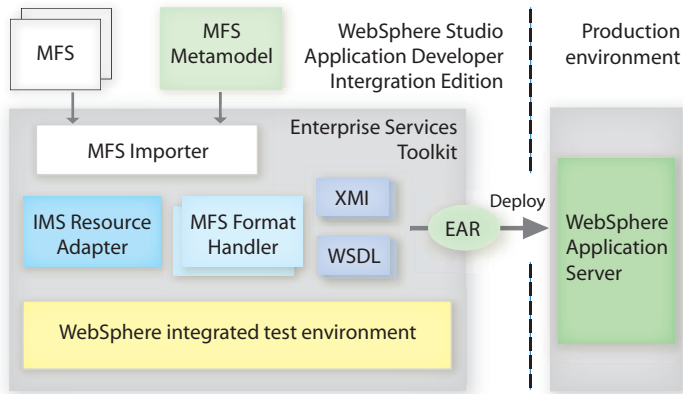


Figure 2

output byte arrays for the IMS host application. Supported features include function keys and data fields, as well as others. IMS MFS Web Services support IBM WebSphere Application Server Version 5 and Version 6 on both the Windows and z/OS platforms. To address new user requirements, we are exploring how we can expand the capabilities of MFS Web Services to support the new Service Component Architecture and Service Data Object programming model with Rational tooling.

Figure 2 shows how WSADIE generates and packages the MFS Web Service into an EAR that contains WSDL files, XMI metadata, java beans, and format handlers.

Figure 3 shows the runtime processing of an MFS Web Service, which can be invoked from an EJB proxy or a Java proxy.

(WSADIE) Version 5.1. The MFS Importer Wizard invokes the MFS Importer to parse MFS source files and generate the enterprise application archive (EAR), which includes the XMI metadata, WSDL files, java beans, and format handlers. The IMS MFS Web Services runtime component is packaged with IMS Connector for Java. The runtime component includes the MFS Adaptor and the MFS Format Handler, which work together with the input and output format handlers to perform data transformation and to generate and process input and

IMS MFS Web Enablement

The IMS MFS Web Enablement solution allows users to facilitate the operation of MFS-based IMS transactions from a Web browser. The MFS Web Enablement tooling component, the MFS XML Utility, is a stand-alone DOS-based application that invokes the MFS Importer to parse MFS source files and generate XML metadata and Web application archive (WAR) files. The runtime support component includes the MFS Servlet,

continued on page 11

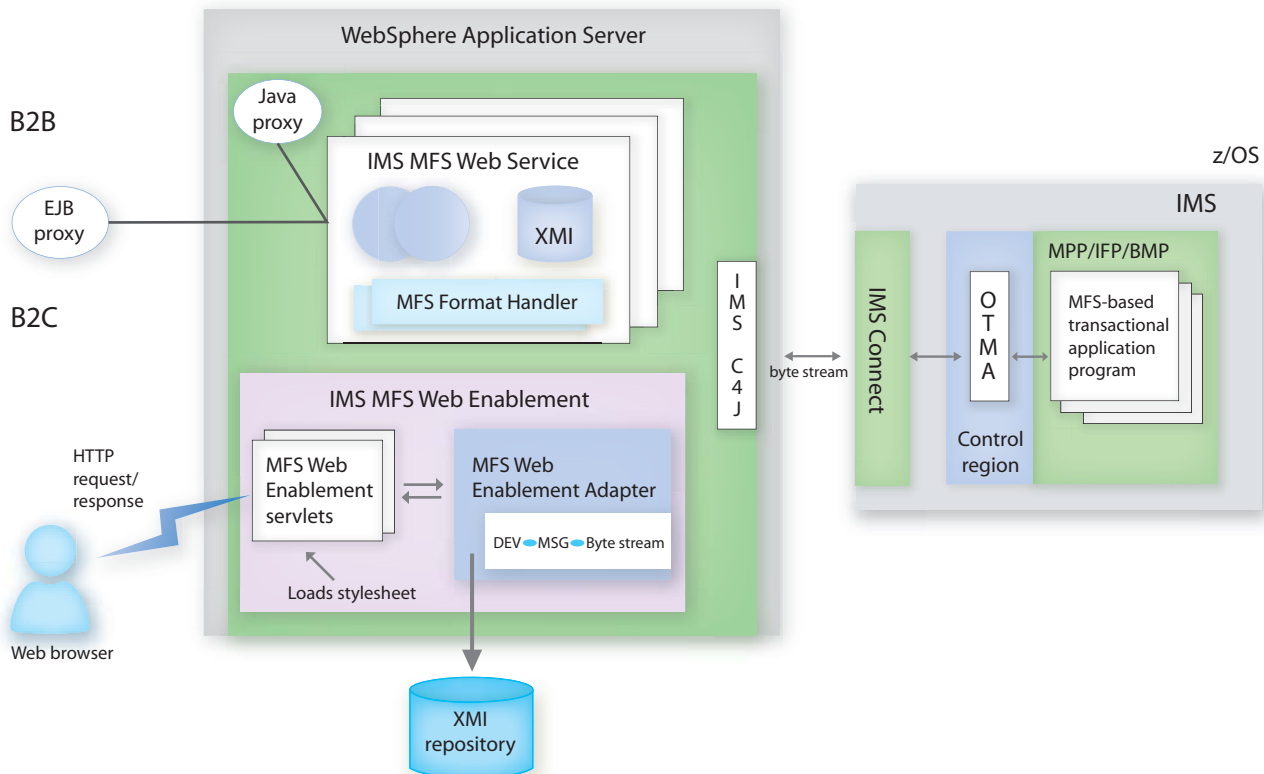


Figure 3



2005 IMS Technical Conference

October 10-13, 2005

San Jose Convention Center

San Jose, California

The 2005 IMS Technical Conference offers a one-of-a-kind, technical, skill-building experience that enables you to compete more effectively in serving your clients. Learn how the features of IMS Version 9 provide opportunities to tap emerging e-markets, optimize the performance of your current systems, raise the productivity of your personnel, effectively utilize new enhancements for application integration, and extend the reach of your on demand business environment into the world of services-oriented architecture. Learn how to quickly and easily adapt your IT infrastructure to today's changing conditions. At this event, you'll get your hands on the most advanced IMS tools available to help you push the boundaries of your enterprise.

Our keynote speakers, Jeanette Horan, Vice President, Worldwide Information Management Development and General Manager, IBM Silicon Valley Laboratory, and Greg Lotko, IMS Product Director, will give a joint address: *Information management: an integrated product portfolio for on demand business transformation.*

In-depth technical sessions

Choose from dozens of information-packed elective sessions. These sessions will help you stay current and are designed to increase your IMS skills, no matter what your present knowledge level. Develop knowledge in technical areas that have practical value in the workplace. The spectrum of topics include:

- IMS tools talk: featured session
- Basics and general interest
- Information Integration: connectivity and new application development
- Manageability
- Scalability in performance, capacity, availability, and recovery
- Hands-on sessions

Hands-on sessions

Learn by doing and judge the performance of new IMS tools and technology for yourself. Experience exciting features and functions personally. Stretch your mind and develop skills that you can start using the day you return to the office. Participate in hands-on sessions with the latest IMS technology facilitated by members of IBM's Advanced Technical Support Group and the Data Management Tools Group.



Photo courtesy San Jose Convention & Visitors Bureau

Attend these hands-on sessions to see, use, and be updated on what's new with IMS:

- Analysis solution: IMS Performance Analyzer, IMS Problem Investigator, and IMS Connect Extensions
- Database administration solution: IMS Parameter Manager
- TM management solution: IMS High Performance Sysgen Tools
- Administration solution: HALDB conversion and maintenance aid IMS connectivity, parts 1 and 2

Early-bird sessions included in the conference fee

Add an extra day's value to the conference by attending elective sessions on Sunday! Since you're coming to San Jose anyway, why not arrive on Saturday and start your conference experience early? You might even save on airfare by flying on Saturday and spending a day with the product experts.

Mix and match from the following sessions:

- IMS OTMA implementation, parts 1 and 2
- HALDB database administration, parts 1 and 2
- Using IPCS with IMS
- IMS exploitation of the parallel sysplex®, parts 1 and 2
- Debugging IMS abends with IPCS

IMS usability labs

Your feedback is essential to enhancing the usability of the IMS tools and functions that you already know and use. The IMS design team puts you in the driver seat. Help define the direction of IMS software by previewing what is coming soon from IBM. Tell us your preferences through group discussion, hands-on design explorations, user testing, and surveys.

Product EXPO

Count on the product EXPO for a firsthand look at the latest IMS technologies and solutions. Vendors will be on hand to provide demonstrations and answers to your questions about applications that can benefit your enterprise. Get new ideas and bring a list of your specific challenges. EXPO is a great planning tool for making future purchasing decisions.

Special post-conference workshop

Extend your stay to enhance your learning experience by attending a free post-conference workshop: IMS Version 9 installation and maintenance.

This workshop is being offered at no additional charge if you register for the conference and are one of the first 25 customers to sign up for the workshop. Your instructor is John Butterweck, IMS Technical Support Specialist, IBM Silicon Valley Lab. Special materials created for this class have been used to help workshop attendees through difficult parts of the system installation and migration processes.

The workshop is approximately seven hours in length and will run from 1:00 p.m. to 5:00 p.m. on Thursday, October 13 and Friday, October 14 from 8:30 to noon. It will be held at the IBM Silicon Valley Laboratory and transportation will be provided. Seats are limited, so sign up when you enroll for the conference! To enroll in the workshop, contact John Butterweck at jbutterw@us.ibm.com. Remember, you must sign up for the conference separately.

Mix and mingle with the IMS developers

Enjoy a unique opportunity to network with the the IMS developers who are participating in this key event in their home city of San Jose. This conference will have more developers in attendance than any prior IMS conference. Learn from the IMS developers themselves. Gain tips, tactics, insights, implementation strategies, and get solid advice.

San Jose: A city of fun, diversity, and innovation

San Jose has been named america's safest big city for the third consecutive year, and recently was named the third most fun city in the US. It offers visitors of all ages a variety of high-tech, historic, and fun activities. Whether it's magnificent museums, theme parks, night-life in the South First Area District or shopping at Santana Row, there is plenty to do for all. With 300 sunny days per year, you and your family can enjoy a wonderful time in San Jose, which is conveniently situated 45 minutes south of San Francisco and 1.5 hours north of Monterey. This city of cultural diversity, innovation, and entrepreneurial success will help make your IMS Technical Conference a most memorable 3.5 days you will long remember.

Need more information?

Do you need more information to convince your manager that you need to attend the conference? Check the Web site for

detailed session titles, speakers, and abstracts to justify your enrollment. On the Web site you'll find the detailed agenda, along with a spreadsheet, ready to download, so you can determine the optimal schedule for you while at the conference.

Don't miss all the excitement of this special event held this year in IMS developer country. Join us for skill building, problem solving, and networking this October in San Jose!

For more details and to enroll, visit the conference Web site at ibm.com/training/us/conf/ims

Your registration fee of \$1,895 is already paid if you have a U.S. IBM Education Card. If your company has a U.S. IBM Education Pack - online account, you can pay your registration fee electronically. For more information on these flexible IBM Education Advantage savings programs, please visit ibm.com/training/us/savings.

Amiet Goldman

Marketing Manager

IBM Global Services, supporting IBM IT Education Services

IMS MFS Web Support

continued from page 9

which sends and receives HTTP requests and responses, and the MFS Adapter. The output HTML page is dynamically rendered by the MFS Servlet using a cascading stylesheet. Sample stylesheets are provided and can be customized to fit specific business needs. Supported features of IMS MFS Web Enablement include IMS conversation, /FOR and /EXIT commands, attributes and extended attributes, function keys, PA1 keys, and others. IMS MFS Web Enablement is available as a free download for IMS V9 customers.

Figure 3 also shows the runtime processing of the MFS Web Enablement Servlet and the MFS Web Enablement Adapter. Data flows to and from a Web browser through IMS MFS Web Enablement runtime support to the IMS host application.

IMS MFS Web solutions are available as part of the IMS Integration Solutions Suite at ibm.com/ims.

Shyh-Mei Ho

IMS On Demand Integration Architecture

Jenny Hung

IMS Development



IMS High Performance Sysgen Tools for z/OS Version 2.1

If you hear that the latest release of the IMS High Performance Sysgen Tools for z/OS is much improved, you can be sure that it's not just marketing hype! The new features included in Version 2.1 mean that database administrators and system programmers have reliable, easy-to-use IMS management tools.

The Version 2.1 enhancements enable you to:

- Use the new ISPF interface to easily manage IMS system generation (sysgen) resources
- Use the new ISPF resource update list editor to quickly and easily define updates to the IMS system generation resources
- View IMS system generation resource definition attributes
- Recreate IMS system generation and security generation source from the MODBLKS and MATRIX data sets
- Review resource changes that have been made, and back out selected resource changes

Managing resources using the new ISPF interface

The IMS High Performance Sysgen Tools includes a new ISPF interface. As an integrated component, the ISPF interface makes it easier for you to manage your IMS system generation resources by allowing you to quickly review and change your application resources definitions.



Using the new ISPF interface, you can display current IMS resource definitions, create incremental system generation changes (resource update list), install a resource update list, check the syntax of your IMS stage 1 system generation source, perform a fast IMS system generation (MODBLKS system generation), recreate IMS stage1 source or security generation source from your MODBLKS or MATRIX data sets, view a historical log of HP system generation activity, and issue an IMS command and view the command output.

Resource update lists

The new resource update list included in Version 2.1 enables you to define a set of IMS system definition changes (for example, update, add, or delete databases, programs, transactions, or route codes), verify that the changes are compatible with a specific IMS subsystem's configuration, and install the changes in an IMS subsystem.

Figure 1 shows how resource update list entries are maintained.

```

VIEW      IMS HP Sysgen Tools - Program List      Row 2,842 to 2,857 of 9,904
Command ==> _____ Scroll ==> CSR

Primary Commands:          Line Commands:
SORT Sort the List         S View a Program Definition
L  Locate an Entry         T View Associated Transactions

  PSBNAME  RESIDENT  DOPT  PGMTYPE  SCHDTYP  GPSB  LANGUAGE  FPATH
  -----  -
  DBB4011A          BATCH  SERIAL
  DFSIVPA          BATCH  SERIAL
  DFSIVPB          BATCH  SERIAL
  DFSIVPC          BATCH  SERIAL
  DFSIVP1           TP     SERIAL
  DFSIVP2           TP     SERIAL
  DFSIVP3           TP     SERIAL
  DFSIVP4           TP     SERIAL
  DFSIVP5  RESIDENT  TP     SERIAL
  DFSIVP6          BATCH  SERIAL
  DFSIVP7          BATCH  SERIAL
  DFSIVP8          BATCH  SERIAL
  DFSIVP9          BATCH  SERIAL
  DFSSAM02         TP     PARALLEL
  DFSSAM03         TP     PARALLEL
  DFSSAM04         TP     SERIAL
    
```

To create an entry for a resource update list, specify the function (add, delete, or update) and the resource type (database, program, transaction, or route code), and a resource name. After you specify these items, as shown in figure 2, you must specify values for all the attributes of the resource definition. For example, when updating a transaction, the current attribute values are displayed, along with the desired attributes. You can copy the attributes of an existing transaction using the COPY command, which replaces the desired attribute values with the values that are associated with the copied transaction, or you can simply update the attribute values.

Figure 1

When editing the resource update list entry attribute values, the IMS High Performance Sysgen Tool checks to ensure that your specifications are valid IMS values and not values valid for any particular IMS subsystem. You can verify the validity of the resource update list after it is created using the Verify option, which allows you to verify the attribute values for all resources in the resource update list against a specified IMSID. Inconsistent or invalid attribute values will be shown at the terminal, which allows you to edit the resource update list to correct any problems.

After you create and verify a resource update list, you can install it by selecting the list name and specifying the IMSID where the changes will

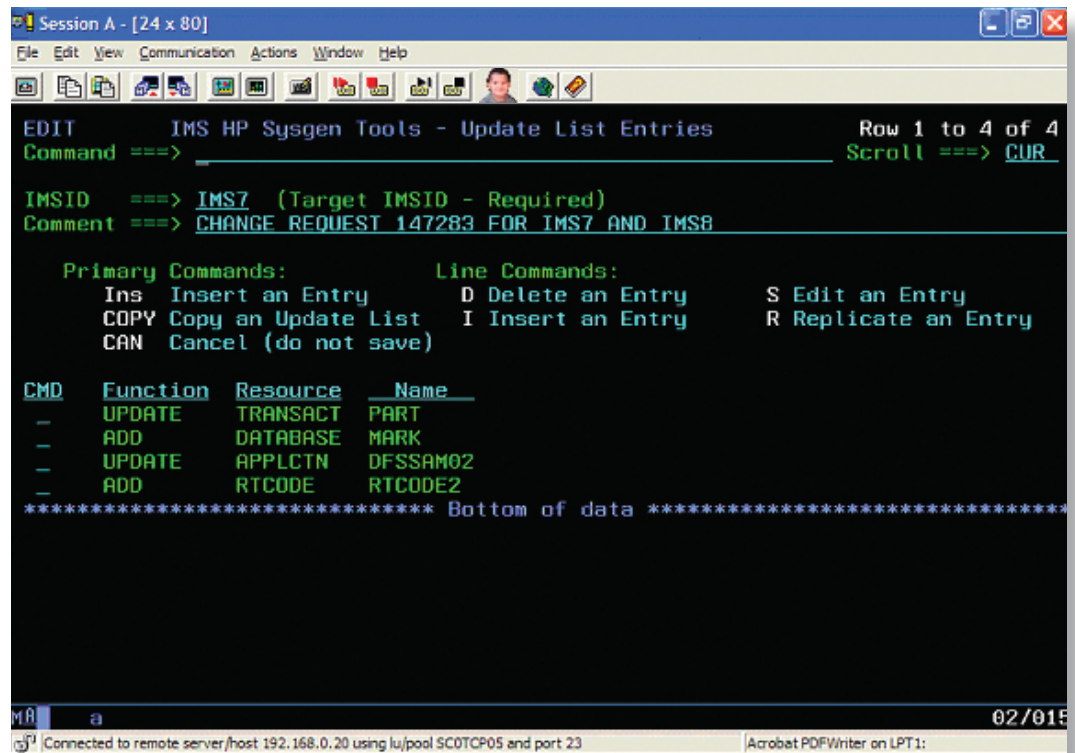


Figure 2

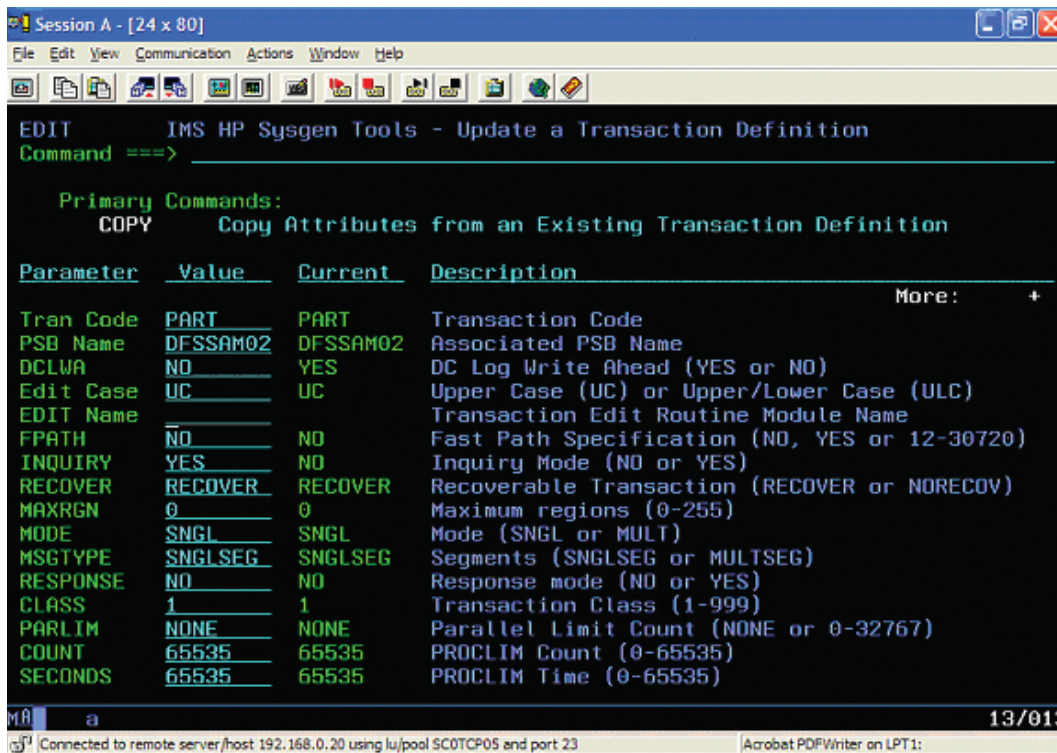


Figure 3

be implemented. The IMS High Performance Sysgen Tool then updates the IMS control blocks and the MODBLKS and MATRIX data sets to reflect the specified changes. Because the MODBLKS and MATRIX data sets are updated when the resource update list is installed, changes are carried across IMS restarts and cold starts. IMS system generation source is not used during this process. The IMS High Performance Sysgen Tool uses the internal IMS control blocks, MODBLKS data sets, and MATRIX data sets during the installation process.

continued on page 14

IMS High Performance Sysgen Tools for z/OS Version 2.1

continued from page 13

Viewing IMS system generation resource definitions

The new VIEW command lets you view IMS resource definitions from either the active MODBLKS data set or the internal IMS control blocks and displays the resources in an ISPF table. Figure 3 shows the new table format where you can sort entries by any attribute. For example, you can find all Fast Path PSBs by entering the command SORT FPATH on the command line. You also can display a specific PSB, which shows the source IMS sysgen macro that will create the resource definition. When viewing transactions, route codes, or programs, you can also display all of the resources that are associated with the selected resource, including the IMS system generation definition for the PSB, and all transactions and route codes that are associated with that PSB.

Backing out of a system generation update

The IMS High Performance Sysgen Tool provides a way to back out of incorrect system generation updates using the history log. The history log tracks all updates applied to each IMS subsystem. You can review the history log entries, which show installation details such as the date, time, and user ID that requested the installation, and the before and after resource attributes associated with the log entry.

To back out of one or more entries in the history log, simply select the entries that you want to back out. The IMS High Performance Sysgen Tool creates a resource update list that will reverse the effects of the history log entries. To undo the selected changes, install the generated resource update list and your selected changes will be restored.

You can also use the history log to maintain your IMS system generation source by generating IMS system generation source macros from the history log entries so that your source matches your online environment. The history log also provides a facility to document which log entries have been integrated into your IMS sysgen source, which makes maintaining the IMS system generation source easier.

Recreating IMS stage 1 and security generation source

The recreation of IMS stage 1 and security generation source code is critical when you need to run an IMS system generation or investigate the source of definition change problems. IMS High Performance System Generation Version 2.1 includes a new a reverse system generation and MATRIX function that allows you to generate IMS system generation source or security generation source that matches sysgen output data sets for the currently executing IMS environment. In addition, you can use the reverse system generation and MATRIX function to create new IMS sysgen source or to create source from an existing MODBLKS data set.

To recreate IMS stage 1 source, you can select from three kinds of data sources: the active MODBLKS data set, a user-specified MODBLKS data set, or the internal control blocks for a running IMS region. From these data sources, you can choose to generate source for either the databases, or the programs, transactions, and route codes.

To recreate security generation source, the reverse system generation and MATRIX function allows the same flexibility for selection of the data sources, which allows you to create IMS security generation source for an existing backup MATRIX data set.

For more information see *IBM IMS Version 2.1 High Performance System Generation Tools for z/ OS*.

Janet LeBlanc

IMS Database Tools Sales and Migrations

The IMS Newsletter Fall 2005

Editor-in-Chief Glenn N. Galler
Creative Director Robert Lee
Graphic Designer Kelly Lisai
Managing Editor Lara Hamilton Cone
Consulting Editor Moira McFadden Lanyi
Contributing Editor Sandy Sherrill
Contributing Author Joel Hamill
Senior Web Consultant Michael Morrison


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

Or contact us in any of the following ways:

Email: imsmt@us.ibm.com
Fax: (408) 463-4101
US Mail: IMS Newsletter
IBM Corporation
DQY/A2
555 Bailey Avenue
San Jose, CA 95141-1003



Hot Websites



Information Management Software for
z/OS Solutions Information Center
<http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/index.jsp>

IBM Redbooks™
ibm.com/redbooks/

IBM Software for z/Series®
ibm.com/software/zseries/

IMS
ibm.com/ims/

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

IMS Examples Exchange
ibm.com/software/data/ims/examples/exHome.html

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

IMS Roadmap
ibm.com/software/data/ims/roadmap/

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

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

IMS Tools Library page:
ibm.com/software/data/db2imstools/imstools-library.html

DB2 and IMS Tools Support page:
ibm.com/software/data/db2imstools/support.html

New Book: *An Introduction to IMS*
www.amazon.com/exec/obidos/tg/detail/-/0131856715/103-4160926-2088609?v=glance



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





They must be in here somewhere...

Once again you can find my children's pictures hidden somewhere in this issue. We will send an IBM IMS button to every person who can correctly find the two pictures of my children in this issue of the IMS Newsletter.

If you would like to receive an IBM IMS button, please e-mail your external postal address and the picture locations of my children to: ismkt@us.ibm.com

Or write to:

Editor-in-Chief, IMS Newsletter
IBM Corporation
DQY/A2
555 Bailey Avenue
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

© International Business Machines Corporation
Printed in the U.S.A. 9/05
All Rights Reserved

DB2, IBM, IMS, MVS, OMEGAMON, OS/390, Parallel Sysplex, RACF, Redbooks, Tivoli, VTAM, WebSphere, z/OS, and zSeries are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

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

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

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

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

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

GC26-9655-17



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

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