



IBM Software Group




IMS Version 11

Feature Overview
August 2009

 Alison Coughtrie
alison_coughtrie@uk.ibm.com



Information Management software



© Copyright IBM Corporation 2009



Important Disclaimer

© Copyright IBM Corporation 2009. All rights reserved.

U.S. Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS AND/OR SOFTWARE.

IBM, the IBM logo, ibm.com, and IMS are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml

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




IMS 11 Highlights

Database Manager	System	Transaction Manager	DBRC
<ul style="list-style-type: none">IMS Open DatabaseDatabase QuiesceACBLIB UsabilityDatabase RASOLR PerformanceFast Path 64-bit Buffer ManagerFast Path Usability	<ul style="list-style-type: none">IMS Connect EnhancementsUser Exit InterfaceDump FormatterSyntax Checker & IVPLSQA Storage ReductionKBLA/DIAGNOSE Command	<ul style="list-style-type: none">Type-2 Query TM CommandsOTMA EnhancementsOTMA Type-2 CommandsTransaction ExpirationShared Queues Affinity Routing	<ul style="list-style-type: none">BPE-Based DBRCSecurity Override for Non-Production RECONUnconditional deletion of PRILOG InformationDBRC Migration/Coexistence from IMS 9 and 10

3 | | IMS 11 Overview | August 2009 | © Copyright IBM Corporation 2009




IBM Software Group

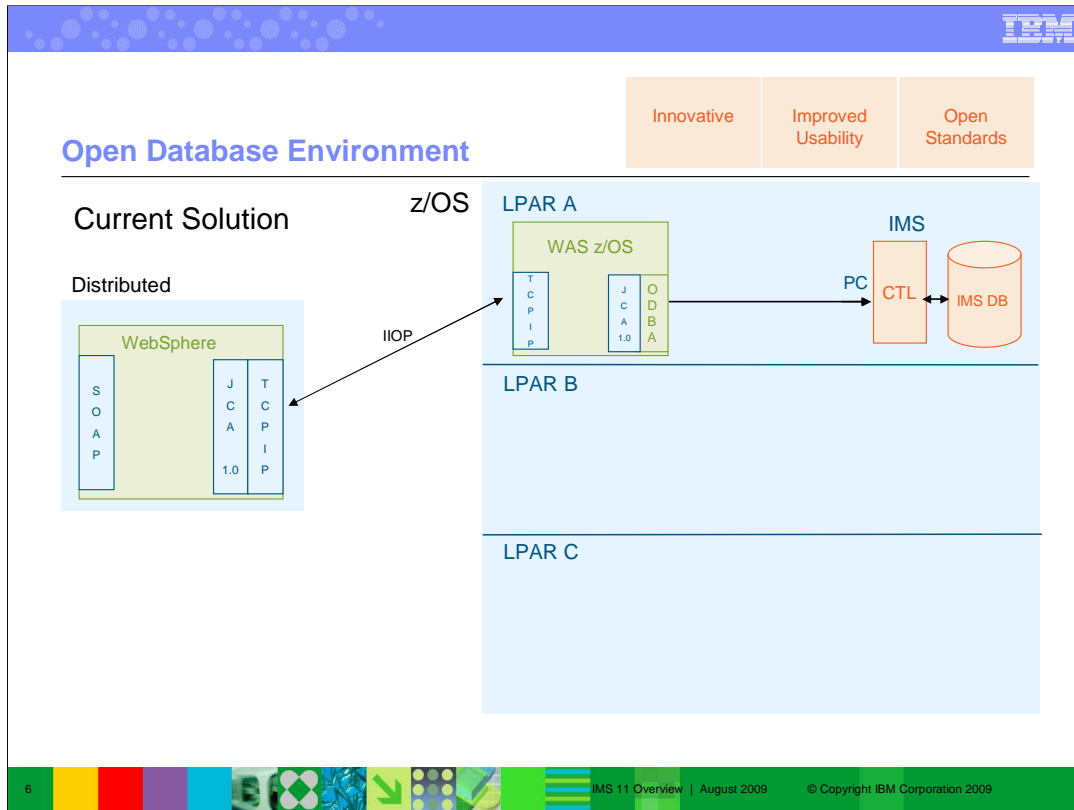


Database Enhancements

Information Management software

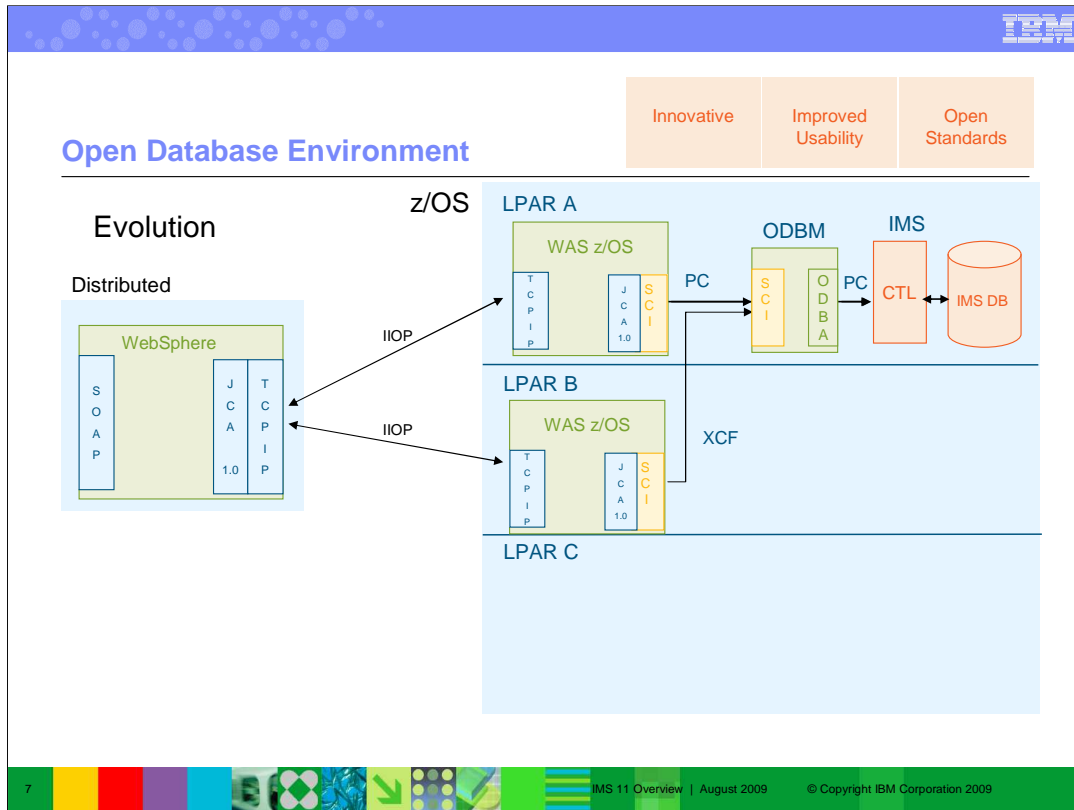


© Copyright IBM Corporation 2009

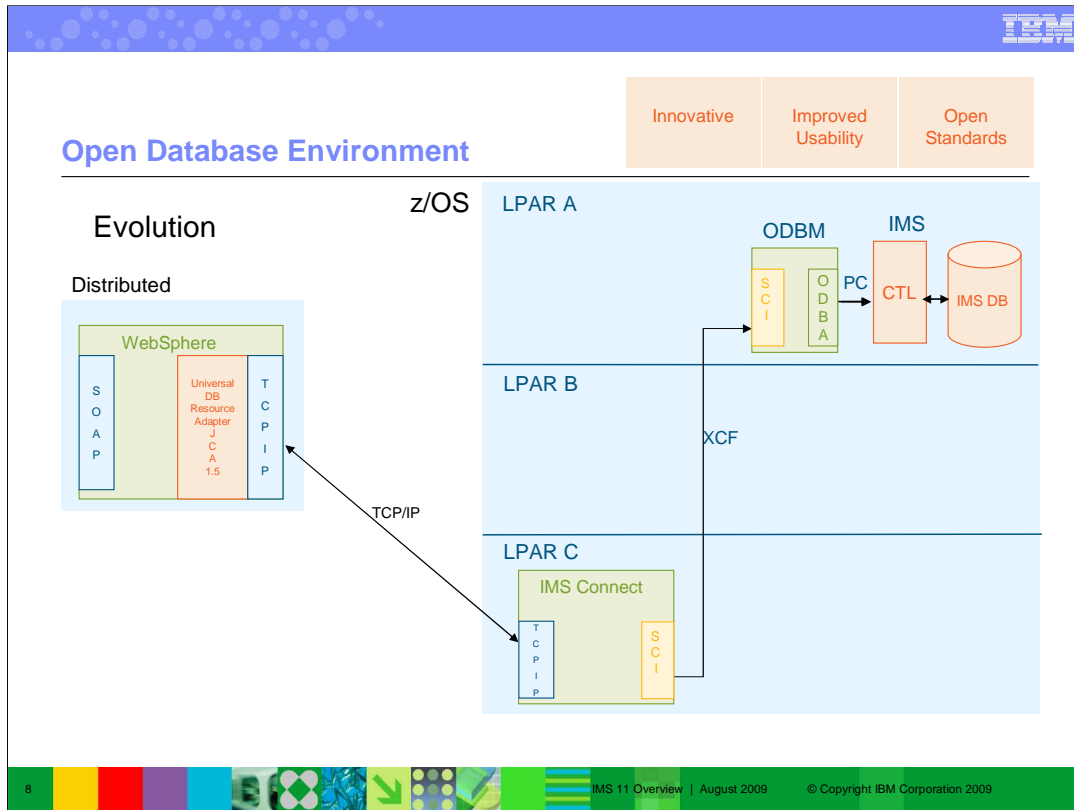


The intent of the following 4 charts is to show the current topology and illustrate the evolution to the new one, pointing out the enhancements at each step. As a point of fact, WAS z/OS cannot take advantage of ODBM's cross-LPAR feature unless WAS itself embraces SCI. Applications can use the out-of-the-box compatibility mode to use AERTDLI and have those calls routed to an ODBM which will still prevent the U113 abend – but WAS and the ODBM address space will still need to be on the same LPAR. It is just an illustrative example showing what can be possible with WAS z/OS as an ODBM client.

The current solution (whether or not we are talking about distributed or local access to IMS DB) leverages ODBA as the API to access IMS database resources. ODBA is capable of making address space to address space calls (PC calls) in the same logical partition. The net effect of this is that the ODBA modules need to be on the same LPAR as the IMS CTL region. These modules (ODBA) are loaded in the address space of the application, which is in turn loaded in the address space of the container. In this case the container is WebSphere AS. The result of this is that the WAS installation has to be on the same LPAR as the IMS DB itself. There is no isolation.

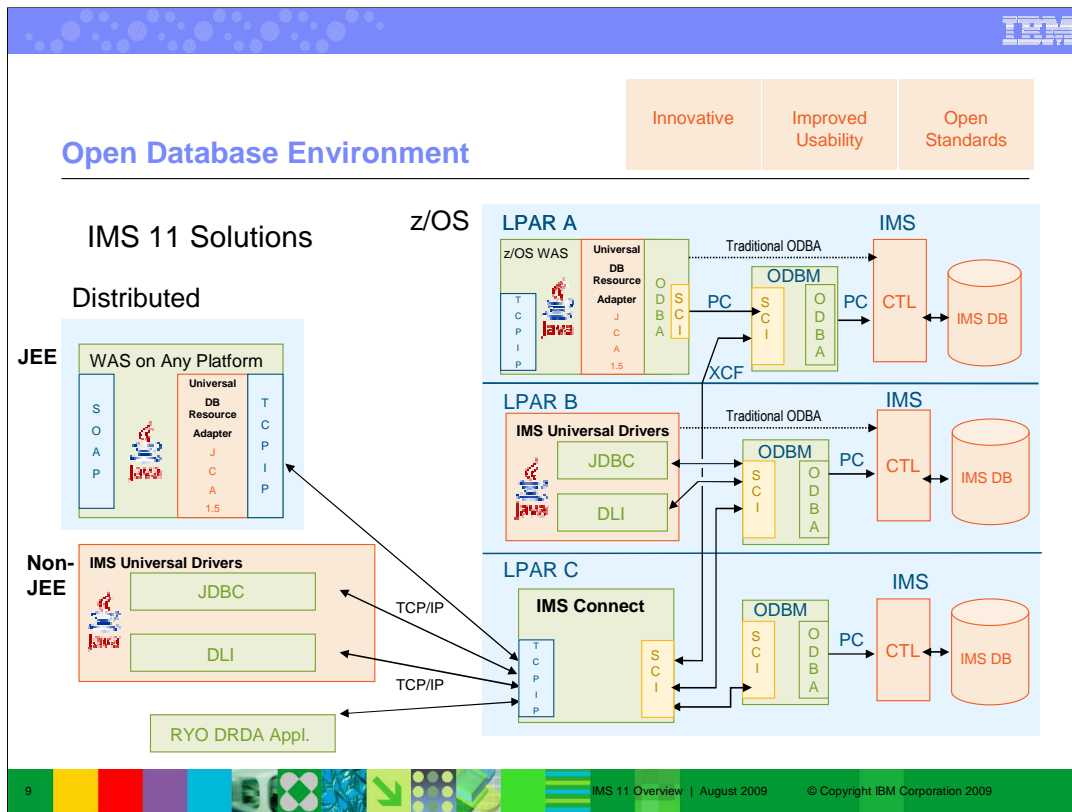


By leveraging SCI, the applications can be on any LPAR in an IMSplex. SCI uses either PC or XCF calls to communicate with other SCI components. XCF allows calls to go across LPARs in an IMSplex. This allows applications (and their containers) to be isolated on their own LPARs.



What we are doing is creating a new CSL address space to house the ODBA modules. This interface will use SCI as its communication mechanism. The ODBA modules are no longer tightly coupled with the applications themselves (and therefore the containers).

By leveraging SCI, the applications can be on any LPAR in an IMSplex. SCI uses either PC or XCF calls to communicate with other SCI components. XCF allows calls to go across LPARs in an IMSplex. This allows applications (and their containers) to be isolated on their own LPARs.



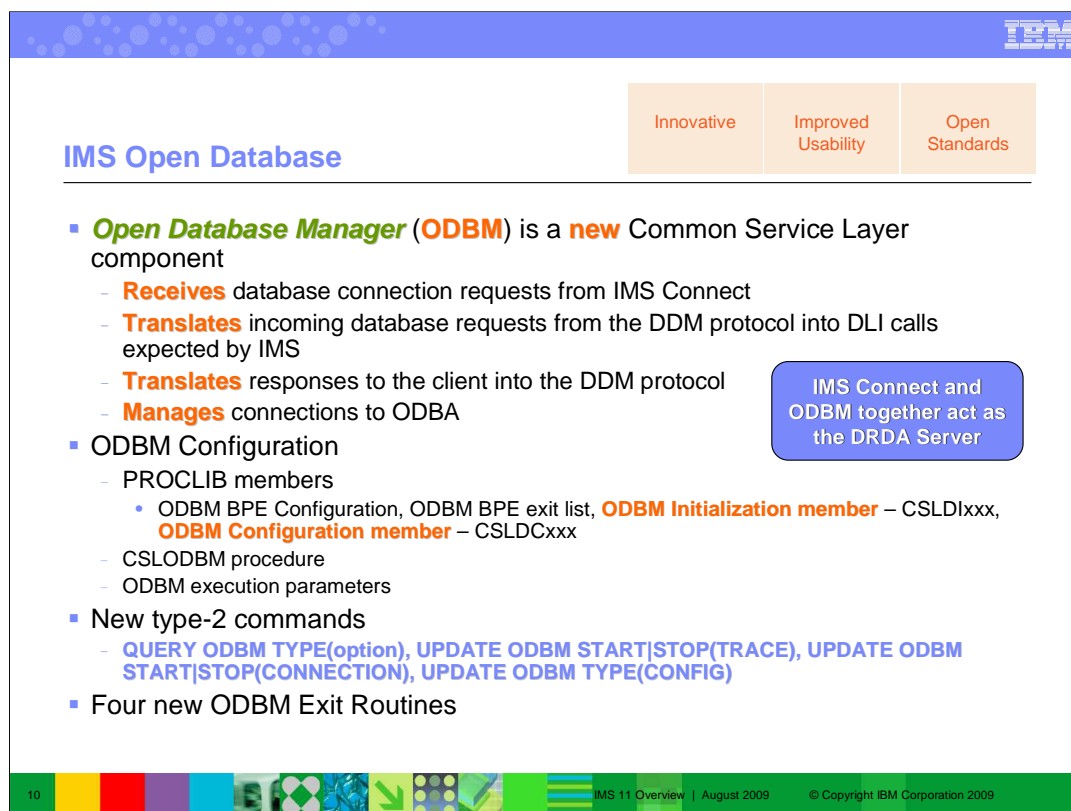
This leads us to our real goal, which is to leverage IMS Connect as the complete gateway solution for IMS TM, OM, and now DB. IMS Connect will be augmented to be an ODBM client. This will allow distributed applications to leverage the TCP/IP protocol to communicate with IMS Connect, which can then access any database in the entire IMSplex.

IMS Connect becomes the IMS Gateway to both IMS TM and IMS DB.

WebSphere and DB2 Stored Procedures no longer have to be on the same LPAR with IMS when they interface with the IMS ODBM (Open Database Manager) address space. The ODBM address space must be on the same LPAR with IMS due to the use of the ODBA (Open Database Access) interface.

Distributed clients would now have the option of going directly to IMS Connect for IMS DB requests.

Existing DB Resource Adapter applications are unaffected by Open Database. In order to **exploit** Open Database from existing DB Resource Adapter applications, a migration to the JCA 1.5 programming model would have to be done.



Innovative
Improved Usability
Open Standards

IMS Open Database

- **Open Database Manager (ODBM)** is a **new** Common Service Layer component
 - **Receives** database connection requests from IMS Connect
 - **Translates** incoming database requests from the DDM protocol into DLI calls expected by IMS
 - **Translates** responses to the client into the DDM protocol
 - **Manages** connections to ODBA
- ODBM Configuration
 - PROCLIB members
 - ODBM BPE Configuration, ODBM BPE exit list, **ODBM Initialization member** – CSLDIxxx, **ODBM Configuration member** – CSLDCxxx
 - CSLODBM procedure
 - ODBM execution parameters
- New type-2 commands
 - **QUERY ODBM TYPE(option), UPDATE ODBM START|STOP(TRACE), UPDATE ODBM START|STOP(CONNECTION), UPDATE ODBM TYPE(CONFIG)**
- Four new ODBM Exit Routines

IMS Connect and ODBM together act as the DRDA Server

10 | IMS 11 Overview | August 2009 | © Copyright IBM Corporation 2009

One ODBM instance must be defined in the IMSplex to use ODBM functions. Each z/OS image can have more than one ODBM. If multiple instances of ODBM are defined in the IMSplex, any ODBM instance can perform work from any z/OS image in the IMSplex.

ODBM routes the database connection requests received from IMS Connect to the IMS systems that are managing the requested database. Before establishing the connection to the IMS system, ODBM translates the incoming database requests from the DDM protocol submitted by the IMS-provided connectors and user-written DRDA applications into the DL/I calls expected by IMS. When ODBM returns the IMS output to the client, ODBM translates the response to the DDM protocol. From the ODBM perspective, application programs that interact directly with ODBM, such as IMS Connect, are ODBM clients. Users can create their own ODBM clients by using the new ODBM CSLDMI API. ODBM client application programs can access databases that are managed by IMS DB on any LPAR in an IMSplex.

Use the CSLDIxxx IMS.PROCLIB member to specify parameters that initialize the ODBM address space. Certain parameters within CSLDIxxx can be overridden with the ODBM execution parameters.

Use the CSLDCxxx IMS.PROCLIB member to configure the ODBM datastore connections to IMS systems.

The following IMS exit routines are new:

- CSL ODBM Initialization and Termination user exit routine - Called when an ODBM address space or IMSplex initializes or terminates.
- CSL ODBM Input user exit routine - Called when ODBM receives a call to the ODBM callable interface.
- CSL ODBM Output user exit routine - Called when ODBM returns the results from a call to the ODBM callable interface.
- CSL ODBM Client Connect and Disconnect user exit routine - Called when a client registers to or de-registers from ODBM.



IMS Open Database – IMS Connect

Innovative Improved Usability Open Standards

- **IMS Connect** has the following enhancements in support of IMS Open Database:
 - IMS Connect Configuration member HWSCFGxx
 - New **ODACCESS** statement
 - DRDA ports, timeout value, IMSplex name etc.
 - Changes to existing **commands**
 - VIEWHWS, VIEWDS, VIEWPORT
 - **New Commands**
 - STARTOD, STOPOD, STARTIA, STOPIA, VIEWIA, SETOAUTO
 - **New User Exits**
 - HWSROUT0 – Routing Exit for ODBM
 - can override the IMS alias and/or select the ODBM target
 - HWSAUTH0 – Security Exit for ODBM
 - can perform the authentication of the userid

11 IMS 11 Overview | August 2009 © Copyright IBM Corporation 2009

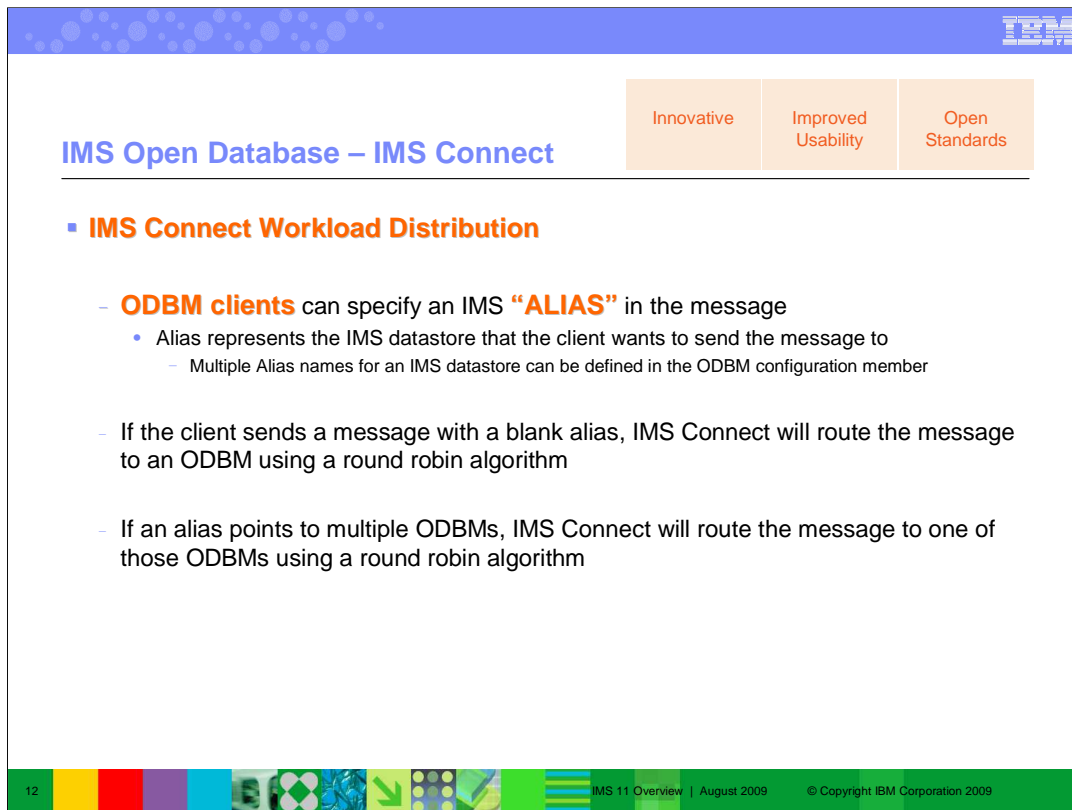
IMS Connect, which previously provided access to only IMS Transaction Manager (IMS TM), through Open Transaction Manager Access (OTMA), now also provides access to IMS Database Manager (IMS DB), through ODBM. IMS Connect can also provide workload distribution by routing database connection requests to ODBM based on an alias name that is submitted by the client application program. To the client application, the alias name represents the IMS system, or data store, to which the application program connects. Depending on the value of the alias name submitted, IMS Connect either routes the incoming connection request to a specific ODBM instance or distributes the incoming connection request to any available instance of ODBM in an IMSplex.

IMS Connect will invoke the routing user exit first to allow the exit to select an ODBM and/or override the IMS ALIAS.

- If the routing exit selects an ODBM, IMS Connect will use that ODBM and will not perform the round robin routing method.
- If the routing exit does not select an ODBM, IMS Connect will select an ODBM and perform the round robin routing method based upon the explicit ALIAS name as well as the blanked ALIAS name.
- If the routing exit overrides the IMS ALIAS, IMS Connect will use that IMS ALIAS.
- IMS Connect will validate the ALIAS and the ODBM upon returning from the exit.

IMS Connect will always call the new security user exit HWSAUTH0 independently of the RACF= parameter in the IMS Connect configuration member.

- IMS Connect calls HWSAUTH0 before invoking any installed security facility such as RACF.
- HWSAUTH0 will perform the authentication of the userid/passticket of the ODBM client.
- HWSAUTH0 can override the input userid with a different userid and to provides a RACF groupid to be authenticated further by IMS Connect. HWSAUTH0 can be refreshable.
- Password is passed in the clear.



The slide features a blue header with a pattern of white dots. Below the header, the title "IMS Open Database – IMS Connect" is displayed in blue. To the right of the title are three orange boxes containing the words "Innovative", "Improved Usability", and "Open Standards". The main content area contains a bulleted list under the heading "IMS Connect Workload Distribution". The list includes three main points: ODBM clients specifying an alias, routing with a blank alias using a round robin algorithm, and routing with a specific alias using a round robin algorithm. The footer of the slide is green and contains the number "12", a series of colorful icons, and the text "IMS 11 Overview | August 2009 © Copyright IBM Corporation 2009".

IMS Open Database – IMS Connect

Innovative Improved Usability Open Standards

- **IMS Connect Workload Distribution**
 - **ODBM clients** can specify an IMS **"ALIAS"** in the message
 - Alias represents the IMS datastore that the client wants to send the message to
 - Multiple Alias names for an IMS datastore can be defined in the ODBM configuration member
 - If the client sends a message with a blank alias, IMS Connect will route the message to an ODBM using a round robin algorithm
 - If an alias points to multiple ODBMs, IMS Connect will route the message to one of those ODBMs using a round robin algorithm


12 IMS 11 Overview | August 2009 © Copyright IBM Corporation 2009


IMS Connect can also provide workload distribution by routing database connection requests to ODBM based on an alias name that is submitted by the client application program. To the client application, the alias name represents the IMS system, or data store, to which the application program connects. Depending on the value of the alias name submitted, IMS Connect either routes the incoming connection request to a specific ODBM instance or distributes the incoming connection request to any available instance of ODBM in an IMSplex.

Innovative
Improved Usability
Open Standards

IMS Open Database

- Comparison of programming approaches for accessing IMS data:

Application Platform	Data Access Method	Transaction Processing Required	Recommended Approach
	CCI programming interface to perform SQL or DL/I data operations.	Local transaction processing only.	Use the IMS Universal DB resource adapter with local transaction support (imsudbLocal.rar), and make SQL calls with the SQLInteractionSpec class or DL/I calls with the DLIInteractionSpec class.
		Two-phase (XA) commit processing or local transaction processing.	Use the IMS Universal DB resource adapter with XA transaction support (imsudbXA.rar), and make SQL calls with the SQLInteractionSpec class or DL/I calls with the DLIInteractionSpec class.
	JDBC programming interface to perform SQL data operations.	Local transaction processing only.	Use the IMS Universal JCA/JDBC driver version of the IMS Universal DB resource adapter with local transaction support (imsudbJLocal.rar), and make SQL calls with the JDBC API.
		Two-phase (XA) commit processing or local transaction processing.	Use the IMS Universal JCA/JDBC driver version of the IMS Universal DB resource adapter with XA transaction support (imsudbJXA.rar), and make SQL calls with the JDBC API.

13

IMS 11 Overview | August 2009
© Copyright IBM Corporation 2009

We are providing different API layers that leverage the DRDA protocol. The DLI Java API is an API that mimics the traditional DLI API plus a few more advanced concepts. We are also offering XA capabilities for applications leveraging our framework. (Definition of XA is below.)

For Open Database, we are offering a type 4 JDBC driver built on top of this framework. Ultimately we will offer a J2EE (Java Enterprise Edition) platform solution which is packaged as a JCA (Java Connection Architecture) 1.5 resource adapter for deployment in WebSphere Application Server (which is a J2EE server). (Editorial note: Sun dropped the '2' so it is really now JEE and not J2EE but most folks still keep the '2'.) **This JCA resource adapter will leverage all of the robust features of the J2EE platform** (connection management, transaction management, security management, XA support, logging, etc.) and will allow for direct distributed access to IMS database resources.

These interfaces will provide a base set of APIs that future clients (such as XQuery) can be built on top of (we have an XQuery implementation today, but now with ODB these APIs will be distributed).

The XA standard is an X/Open specification for distributed transaction processing (DTP). It describes the interface between the global transaction manager and the local resource manager. We are the resource manager (from JDBC driver through IMS connect, ODBM and IMS).

Innovative
Improved Usability
Open Standards

IMS Open Database

- Comparison of programming approaches for accessing IMS data:

Application Platform	Data Access Method	Transaction Processing Required	Recommended Approach
Standalone Java application (outside a Java EE application server) that resides on a distributed platform or a z/OS platform 	JDBC programming interface to perform SQL data operations.	Two-phase (XA) commit processing or local transaction processing.	Use the IMS Universal JDBC driver (imsudb.jar), and make SQL calls with the JDBC API.
	Traditional DL/I programming semantics to perform data operations.	Two-phase (XA) commit processing or local transaction processing.	Use the IMS Universal DL/I driver (imsudb.jar), and make DL/I calls with the PCB class.
Standalone non-Java application that resides on a distributed platform or a z/OS platform	Data access using DRDA protocol.	Two-phase commit processing or local transaction processing.	Use a programming language of your choice to issue DDM commands to IMS Connect. The application programmer is responsible for implementing the two-phase commit mechanism.

14
IMS 11 Overview | August 2009
© Copyright IBM Corporation 2009

IMS Open Database

- Distributed Syncpoint (global transaction) requires RRS on z/OS
- Use of RRS with ODBM is optional
 - RRS=Y|N parm for ODBM start-up
 - If RRS=Y (also the default), ODBM will use the ODBA interface (i.e. AERTDLI)
 - If RRS=N, ODBM will use the DRA interface like CICS
 - Global transactions are not supported if RRS=N)

Innovative

Improved Usability

Open Standards

15
IMS 11 Overview | August 2009
© Copyright IBM Corporation 2009

RRS=Y must be specified in the configuration for a global transaction.

This slide shows how Open Database achieves cross LPAR transaction management.

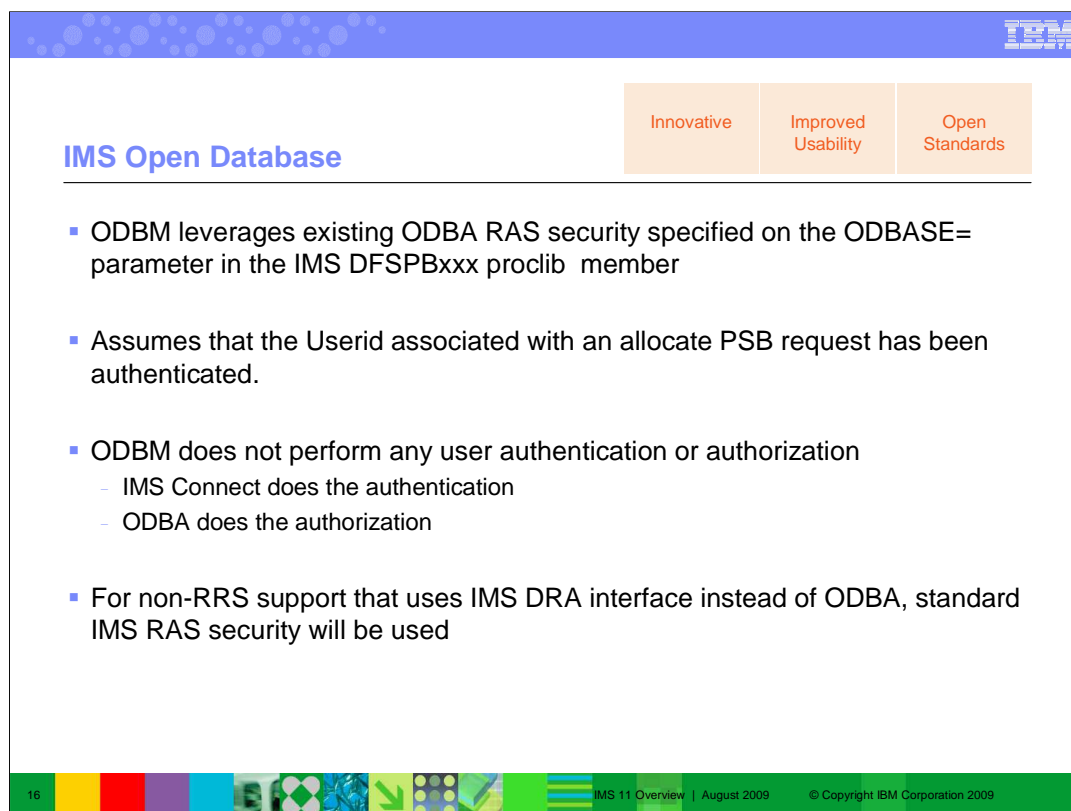
When a client establishes a connection through IIMS Connect to ODBM several things are done to establish a coordinated Unit of Work.

First IMS Connect creates the parent Unit of Recovery.

IMS Connect then sends the UR token to ODBM.

ODBM then expresses interest in the UR as a child.

At this point we have a coordinated Unit of Recovery established.



The slide features a blue header with a pattern of white dots. Below the header, the title "IMS Open Database" is displayed in blue. To the right of the title are three orange boxes containing the words "Innovative", "Improved Usability", and "Open Standards". The main content area contains a bulleted list of four items. At the bottom of the slide, there is a green footer bar with a decorative pattern of colorful squares and icons on the left, and the text "16 IMS 11 Overview | August 2009 © Copyright IBM Corporation 2009" on the right.

IMS Open Database

- ODBM leverages existing ODBA RAS security specified on the ODBASE= parameter in the IMS DFSPBxxx proclib member
- Assumes that the Userid associated with an allocate PSB request has been authenticated.
- ODBM does not perform any user authentication or authorization
 - IMS Connect does the authentication
 - ODBA does the authorization
- For non-RRS support that uses IMS DRA interface instead of ODBA, standard IMS RAS security will be used

For each request to access IMS data, connection information on the IMS host system, port number, and a valid user ID and password must be supplied in order to establish communication with IMS. A socket connection is first established to connect to the host IMS Connect system. When IMS Connect receives the request, it proceeds to authenticate the user based on the supplied user ID and password. After successful authentication, necessary information on the socket, such as PSB name and IMS alias (database subsystem) is sent to ODBA in order to allocate the PSB to connect to the database. An actual connection to an IMS database is only established when a PSB is allocated. Authorization for a particular PSB is done by the ODBA component during the allocation of a PSB.



Other Java Enhancements

Innovative

Improved Usability

Open Standards

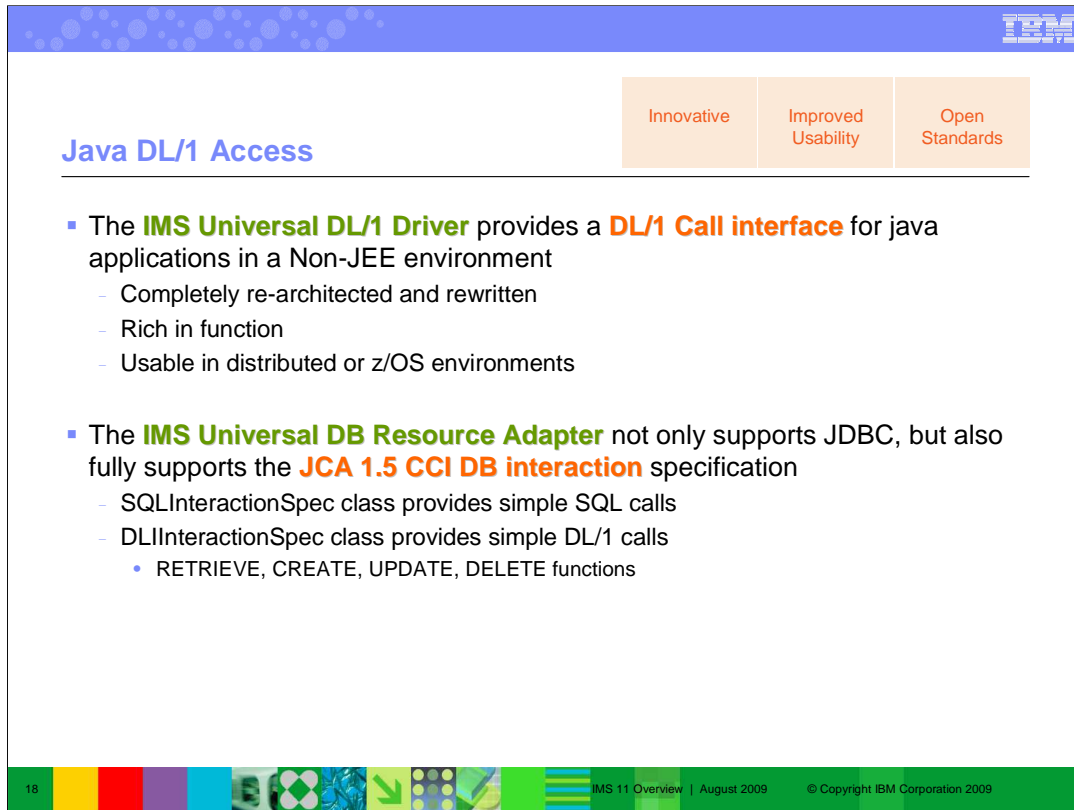
- The two Universal drivers for JDBC - **IMS Universal DB Resource Adapter** and **IMS Universal JDBC Driver** - offer a greatly enhanced **JDBC implementation**
 - JDBC 3.0
 - Keys of parent segments are included in table as foreign keys
 - Allows Standard SQL implementation (including for INSERT)
 - Updatable result sets
 - Metadata discovery API implementation
 - Uses metadata generated by DLIModel Utility as “catalog data”
 - Enables JDBC tooling to work with IMS DBs just as they do with DB2 DBs

- IMS 11 Java requires JDK 5.0 or greater
 - JMP and JBP online regions require JDK 6.0
 - JDK 6.0 offers significantly better performance



17

IMS 11 Overview | August 2009
© Copyright IBM Corporation 2009



Java DL/1 Access

Innovative Improved Usability Open Standards

- The **IMS Universal DL/1 Driver** provides a **DL/1 Call interface** for java applications in a Non-JEE environment
 - Completely re-architected and rewritten
 - Rich in function
 - Usable in distributed or z/OS environments
- The **IMS Universal DB Resource Adapter** not only supports JDBC, but also fully supports the **JCA 1.5 CCI DB interaction** specification
 - SQLInteractionSpec class provides simple SQL calls
 - DLInteractionSpec class provides simple DL/1 calls
 - RETRIEVE, CREATE, UPDATE, DELETE functions

18 IMS 11 Overview | August 2009 © Copyright IBM Corporation 2009

Improved Usability
Improved Availability
Reduced Complexity

ODBA Enhancements

- New **CIMS CONNECT** call which allows ODBA applications to connect to **multiple IMS subsystems** with a **single call**
 - Previously CIMS INIT was the only function available, which only allowed the caller to connect to a single IMS
 - Multiple CIMS INIT calls were required to connect to multiple IMS DB systems

- ODBA application programs can use the **ODBM address space** to manage the ODBA interface
 - Gives **protection from potential U113 abends** when ODBA applications are stopped during DLI processing
 - No changes required for ODBA applications
 - Need to add the IMSplex and ODBMNAME keywords to DFSPRP macro
 - Recompile and rebind the DFSxxxx0 load module

- **Reduces complexity** and increases **availability**

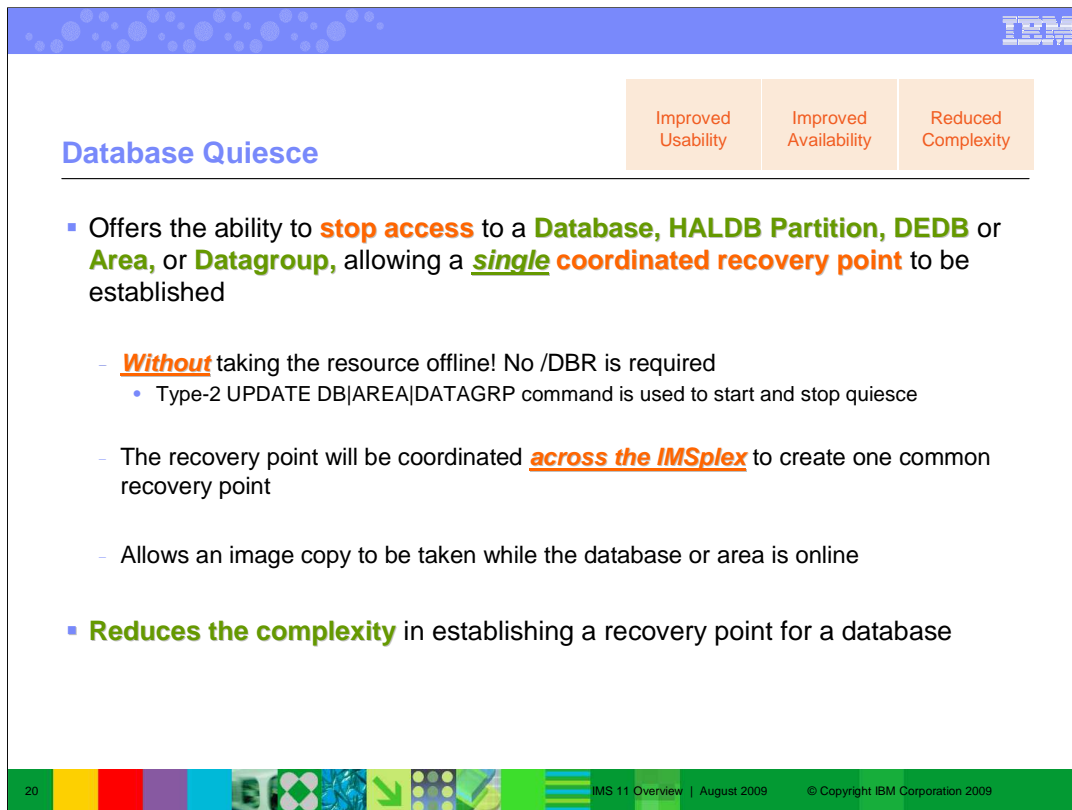
19

IMS 11 Overview | August 2009

© Copyright IBM Corporation 2009

The ODBA interface has a new command for IMS Version 11: CIMS CONNECT. This command will initialize the ODBA interface and will connect to multiple IMS DB systems. Prior to IMS Version 11, ODBA applications used the CIMS INIT command to initialize the ODBA interface and connect to a single IMS DB system. If the application wanted to connect to multiple IMS DB systems, it had to issue multiple CIMS INIT commands, one for each IMS DB system. The ODBA interface is enhanced in IMS Version 11 to support a new command, CIMS CONNECT, which allows the ODBA application to initialize the ODBA interface and connect to multiple IMS DB systems with a single command.

Existing ODBA applications can also use ODBM to protect IMS from abends that are caused by the unexpected termination of the ODBA applications during DLI processing. ODBM support for ODBA application programs requires no changes to the application program, but does require minor changes to the DFSPRP macro, which defines the connections to IMS DB, and a recompile and rebinding of the DFSxxxx0 load module. You can modify your existing ODBA application servers to use ODBM by adding the IMSPLEX and ODBMNAME keywords to the DFSPRP macro. After adding these keywords, you must recompile and rebind the DFSxxxx0 load module (xxxx is the DRA start-up table name specified on the APSB call in the AIBRSNM2 field of the AIB). FUNCLV=2 should be specified when using the V11 DFSPRP macro that contains these two new parameters. FUNCLV=2 is set as the default in the shipped V11 DFSPRP macro.



Database Quiesce

- Improved Usability
- Improved Availability
- Reduced Complexity

- Offers the ability to **stop access** to a **Database, HALDB Partition, DEDB** or **Area, or Datagroup**, allowing a **single coordinated recovery point** to be established
 - **Without** taking the resource offline! No /DBR is required
 - Type-2 UPDATE DB|AREA|DATAGRP command is used to start and stop quiesce
 - The recovery point will be coordinated **across the IMSplex** to create one common recovery point
 - Allows an image copy to be taken while the database or area is online
- **Reduces the complexity** in establishing a recovery point for a database

20 IMS 11 Overview | August 2009 © Copyright IBM Corporation 2009

Database quiesce enables you to create a coordinated recovery point across an IMSplex for IMS Fast Path data entry databases (DEDBs), Fast Path areas, full-function databases (including High Availability Large Databases - HALDBs), and database groups, without taking your resources offline or causing applications to encounter an unavailable database.

When the quiesce function is invoked, the point of consistency is reached when all updates that are in progress are committed. After all updates are committed, and stored on DASD, the database data sets reflect the current database information because no updates are pending. At this time, the point of consistency has been reached locally on a single IMS. This process must occur on every IMS in the IMSplex that is actively using the database. When all activity has been quiesced, a new recovery point is recorded.

You can create a point of consistency on a database without making the database unavailable to applications by using a new form of the existing type-2 UPDATE command (for example, UPDATE DB START (QUIESCE)). This command quiesces all work in progress so that a new recovery point can be created for a database while it is actively in use in the IMSplex.

This enhancement should reduce the complexity of establishing a recovery point for a database. The reduced complexity should allow more frequent recovery points to be created for a database resource. We see value to your business with this enhancement by requiring less complexity in establishing a recovery point. Creating frequent recovery points could reduce the impact of an outage on a database since there would be less data to be processed during recovery.

Database Quiesce...

Improved Usability

Improved Availability

Reduced Complexity

- Creating a recovery point **without** Database Quiesce...
 - Issue /DBR on each IMS to establish a recovery point
 - Each IMS will record a DEALLOC time in the RECON
 - Database resource is closed and data sets are deallocated
 - Check to make sure the database or area was successfully taken offline on each IMS
 - Issue /STA on each IMS to allow database to be used
 - OPEN option on /STA command is required to allocate and open the data sets
 - Otherwise first access will open the data sets

```

graph LR
    User[User] --> IMS1[IMS]
    User --> IMS2[IMS]
    User --> IMS3[IMS]
    IMS1 --> RECON[RECON]
    IMS2 --> RECON
    IMS3 --> RECON
  
```

21
IMS 11 Overview | August 2009
© Copyright IBM Corporation 2009

Before going into the detail of quiesce, let's look at an illustration of the current process required to create a recovery point for a database resource. In this picture there is a simple IMSplex with three different IMS systems in the DBRC sharing group. The database resource is in use on the three different IMS systems.

First a user needs to take the resource offline. To do this the /DBR command is issued for the resource and routed to each IMS or issued on each IMS. On each IMS, the command checks for activity on the resource, if there is none then the command can be processed. Each IMS will go to the RECON and record a DEALLOC time. The database resource is closed and the data sets are deallocated.

Next the user needs to check to make sure that the command was successful on each IMS system.

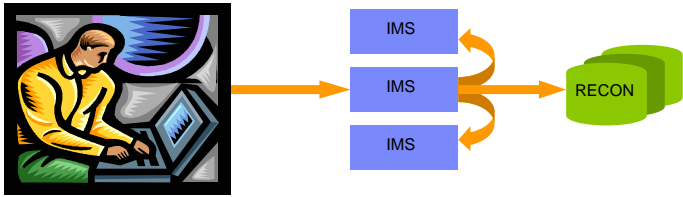
Now a recovery point has been created for the database.

Finally the user needs to make the resource available online. To do this the /STA command is issued for the resource and routed to each IMS or issued on each IMS. On each IMS, the command makes the database available. The database data sets are not allocated, and they are closed after the /DBR command. The OPEN option would be required to allocate and open the data sets on the /STA command, otherwise the first access will open the data sets.

Database Quiesce...

Improved Usability
Improved Availability
Reduced Complexity

- Creating a recovery point **with** Database Quiesce...
 - Issue UPDATE DB|AREA|DATAGRP START(QUIESCE) command
 - Command is processed by one IMS (command master) in the IMSplex
 - Databases and data sets are not closed and deallocated
 - Command waits for active use of the resource to reach a commit point and then quiesces the resource
 - Coordinated across the IMSplex, OLDS are switched on each IMS by default
 - RECON updated by one IMS with common DEALLOC timestamp on behalf of all the IMS systems
 - New allocations for FP areas done on local IMS
 - For FF databases, a new ALLOC is created at first database update after releasing the quiesce
 - Command master communicates with IMS systems to make resource available again



22
IMS 11 Overview | August 2009
© Copyright IBM Corporation 2009

This is an illustration of the process to create a recovery point for a database resource with Database Quiesce. In this picture there is a simple IMSplex with three different IMS systems in the DBRC sharing group. The database resource is in use on the three different IMS systems.

To create the recovery point issue the quiesce command. The command will be processed by a single IMS, the command master. That command master will communicate with the other IMS systems to get the resource quiesced. If there is current activity on the resource the command will wait for a point of consistency to be reached. After that point the access on the database is stopped. The IMS systems will communicate back that the resource is quiesced. The command master will update the RECON on behalf of all the IMS systems, creating a common DEALLOC timestamp. The OLDS are switched on each IMS (as the default). After the quiesce point is reached, the quiesce on the resource is released. When the quiesce is released then a new ALLOC will be recorded for DEDB areas. No ALLOC is recorded for Full Function databases, instead, the first update access to the database will create a new ALLOC in the RECON. The command master communicates with all IMS systems to make the resource available again.

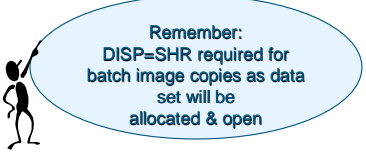
Database Quiesce...

Improved Usability

Improved Availability

Reduced Complexity

- Two types of Quiesce
 - **Quiesce and Hold - UPDATE DB|AREA|DATAGRP START(QUIESCE) OPTION(HOLD)**
 - Reach the quiesce point and **keep the resource quiesced**
 - Allows IC utilities to be run after DB has been quiesced
 - Must issue UPDATE DB|AREA|DATAGRP **STOP(QUIESCE)** to release the quiesce
 - **Quiesce and Go - UPDATE DB|AREA|DATAGRP START(QUIESCE)**
 - **Reach** the quiesce point **and release** the quiesce **in a single command**
 - Used to quickly establish a new recovery point for a resource
 - Could be issued before running a concurrent image copy using IC2
 - Only a small amount of updates would need to be applied to restore DB to this point in time



Remember:
DISP=SHR required for
batch image copies as data
set will be
allocated & open

23

IMS 11 Overview | August 2009

© Copyright IBM Corporation 2009

The quiesce function will be implemented with two different forms of quiesce. A quiesce and hold and a quiesce and go type of command.

The quiesce and hold will allow the database resource to be quiesced and stay in the quiesce state. This would allow utilities like IC2 to be run against the quiesced resource. To release the quiesce, a second command must be issued. This would be similar to issuing a /DBR or /DBD today then running utilities and then issuing a /STA command to allow normal access to the resource to resume. This process will create a recovery point for the database which doesn't require any logs to be applied to recover the database. Note that in order for image copies to be allowed to execute when the database or area is in a QUIESCE HELD state, the DISP for the dataset to be copied must be DISP=SHR (instead of DISP=OLD) for batch image copies since the dataset is open and allocated to the online IMS subsystems.

The quiesce and go is a way of quickly creating a new recovery point for a database resource and then making the resource available again immediately in a single command execution. The quiesce and go may be used for example when a recovery point is created outside of the normal recovery period, when there is more activity on the IMS system. After the quiesce and go has completed, a concurrent image copy of the database could be started using IC2. This process will create a recovery point for the database which only requires a small amount of updates to be applied, from the point after the quiesce until the image copy, to restore the database to this point in time.

Database Quiesce...			
	Improved Usability	Improved Availability	Reduced Complexity

- Databases can be quiesced when **applications reach a sync-point**
 - **IMS holds applications at that point** until the quiesce is released
- Quiesce **“all or none”**
 - All resources specified on the command will be quiesced together
 - If any of the resources cannot be quiesced, the entire list will fail
- Quiesce will have a **timeout value**
 - Timeout will occur after **30 seconds** if the resource is not quiesced
 - User can override the timeout value on the command
 - User can override the default of 30 seconds in the DFSCGxxx or DFSDfxxx PROCLIB members
 - DBQUIESCETO=1-999 seconds

Applications using PSBs with update intent for the resource will be quiesced

24

IMS 11 Overview | August 2009 © Copyright IBM Corporation 2009

When the quiesce function is invoked, the point of consistency is reached when all updates that are in progress are committed. After all updates are committed, and stored on DASD, the database data sets reflect the current database information because no updates are pending. At this time, the point of consistency has been reached locally on a single IMS. This process must occur on every IMS in the IMSplex that is actively using the database. When all activity has been quiesced, a new recovery point is recorded. During the time that the database is quiesced, only read-only applications can access the database. All other applications are held at the time when they attempt to access the database with a DL/I call, until the quiesce is released. Batch jobs with update access or above will fail authorization while the database is quiesced.

When the quiesce is released, the applications are able to access the database again. The quiesce leaves the database data sets in the same state that they were in at the start of the quiesce process.

The DFSCGxxx PROCLIB member and the COMMON_SERVICE_LAYER section of the DFSDfxxx PROCLIB member are updated with the new DBQUIESCETO keyword, where you can specify the maximum amount of time (from one to 999 seconds) to allow a database quiesce process to complete before timing out. The default value is 30 seconds.

Improved Usability
Improved Availability
Reduced Complexity

Database Quiesce...

- Example of Quiesce and hold for a database

```
File Action Manage resources SPOC View Options Help
IMS11          IMS Single Point of Control
Command ==> UPDATE DB NAME(DBXYZ) START(QUIESCE) OPTION(HOLD)

Plex . . . . . Route . . . . . Wait . . . . .

Response for: UPDATE DB NAME(DBXYZ) START(QUIESCE) OPTION(HOLD)
DBName  MbrName  CC
DBXYZ   IM02    0
DBXYZ   IM01    0
DBXYZ   IM03    0
```

```
File Action Manage resources SPOC View Options Help
IMS11          IMS Single Point of Control
Command ==> QRY DB NAME(DBXYZ) SHOW(STATUS)

Plex . . . . . Route . . . . . Wait . . . . .

Response for: QRY DB NAME(DBXYZ) SHOW(STATUS)
DBName  MbrName  CC      TYPE      LclStat
DBXYZ   IM02    0      DLI      ALLOCS,OPEN,QUIESCED
DBXYZ   IM01    0      DLI      ALLOCS,OPEN,QUIESCED
DBXYZ   IM03    0      DLI      ALLOCS,OPEN,QUIESCED
```

25
IMS 11 Overview | August 2009
© Copyright IBM Corporation 2009

There are three IMS systems in the IMSplex, IM01, IM02, and IM03.

The TSO SPOC was used to issue the quiesce and hold command.

The output of the command shows it was successful on three IMS systems.

After the quiesce and hold was performed a QUERY command was issued to display the status of the database resource.

The command shows that the resource is quiesced, but notice that the database is still allocated and open. This is because the quiesce doesn't deallocate the database data sets or close the database data sets.

Improved Usability
Improved Availability
Reduced Complexity

Database Quiesce...

- Example of Releasing the Quiesce

```
File Action Manage resources SPOC View Options Help
IMS11          IMS Single Point of Control
Command ==> UPDATE DB NAME(DBXYZ) STOP(QUIESCE)

Plex . . _____ Route . . _____ Wait . . _____
Response for: UPDATE DB NAME(DBXYZ) STOP(QUIESCE)
DBName  MbrName  CC
DBXYZ   IM02    0
DBXYZ   IM01    0
DBXYZ   IM03    0
```

```
File Action Manage resources SPOC View Options Help
IMS11          IMS Single Point of Control
Command ==> QRY DB NAME(DBXYZ) SHOW(STATUS)

Plex . . _____ Route . . _____ Wait . . _____
Response for: QRY DB NAME(DBXYZ) SHOW(STATUS)
DBName  MbrName  CC      TYPE      LclStat
DBXYZ   IM02    0      DLI      ALLOCS,OPEN
DBXYZ   IM01    0      DLI      ALLOCS,OPEN
DBXYZ   IM03    0      DLI      ALLOCS,OPEN
```

26
IMS 11 Overview | August 2009
© Copyright IBM Corporation 2009

Now release the quiesce on the resource by issuing another UPDATE command.

The output of the command shows it was successful on three IMS systems.

After the quiesce and hold was released, a QUERY command was issued to display the status of the database resource.

The command shows that the resource is allocated and open.

Database Quiesce...

Improved Usability

Improved Availability

Reduced Complexity

- **Type-1** DISPLAY commands are enhanced to show database status after quiesce
 - **DIS DB/AREA QSC**
 - **DIS DB STATUS**
 - **DIS STATUS**

```

File Action Manage resources SPOC View Options Help
IMS11          IMS Single Point of Control
Command ==>   DIS DB QSC

Log for . . . : DIS DB QSC          Flex . . .   Route . . .   Wait . . .

IMSplex . . . . . : IMS11
Routing . . . . . :
Start time. . . . . : 2008.255 16:02:18.08
Stop time . . . . . : 2008.255 16:02:18.09
Return code . . . . : 00000000
Reason code . . . . : 00000000
Reason text . . . . :
Command master. . . : IM0A

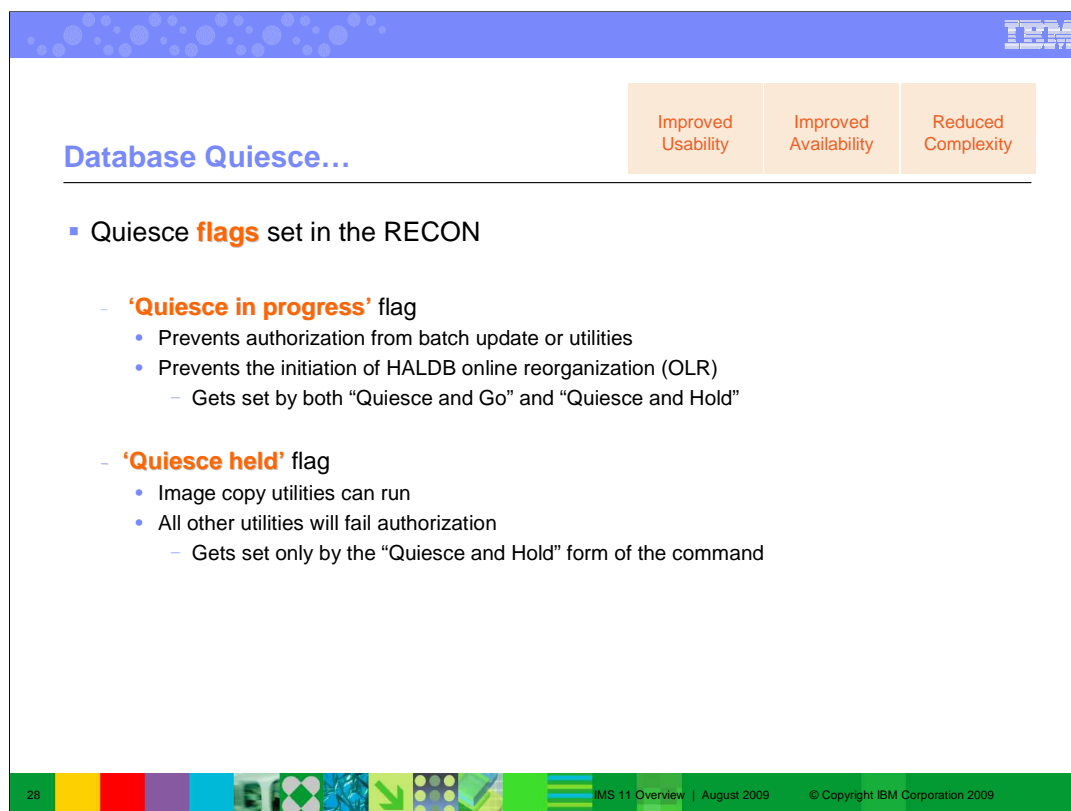
MbrName  Messages
IM0A     DATABASE TYPE  TOTAL      UNUSED     TOTAL      UNUSED     ACC      CONDITIONS
IM0A     DBXYZ1  DL/I                UNUSED     TOTAL      UNUSED     UP       ALLOCS ,QUIESCING
IM0A     DBXYZ3  DL/I                UNUSED     TOTAL      UNUSED     UP       ALLOCS ,QUIESCED
            
```

27
IMS 11 Overview | August 2009
© Copyright IBM Corporation 2009

The following Type-1 commands are enhanced to show quiesce status for a database resource:

- DIS AREA with QSC attribute – displays all FP areas that have status of ‘quiescing’ (indicating quiesce in progress) or ‘quiesced’ (indicating the resource is quiesced)
- DIS DB with QSC attribute – displays all databases that have status of ‘quiescing’ (indicating quiesce in progress) or ‘quiesced’ (indicating the resource is quiesced)
- DIS DB STATUS
- DIS STATUS

The command specified in the example is to show database resources that have status of ‘quiescing’ or ‘quiesced’.



Database Quiesce...

- Improved Usability
- Improved Availability
- Reduced Complexity

- Quiesce **flags** set in the RECON
 - **'Quiesce in progress'** flag
 - Prevents authorization from batch update or utilities
 - Prevents the initiation of HALDB online reorganization (OLR)
 - Gets set by both "Quiesce and Go" and "Quiesce and Hold"
 - **'Quiesce held'** flag
 - Image copy utilities can run
 - All other utilities will fail authorization
 - Gets set only by the "Quiesce and Hold" form of the command

When quiesce is started, with either form of the command, the 'Quiesce in progress' flag in the RECON will be set on. While the flag is on, any new batch update or utility authorization of the database or area will be prevented. All online IMS subsystems including update access and batch jobs requesting read access will be allowed authorization. All utilities will fail authorization. Batch jobs with read access only and online IMS subsystems with read access only will be allowed access to the database or area. IMS systems which didn't participate in the quiesce will be allowed authorization but will be prevented from accessing the database or area while the flag is on in the RECON data sets.

Also the quiesce function will not be allowed during integrated HALDB online reorganization (OLR).

Database Quiesce...

- Improved Usability
- Improved Availability
- Reduced Complexity

- For **DBRC registered databases** the quiesce point **is** recorded in the RECON
 - Same DEALLOC timestamp is recorded for all open ALLOC records by a single IMS
 - For DEDB Areas, a new ALLOC is created
 - For Full Function Databases, a new ALLOC is created at first database update after releasing the quiesce
 - New USID/DSSN created after quiesce point is reached
- For **non-registered databases** the quiesce point **is not** recorded in the RECON
 - A type x'4C' log record is created for databases, x'59' for areas
 - After the log record is created, recovery point is recorded in the RECON

29 IMS 11 Overview | August 2009 © Copyright IBM Corporation 2009

The quiesce point will be a common DEALLOC time in the RECON for all open ALLOCs.

When the quiesce is achieved any open ALLOC record in the RECON data sets is updated. The open ALLOC record will be closed with a DEALLOC time which corresponds to the point the quiesce was achieved across the IMSplex. The closed ALLOC record will show that the deallocation was the result of a quiesce command.

When the quiesce is reached, the command master will update the open ALLOCs in the RECON on behalf of all the IMS systems.

When the quiesce is released (either quiesce and go or release of quiesce after quiesce and hold) then a new ALLOC will be recorded for DEDB Areas. No ALLOC is recorded for Full function databases, instead the first access to the database will create a new ALLOC in the RECON.

If the database is not registered with DBRC, then the assumption is that it is not being shared with other systems. Since it is not being shared, there is no coordination. Also since there is no RECON to update the quiesce point can only be recorded on the log. When the quiesce has caused a point of consistency to be reached, a log record is created. The type X'4C' log record is created for databases. The type X'59' log record is created for areas. The log records are used by each participant in the quiesce to note that the database has been quiesced. After the log record is created, the recovery point is recorded in the RECON data sets.

When a quiesce and hold had been issued, the quiesce status is maintained across restarts of IMS. A call is made to DBRC during restart processing to recover the status.

The quiesce point will be a common DEALLOC time in the RECON for all open ALLOCs.

When the quiesce is achieved any open ALLOC record in the RECON data sets is updated. The open ALLOC record will be closed with a DEALLOC time which corresponds to the point the quiesce was achieved across the IMSplex. The closed ALLOC record will show that the deallocation was the result of a quiesce command.

When the quiesce is reached, the command master will update the open ALLOCs in the RECON on behalf of all the IMS systems.

When the quiesce is released (either quiesce and go or release of quiesce after quiesce and hold) then a new ALLOC will be recorded for DEDB Areas. No ALLOC is recorded for Full function databases, instead the first access to the database will create a new ALLOC in the RECON.

If the database is not registered with DBRC, then the assumption is that it is not being shared with other systems. Since it is not being shared, there is no coordination. Also since there is no RECON to update the quiesce point can only be recorded on the log. When the quiesce has caused a point of consistency to be reached, a log record is created. The type X'4C' log record is created for databases. The type X'59' log record is created for areas. The log records are used by each participant in the quiesce to note that the database has been quiesced. After the log record is created, the recovery point is recorded in the RECON data sets.

When a quiesce and hold had been issued, the quiesce status is maintained across restarts of IMS. A call is made to DBRC during restart processing to recover the status.

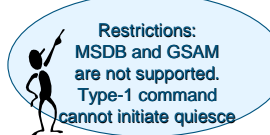
Database Quiesce...

Improved Usability

Improved Availability

Reduced Complexity

- Software prerequisites
 - **Common Service Layer** (OM, RM, SCI)
 - RM with resource structure recommended but not required
 - RM is used to coordinate the quiesce process across the IMSplex
 - RM is not required (RMENV=N) when using the DB quiesce function in a single-IMS IMSplex
 - **MINVERS('11.1')** must be set in RECON



Restrictions:
MSDB and GSAM are not supported.
Type-1 command cannot initiate quiesce

30

IMS 11 Overview | August 2009

© Copyright IBM Corporation 2009

The CSL Resource Manager (RM) is required when using the DB quiesce function in a multi-IMS IMSplex. An RM resource structure is recommended, but not required. RM is not required (RMENV=N) when using the DB quiesce function in a single-IMS IMSplex.

ACBLIB Usability Enhancements

- Improved Usability
- Improved Availability
- Improved Performance

- The ACBLIB Usability enhancements enable you to:
 - Load the ACB members into **64-bit storage**
 - Applies to **Non-resident PSBs** and **DMBs** (DEDBs are not loaded into 64-bit)
 - Reducing read I/O to the ACBLIB data set
 - Create **DFSMDA** members for the **dynamic allocation** of the **ACBLIB datasets**
 - A Restart of IMS is no longer required to:
 - Increase the size of the ACBLIB data sets
 - Correct errors with the inactive ACBLIB
 - Add additional data sets to the ACBLIB concatenation

31 | IMS 11 Overview | August 2009 | © Copyright IBM Corporation 2009

These enhancements add options to:

- Load the ACB members into 64-bit storage. At application scheduling time, the ACB members are retrieved from 64-bit storage rather than from the ACBLIB data set (unless they are not currently in the 64-bit cache). This option reduces the amount of read I/O to the ACBLIB data set. By changing the resident option for a PSB, the user could have a smaller resident pool and also realise some CSA reduction. There will be no performance degradation at scheduling since the PSB is retrieved from the 64-bit pool.

- Use the Dynamic Allocation macro (DFSMDA) macro to create DFSMDA members for the dynamic allocation of the ACBLIB data sets. You can specify how the ACBLIB data sets are allocated with either JCL or DFSMDA members. DFSMDA members provide the following benefits:

- You can increase the size of the ACBLIB data sets without stopping and restarting IMS.
- You can correct errors with the inactive ACBLIB without stopping and restarting IMS.
- You can add additional data sets to the ACBLIB concatenation without stopping and restarting IMS.

The dynamic allocation of the ACBLIB data sets and the 64-bit storage pool for ACB members is supported in the following online IMS system configurations: DB/TM, DBCTL, DCCTL, XRF, and FDBR. Dynamic allocation of ACBLIB datasets and the 64-bit storage pool for ACB members are not supported in IMS batch.

Dynamic Allocation of ACBLIB data sets

Improved Usability

Improved Availability

- To implement **DFSMDA** members for the **dynamic allocation** of the **ACBLIB datasets**:
 - Create a DFSMDA member for each of the **IMSACBA** and **IMSACBB** data set concatenations
 - Remove the IMSACBA and IMSACBB JCL statements from the IMS JCL & DL/I SAS procedures
- *Example:*

```

DFSMDA TYPE=INITIAL
DFSMDA TYPE=IMSACBA
DFSMDA TYPE=DATASET,DSNAME=IMS.ACBLIB1
DFSMDA TYPE=DATASET,DSNAME=IMS.ACBLIB2
DFSMDA TYPE=FINAL
DFSMDA TYPE=INITIAL
DFSMDA TYPE=IMSACBB
DFSMDA TYPE=DATASET,DSNAME=IMS.ACBLIB3
DFSMDA TYPE=DATASET,DSNAME=IMS.ACBLIB4
DFSMDA TYPE=FINAL
                    
```
- */DIS MODIFY changed to show status of "U" for unallocated data sets:*

```

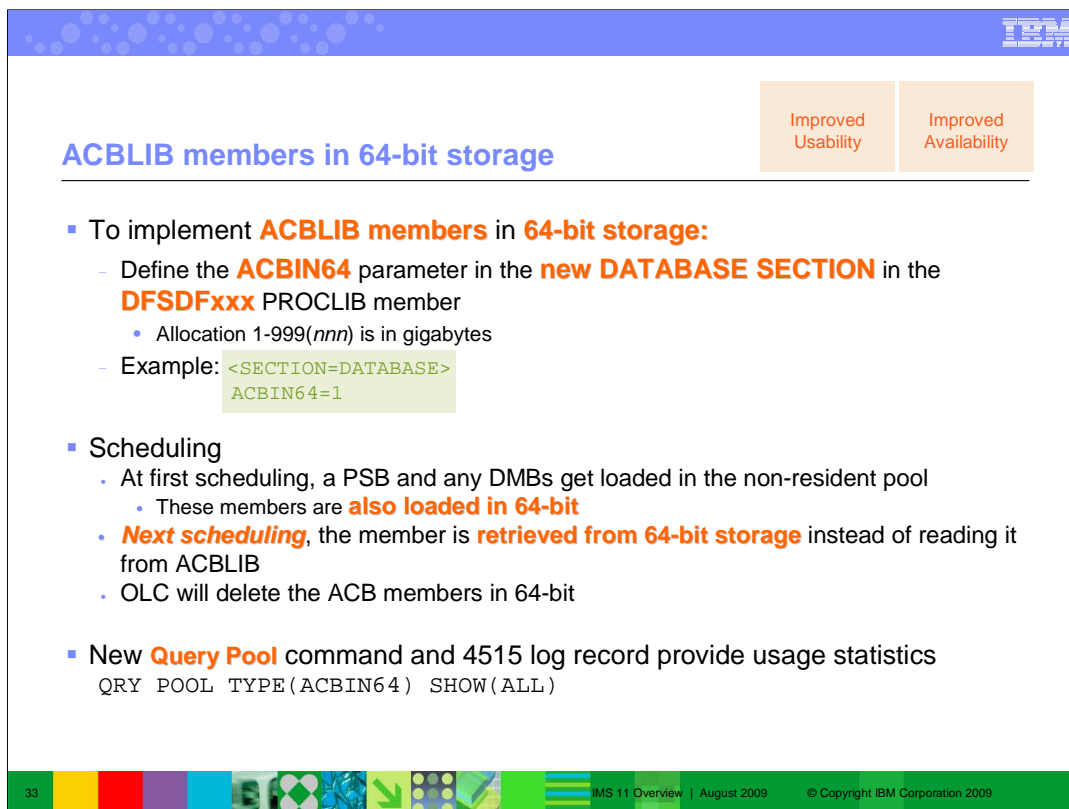
LIBRARY IMSACBA (A) IMS.ACBLIB1
LIBRARY IMSACBA (A) IMS.ACBLIB2
LIBRARY FORMATA (A) IMSTESTG.MFS.FORMAT1
LIBRARY FORMATA (A) IMSTESTG.MFS.FORMAT2
LIBRARY FORMATA (A) IMSQA.FMT1
LIBRARY MODBLKSA (A) IMSBLD.I11ATS17.COMBLKS1
LIBRARY IMSACBB (U) IMS.ACBLIB3
LIBRARY IMSACBB (U) IMS.ACBLIB4
LIBRARY FORMATB (I) IMSTESTG.MFS.FORMAT3
LIBRARY FORMATB (I) IMSTESTG.MFS.FORMAT4
LIBRARY FORMATB (I) IMSQA.FMT1
LIBRARY MODBLKSB (I) IMSBLD.I11ATS17.COMBLKS2
DISPLAY MODIFY COMPLETE *08003/110121* SYS3
                    
```

32

IMS 11 Overview | August 2009

© Copyright IBM Corporation 2009

The /DIS MODIFY command is changed to show the inactive ACBLIB libraries when MDA members are used. IMS will dynamically allocate the Active ACBLIB during initialization and will deallocate it during Online Change. The inactive ACBLIB datasets are not allocated until an OLC process. The datasets will show a status of "U". This indicates that these datasets are unallocated.



ACBLIB members in 64-bit storage

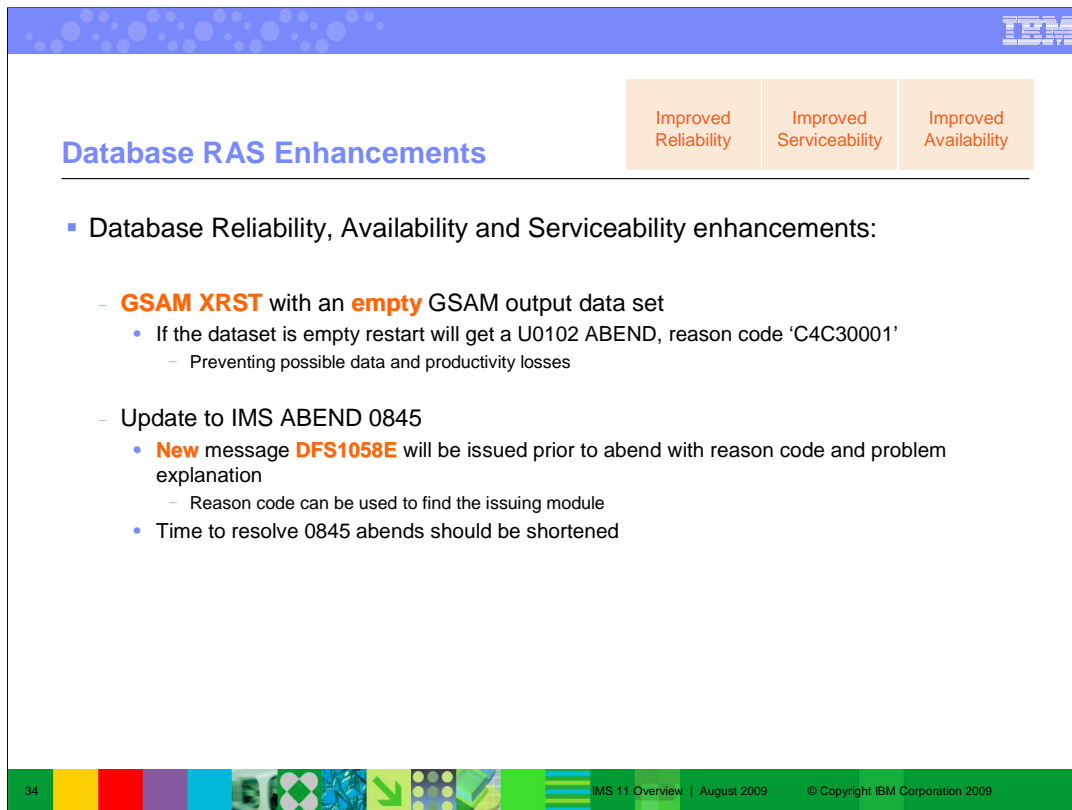
- To implement **ACBLIB members** in **64-bit storage**:
 - Define the **ACBIN64** parameter in the **new DATABASE SECTION** in the **DFSDFxxx** PROCLIB member
 - Allocation 1-999(*nnn*) is in gigabytes
 - Example:

```
<SECTION=DATABASE>
ACBIN64=1
```
- Scheduling
 - At first scheduling, a PSB and any DMBs get loaded in the non-resident pool
 - These members are **also loaded in 64-bit**
 - **Next scheduling**, the member is **retrieved from 64-bit storage** instead of reading it from ACBLIB
 - OLC will delete the ACB members in 64-bit
- New **Query Pool** command and 4515 log record provide usage statistics


```
QRY POOL TYPE(ACBIN64) SHOW(ALL)
```

The allocation quantity is in gigabytes.

The QUERY POOL command is a new command for V11. It will be used to display statistics for IMS storage pools. This command will be used to display the storage utilization statistics for the 64-bit ACB pool. It will also display statistics for other IMS pools such as the Fastpath buffer pool. The 4515 log record contains the same statistics that are displayed on the QRY POOL command.



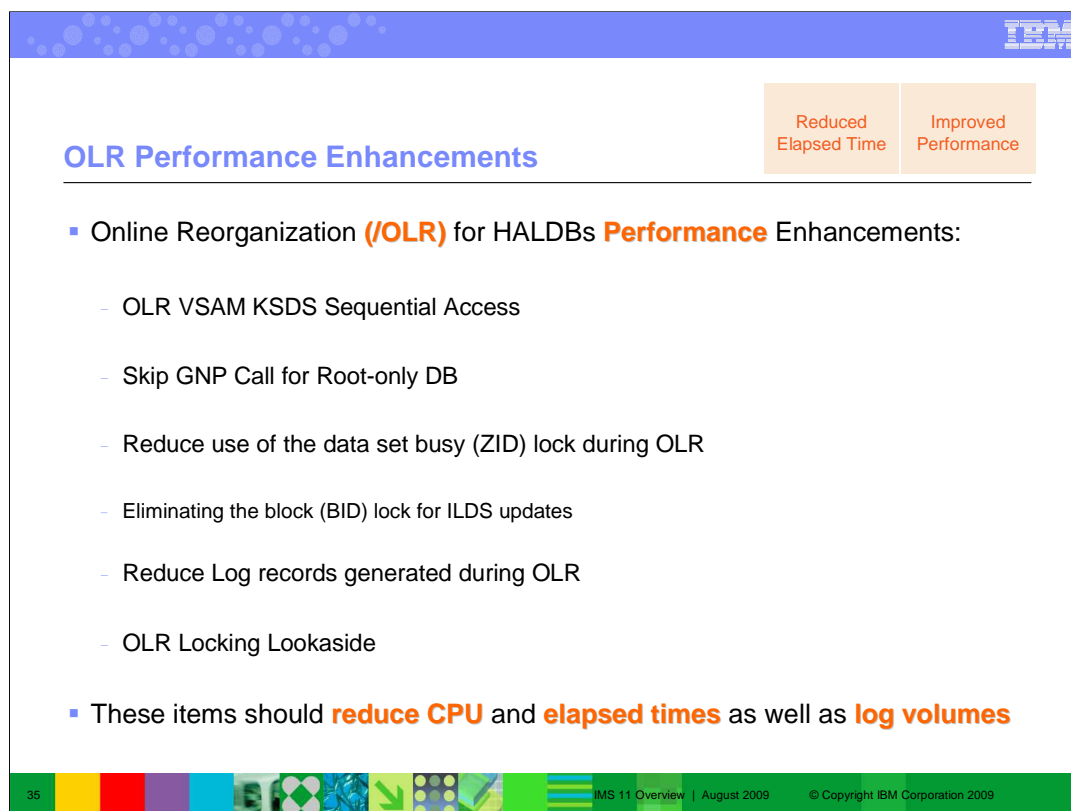
Database RAS Enhancements

- Improved Reliability
- Improved Serviceability
- Improved Availability

- Database Reliability, Availability and Serviceability enhancements:
 - GSAM XRST** with an **empty** GSAM output data set
 - If the dataset is empty restart will get a U0102 ABEND, reason code 'C4C30001'
 - Preventing possible data and productivity losses
 - Update to IMS ABEND 0845
 - New** message **DFS1058E** will be issued prior to abend with reason code and problem explanation
 - Reason code can be used to find the issuing module
 - Time to resolve 0845 abends should be shortened

This improvement prevents you from accidentally pointing to an empty GSAM dataset when you need to do an extended restart (the XRST command). During GSAM XRST processing, a check is made to determine if the GSAM output dataset to be repositioned is empty. If the dataset is empty and the abending job had previously inserted records to the dataset, the restart job will issue an U102 abend with the new reason code 'C4C30001'.

Prior to issuing a U0845 abend, the DFS1058E message will be issued, listing the reason code and a short explanation of the problem. The reason code will also be returned in Reg 1 of the 'Registers at Abend'. The reason codes have been changed and can now be used to identify the abending module.



OLR Performance Enhancements

Reduced
Elapsed Time
Improved
Performance

- Online Reorganization (**OLR**) for HALDBs **Performance** Enhancements:
 - OLR VSAM KSDS Sequential Access
 - Skip GNP Call for Root-only DB
 - Reduce use of the data set busy (ZID) lock during OLR
 - Eliminating the block (BID) lock for ILDS updates
 - Reduce Log records generated during OLR
 - OLR Locking Lookaside
- These items should **reduce CPU** and **elapsed times** as well as **log volumes**

IMS 11 Overview | August 2009 © Copyright IBM Corporation 2009

OLR VSAM KSDS Sequential Access: OLR is enhanced to take advantage of the VSAM sequential access option when issuing KSDS GET requests to retrieve sequentially from the input data set(s). This results in reduction in CPU and elapsed time.

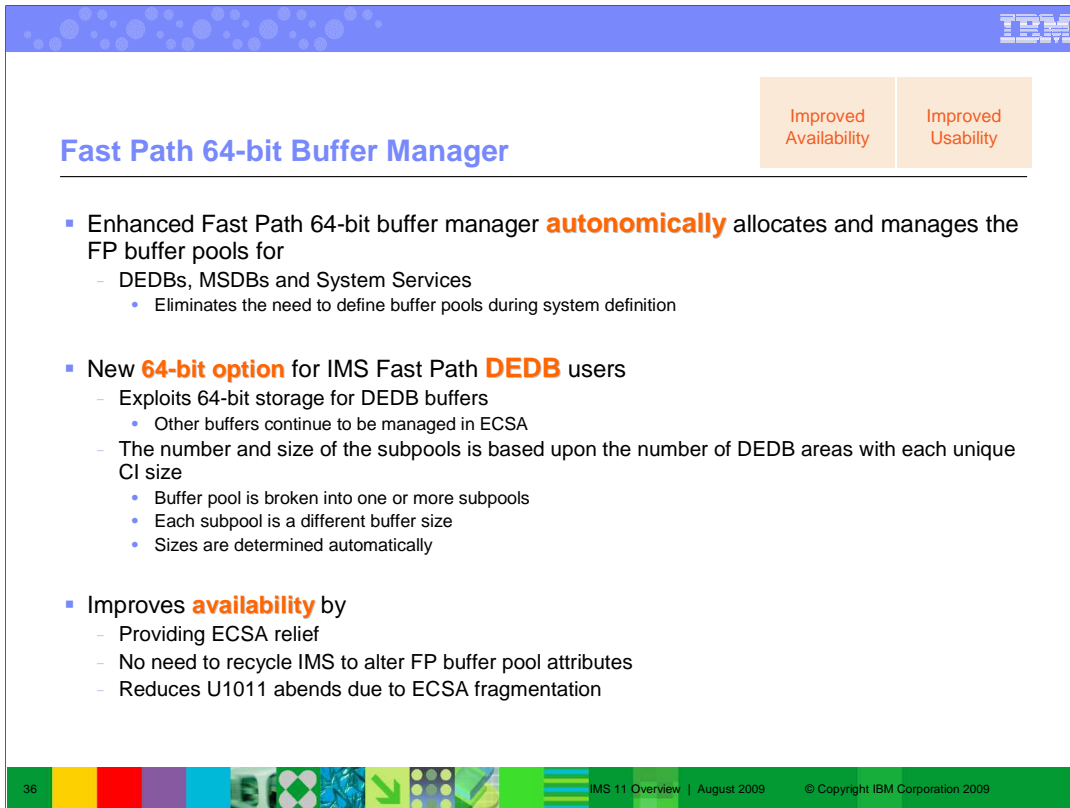
Skip GNP Call for Root-only DB: For a root-only database, there is no reason to issue the GNP call since there is no dependent segment to be read. Skipping the GNP call saves CPU and elapsed time.

Reduce use of the data set busy (ZID) lock during OLR: For a PHIDAM partition undergoing reorganization by OLR the updates for the primary index (KSDS) can be saved for insertion at the end of the unit of reorganization (UOR). By saving all the KSDS updates until the end of the UOR, the usage of the ZID lock for the primary index can be UOR changed. The change will be to obtain the ZID once before starting to insert all the saved primary index updates, and then released once. When the UOR covers many roots, many ZID lock requests will be eliminated which will result in a reduction of CPU usage and elapsed time.

Eliminating the block (BID) lock for ILDS updates: For a HALDB partition with logical relationships the updates for the indirect list dataset (ILDS) by OLR do not need to obtain the block (BID) lock. The BID lock is used for serialization of the updates to a block across IMS data sharers. However the design of HALDB doesn't allow the ILDS to be updated by more than one IMS at a time. Therefore the BID lock does not need to be obtained. However the ZID lock needs to be obtained in order to notify other IMS data sharers in the case of CI/CA splits. To reduce the use of the ZID lock during OLR, the updates for the ILDS can be saved for insertion at the end of the UOR. The ZID lock will be obtained once before starting to insert all the saved ILDS updates, and then released once. When the UOR covers many to the ILDS, many ZID lock request will be eliminated which will result in a reduction of CPU usage and elapsed time.

Log reduction: For a HALDB partition undergoing OLR, the database update log records (type '50'x) will be consolidated when possible into full block updates. By combining all the updates for a full block into a single type '50'x log record, many of the small type '50'x log records will be eliminated. This will result in a reduction in the log volume generated by OLR.

Lock reduction: Whenever a lock request is made by the OLR owning region, a lookaside operation is performed. The lookaside function will determine if this is a new, or already owned, lock. If this is a new lock, normal processing to call IRLM to get the lock is performed. If this is a lock that is already owned, the lock manager returns to the caller. The savings in path length by optimizing the lock manager out of the call flow results in reduction in CPU and elapsed time.

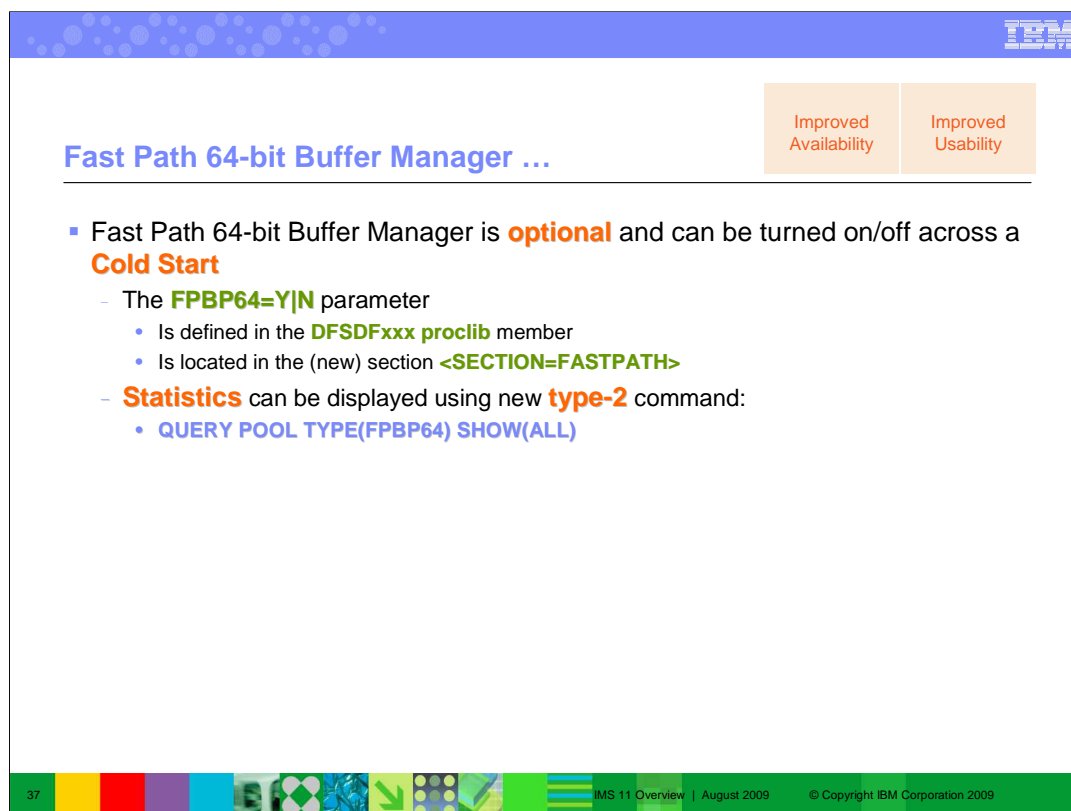


Fast Path 64-bit Buffer Manager

- Enhanced Fast Path 64-bit buffer manager **autonomically** allocates and manages the FP buffer pools for
 - DEDBs, MSDBs and System Services
 - Eliminates the need to define buffer pools during system definition
- New **64-bit option** for IMS Fast Path **DEDB** users
 - Exploits 64-bit storage for DEDB buffers
 - Other buffers continue to be managed in ECSA
 - The number and size of the subpools is based upon the number of DEDB areas with each unique CI size
 - Buffer pool is broken into one or more subpools
 - Each subpool is a different buffer size
 - Sizes are determined automatically
- Improves **availability** by
 - Providing ECSA relief
 - No need to recycle IMS to alter FP buffer pool attributes
 - Reduces U1011 abends due to ECSA fragmentation

The Fast Path 64-bit buffer manager autonomically controls the number and size of Fast Path buffer pools, including buffer pools for data entry databases (DEDBs), main storage databases (MSDBs), and system services, which eliminates the need for users to manually set buffer pool specifications during system definition. The Fast Path 64-bit buffer manager also places the DEDB buffer pools above the bar in 64-bit storage, which reduces the usage of ECSA storage. Fast Path buffer pools for MSDB databases, buffers for inserting sequential dependent (SDEP) segments, buffers for system services, and buffer headers continue to be managed in ECSA storage. The existing mechanism for managing Fast Path buffer pools, which places all Fast Path buffer pools in 32-bit ECSA storage below the bar, requires the number and size of buffer pools to be set during system definition by using the DBBF, DBFX, and BSIZ execution parameters. After IMS is started, the number and size of the Fast Path buffer pools cannot be changed without stopping and restarting IMS.

If a database is added to the online IMS system and none of the active buffer subpools can accommodate the CI size of the database, the Fast Path 64-bit buffer manager allocates a new buffer subpool of the appropriate size.



Fast Path 64-bit Buffer Manager ...

Improved Availability Improved Usability

- Fast Path 64-bit Buffer Manager is **optional** and can be turned on/off across a **Cold Start**
 - The **FPBP64=YJN** parameter
 - Is defined in the **DFSDFxxx proclib** member
 - Is located in the (new) section **<SECTION=FASTPATH>**
 - **Statistics** can be displayed using new **type-2** command:
 - **QUERY POOL TYPE(FPBP64) SHOW(ALL)**

37 IMS 11 Overview | August 2009 © Copyright IBM Corporation 2009

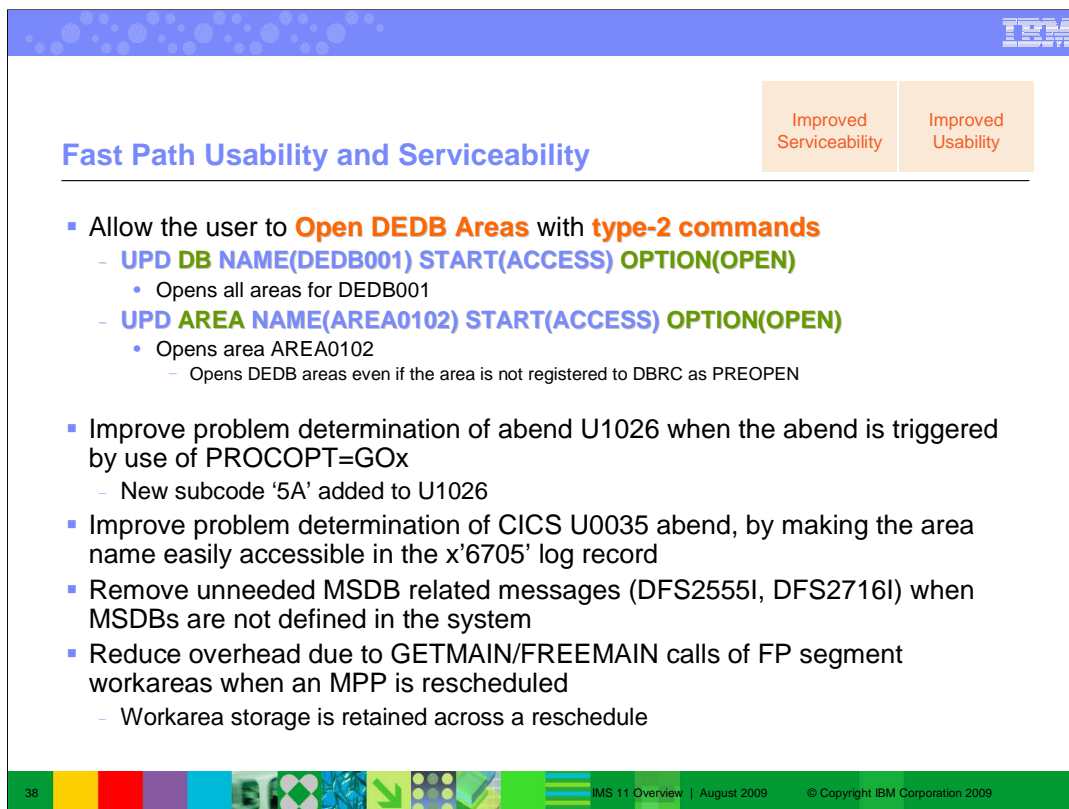
The Fast Path 64-bit buffer manager is enabled by the parameter, FPBP64, in the <SECTION_FASTPATH> of the DFSDFxxx PROCLIB member.

When the Fast Path 64-bit buffer manager is enabled:

- Database administrators can improve the performance of their DEDB databases by changing the CI size of databases, without having to adjust the buffer sizes to match. If no buffer subpools are active that can accommodate a new or changed CI size, the Fast Path 64-bit buffer manager automatically allocates a buffer subpool with the correct CI size.
- Users can display statistics for Fast Path buffers by issuing the new IMS type-2 command QUERY POOL TYPE(FPBP64) SHOW(ALL).
- Multiple application programs can access the overflow buffers (specified with the OBA parameter) in parallel.

The values specified on the overflow buffer allocation (OBA) parameter and on the normal buffer allocation (NBA) parameters affect the Fast Path 64-bit buffer manager. Combined together, they determine the maximum number of buffers that the Fast Path 64-bit buffer manager can allocate to a dependent region.

Users can display statistics for Fast Path buffers by issuing the new IMS type-2 command QUERY POOL TYPE(FPBP64) SHOW(ALL).



Fast Path Usability and Serviceability

- Allow the user to **Open DEDB Areas** with **type-2 commands**
 - **UPD DB NAME(DEDB001) START(ACCESS) OPTION(OPEN)**
 - Opens all areas for DEDB001
 - **UPD AREA NAME(AREA0102) START(ACCESS) OPTION(OPEN)**
 - Opens area AREA0102
 - Opens DEDB areas even if the area is not registered to DBRC as PREOPEN
- Improve problem determination of abend U1026 when the abend is triggered by use of PROCOPT=GOx
 - New subcode '5A' added to U1026
- Improve problem determination of CICS U0035 abend, by making the area name easily accessible in the x'6705' log record
- Remove unneeded MSDB related messages (DFS2555I, DFS2716I) when MSDBs are not defined in the system
- Reduce overhead due to GETMAIN/FREEMAIN calls of FP segment workareas when an MPP is rescheduled
 - Workarea storage is retained across a reschedule

The IMS Fast Path usability and serviceability enhancements provide IMS customers with additional ways to open DEDB areas by using UPDATE DB and UPDATE AREA type-2 commands. Currently /STA area does not open DEDB area logically. With this enhancement, UPD AREA and UPD DB commands can open area logically and save time at first access of applications. By using the UPDATE commands, Fast Path customers can open DEDB areas even if the area is not registered to DBRC as PREOPEN. The DEDB area open process still requires DBRC authorization. The following commands are affected by this enhancement:

- The OPTION(OPEN) keyword on the UPDATE DB command is enhanced to work with Fast Path DEDB areas. Previously, this keyword worked only with full-function databases.
- A new keyword, OPTION(OPEN), is added to the syntax and parameter descriptions of the UPDATE AREA command.

Several enhancements are related to troubleshooting and enable both IMS customers and IMS Service to more efficiently resolve problems that are related to Fast Path.

Problem determination time is reduced by adding a new subcode '5A' to ABEND1026 when it's issued due to PROCOPT=GOx.

Add area name to CICS abendu0035 for problem determination - In analyzing a CICS U0035 which was due to a DEDB area being closed out from under a CICS thread, it was more difficult than needed to find the Area Name. In IMS 11 the area name is added to the x'6705' log record.

Remove MSDB related messages when no MSDBs - Remove these MSDB related messages. When RG suffix does not have MSDB suffix or there is no MSDB defined in MSDBINIT, these messages are issued. DFS2555I NO SUFFIX FOR MSDB MEMBER DEFINED PC

DFS2591I NO MSDB HEADERS FOUND, IMAGE COPY LOAD IGNORED PC

DFS2716I NO MSDBS FOUND - NO MSDB CHECKPOINT TAKEN PC

Reduce GETMAIN/FREEMAIN calls during MPP reschedule process - During each schedule of a dependent region, storage is obtained for a DBFSEG workarea. For IFPs and BMPs, this is acceptable. But for MPPs, which can continue to reschedule while the region is up, this causes performance hits due to the FREEMAIN/GETMAIN overhead. This enhancement changes the processing so the DBFSEG storage is retained across a reschedule. The storage is released if the dependent region terminates, or the existing DBFSEG storage is not large enough for the new DBFSEG blocks.



IBM Software Group

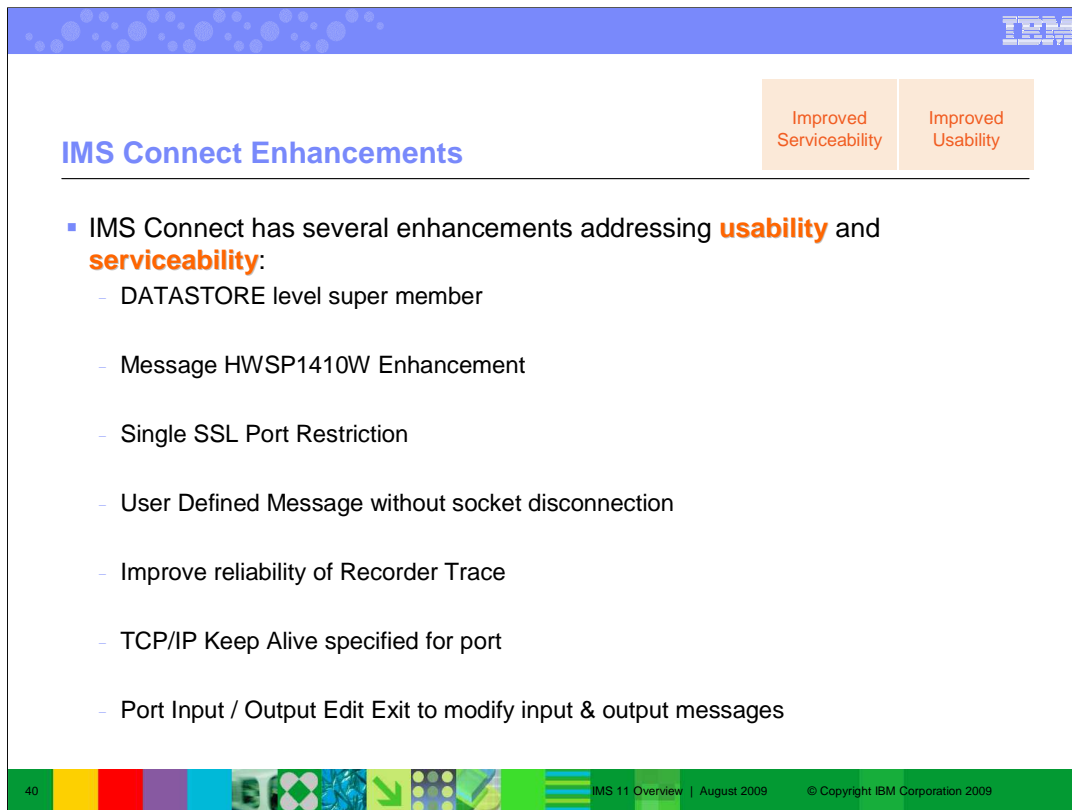


System Enhancements

Information Management software



© Copyright IBM Corporation 2009



The slide features a blue header with the title 'IMS Connect Enhancements' and two orange callout boxes: 'Improved Serviceability' and 'Improved Usability'. The main content is a bulleted list of enhancements. The footer includes a decorative bar with icons, the page number '40', and the text 'IMS 11 Overview | August 2009 © Copyright IBM Corporation 2009'.

IMS Connect Enhancements

Improved
Serviceability

Improved
Usability

- IMS Connect has several enhancements addressing **usability** and **serviceability**:
 - DATASTORE level super member
 - Message HWSP1410W Enhancement
 - Single SSL Port Restriction
 - User Defined Message without socket disconnection
 - Improve reliability of Recorder Trace
 - TCP/IP Keep Alive specified for port
 - Port Input / Output Edit Exit to modify input & output messages

40 IMS 11 Overview | August 2009 © Copyright IBM Corporation 2009

IMS Connect Enhancements – several enhancements for user requirements, usability and serviceability.

DATASTORE level super member – The current IMS Connect support for the OTMA super member feature allows customers to specify only one super member name for each instance of IMS Connect. All datastore connections from IMS Connect must use the same super member name. Customers would like IMS Connect to support multiple super member names, so that, within a single instance of IMS Connect, different groups of client application programs can use different super members. With this enhancement, the super member name can be specified at the connection level so that customers can group client applications by the super member name that they share.

Message HWSP1410W Enhancement - IMS Connect reports errors when freeing storage with message HWSP1410W. This message reports the error return code, type of storage and the module that encountered the error. To further aid problem determination, the address of the storage is added to this message.

Single SSL Port Restriction - IMS Connect does not support multiple SSL ports in an IMS Connect instance. Through this enhancement, IMS Connect will restrict users from specifying more than one SSL port in the IMS Connect configuration member to prevent user errors that will result in potential failure in the Language Environment.

User Defined Message without socket disconnection – This allows customers to modify the user message exit to send a user defined message back to the client application as a reply to the input message with the request to keep the socket connected.

Recorder Trace:

- Improved IMS Connect product reliability by reducing system outages caused by an ABENDS0C4 in the Recorder Trace process.
- Eliminate Recorder Trace dataset full conditions and the resulting loss of diagnostic data.
- Improve the reliability of diagnostic information to shorten or streamline the problem determination process.
- Allow the user more flexibility and control over the amount of diagnostic data recorded.

TCP/IP KeepAlive - Allow IMS Connect customers to set the TCP/IP KeepAlive value at the Port level to override the default stack value.

Port Input/Output Edit Exit - Provide IMS Connect customers with the ability to modify input messages received from TCP/IP before IMS Connect processes them and output messages after IMS Connect has completed processing before they are passed to TCP/IP.

IMS Connect Enhancements...

Improved Serviceability Improved Usability

- IMS Connect has several enhancements addressing **usability** and **serviceability**...
 - Display command enhancements
 - OTMA CM0 (Commit-Then-Send) ACK timeout support
 - IMS TM Resource Adapter Generated Clientid
 - MAXSOC warning messages before limit is reached
 - Cancel Client ID to cancel and re-establish a socket connection
 - TCP/IP Auto Reconnect when network failed and restarted

User message exits HWSIMSO0 and HWSIMSO1 are removed in IMS 11

41 IMS 11 Overview | August 2009 © Copyright IBM Corporation 2009

IMS Connect Enhancements – several enhancements for user requirements, usability and serviceability...

Display Command Enhancements – Provide a “summary” version of the VIEWHWS & QUERY MEMBER commands that will only report the summary information for PORTs without the potentially large amounts of individual socket displays. The addition of the Datastore last accessed to the client information displayed will allow customers to determine to which IMS that request was sent.

OTMA CM0 ACK Timeout Support - Enable IMS Connect customers to specify the name of the Queue to be used by OTMA for rerouting any CM0 messages that timeout, while waiting for the ACK rather than using the OTMA default queue name.

IMS TM Resource Adapter Generated Clientid - Eliminate the requirement that customers use different IMS Connect PORTs for instances of distributed WAS when using ITMRA (IC4J) Shareable Persistent sockets. Generate the HWSxxxx clientid if none is provided. To be exploited by corresponding IMS TM Resource Adapter change.

MAXSOC Warning - Provide IMS Connect customers with warning messages before the max socket limit is reached. Provide the ability to specify the warning level to start issuing the warning messages.

Cancel Client ID - Provide IMS Connect customers with the ability to cancel an existing socket connection and establish a new socket connection with the same Client Id.

TCP/IP Auto Reconnect - Eliminate operator intervention to re-establish the connection to the TCP/IP network where the network has failed and restarted during an instance of the IMS Connect.

Removal of Obsolete User Message Exits - Eliminate the non-source (OCO) User Message Exits (HWSIMSO0 & HWSIMSO1) that have been functionally stabilized and are not modifiable by customers. Exits are replaced by HWSSMPL0/1.

User Exit Interface Enhancements

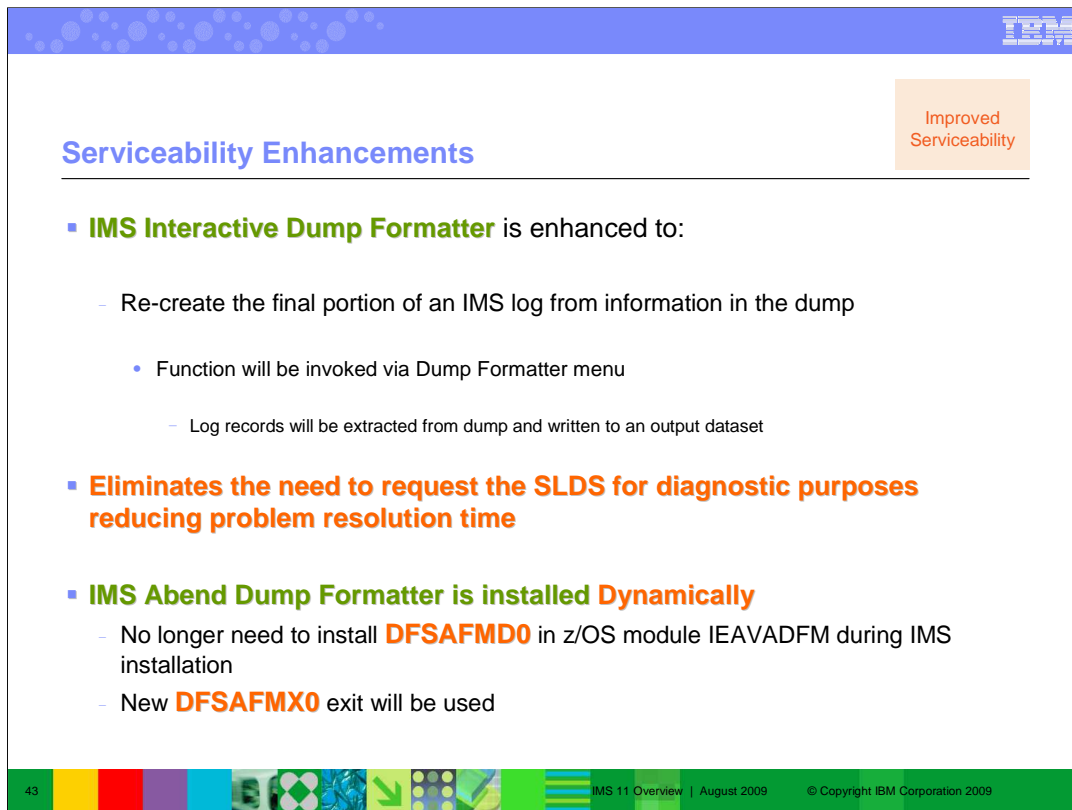
- Enhancements for **selected** IMS Control Region user exits
 - **Multiple instances** of a **user exit type** can be defined
 - User exit modules can be **refreshed** while IMS is active
 - New User Exit types (*which can be refreshed*)
 - EINIT : Early Initialization
 - ICQSEVNT: IMS CQS Event
 - ICQSSTEV: IMS CQS Structure Event
 - Introduces two new Type-2 commands
 - **QUERY USEREXIT**
 - **REFRESH USEREXIT**

To be made available through the Service process after GA

The User exit enhancements line item introduces the following new functions in IMS Version 11 and applies to the new exit routines introduced in IMS 11 (Early Initialization, CQS Event and CQS Structure Event) and the Restart exit which was introduced in IMS 10:

- A way to refresh exit routines online by using the new REFRESH USEREXIT command. Use the type-2 REFRESH USEREXIT command to bring new or modified exit routines online (and delete the old exit routines) without requiring that IMS be stopped and restarted. After the new or modified exit routines are online, any subsequent calls to those exit types will result in the execution of the new or modified exit routines. Only exit types that are specified on the TYPE parameter in the USER_EXITS section of the DFSDFxxx member of the IMS PROCLIB data set are eligible for refreshing. As part of processing the REFRESH USEREXIT command, IMS loads the new or modified exit routines and brings them online before deleting the old exit routines. If there is a failure in processing the command, the old exit routines remain online and are not deleted.

- A way to query information about certain exit routines by using the new QUERY USEREXIT command. Use the type-2 QUERY USEREXIT command to display information about the exit routines that are defined in the USER_EXITS section of the DFSDFxxx member of the IMS PROCLIB data set.



Serviceability Enhancements

Improved Serviceability

- **IMS Interactive Dump Formatter** is enhanced to:
 - Re-create the final portion of an IMS log from information in the dump
 - Function will be invoked via Dump Formatter menu
 - Log records will be extracted from dump and written to an output dataset
- **Eliminates the need to request the SLDS for diagnostic purposes reducing problem resolution time**
- **IMS Abend Dump Formatter is installed Dynamically**
 - No longer need to install **DFSAFMD0** in z/OS module IEAVADFM during IMS installation
 - New **DFSAFMX0** exit will be used

43 IMS 11 Overview | August 2009 © Copyright IBM Corporation 2009

IMS Interactive Dump Formatter is enhanced in IMS Version 11 to re-create the final part of an IMS log from the information that is available in an IMS dump, thus eliminating the need to request the final system log data set (SLDS) for diagnostic purposes. When an IMS customer is working with an IBM Software Support representative, frequently the customer creates a memory dump and the service representative needs IMS log data. To get this log data, the customer requests the final SLDS, compresses (terses) it, and transfers the file to IBM by using FTP. The Service person decompresses (unterses) the data set and then analyzes it. The enhancements built into the IMS Interactive Dump Formatter in IMS Version 11 can build a log data set from the log records that reside in the dump's log buffers, thus avoiding the requests for the SLDS and the extra work associated with these requests.

Syntax Checker & IVP Enhancements

Improved Manageability Improved Usability

- **Syntax Checker** ISPF application supports the following **PROCLIB** members which are new to IMS 11:
 - **CSLDIxxx** - ODBM Initialization member
 - **CSLDCxxx** – ODBM Configuration member
 - **DSPBIxxx** – DBRC Initialization member
- The **Installation Verification Program (IVP)** is enhanced to:
 - Support the **Open Database** enhancements
 - A new step will bring up the ODBM address space
- Improves the **manageability** of IMS release to release migration and the installation process


Reminder:
The Variable Export utility makes it easier to import the IVP variables from a prior release of IMS

44 | IMS 11 Overview | August 2009 | © Copyright IBM Corporation 2009

The Syntax Checker ISPF application assists with IMS release-to-release migrations by providing the ability to convert supported IMS.PROCLIB members from one release to the other.

The IVP is enhanced to support the Open Database enhancements. In a new step, the ODBM address space will be brought up.

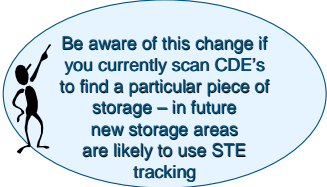
In IMS Version 10 and later, the Variable Export utility can be directly accessed as an option from the IVP Phase Selection panel, which makes it easier to import the IVP variables from a prior release of IMS.




Improved Availability

LSQA Storage Reduction

- The IMS internal storage managing service is enhanced to:
 - Use **64-bit storage** instead of 24-bit private LSQA for tracking certain types of storage allocations
 - For selected storage allocations, IMS no longer builds z/OS Contents Directory Entries (CDEs) control blocks
 - This should improve **availability** by reducing End of Memory (EOM) type system abends that require an IPL to resolve



45

IMS 11 Overview | August 2009
© Copyright IBM Corporation 2009

The IMS internal storage managing service, IMODULE, is updated to use 64-bit storage, instead of 24-bit private storage, for certain IMS functions. This enhancement should reduce the occurrences of end-of-memory (EOM) type IMS abends that require an IPL of the z/OS system to resolve.

IMS provides an internal IMS service called IMODULE, which IMS modules use to allocate and release storage, and load and delete modules. Today, IMODULE keeps track of both storage areas and modules by building a control block structure that is defined by z/OS – CDEs and related blocks. These z/OS blocks must architecturally reside in 24-bit private storage (LSQA storage). 24-bit private storage is a limited resource (limited to a maximum of sixteen megabytes, but more practically in the range of eight to ten megabytes). With the large size of today's address spaces, it is possible to allocate more storage areas in 31-bit storage than it is possible to track using 24-bit CDE structures. When this happens in the IMS CTL or DLI address spaces, it is often the case that z/OS itself cannot get enough storage to perform recovery/termination manager (RTM) processing for the address space. This leads to "end-of-memory" (EOM) type abends, where IMS is unable to cleanup its allocated common storage. This often requires a z/OS IPL to clear up the "orphaned" common storage so that IMS can be restarted on that z/OS.

This line item creates a new internal IMS service and control block structure for tracking storage. The use of this new tracking service for any given storage request is controlled by coding a new optional parameter on the IMODULE macro. The storage tracking elements (STEs) are built in 64-bit storage for the CTL region and DLI region (no LSQA storage is used). Other region types continue to use the CDE tracking technique.

The IMODULE GETMAIN storage requests for an internal type of IMS block called a "BCB IPAGE" are changed to use the new parameter to request tracking by IMS STEs, rather than by z/OS CDEs. BCB IPAGE storage is heavily used for many IMS internal processes and control block structures. Often, run-away conditions lead to the allocation of many IPAGES of storage, and can lead to the out-of-storage and end-of-memory conditions. Thus, moving this one type of IMS storage to be tracked by STEs should address many of the common scenarios that end up leading to EOM and z/OS IPL situations. However, note that the design of the new IMODULE service is such that other IMS storage areas could easily be changed in the future to be tracked with STEs, should the need arise.

For this line item, only DFSBCB IPAGE storage will be changed to be tracked with STEs; however, future use of STE tracking – particularly with new storage areas – will certainly occur.

KBLA Enhancements

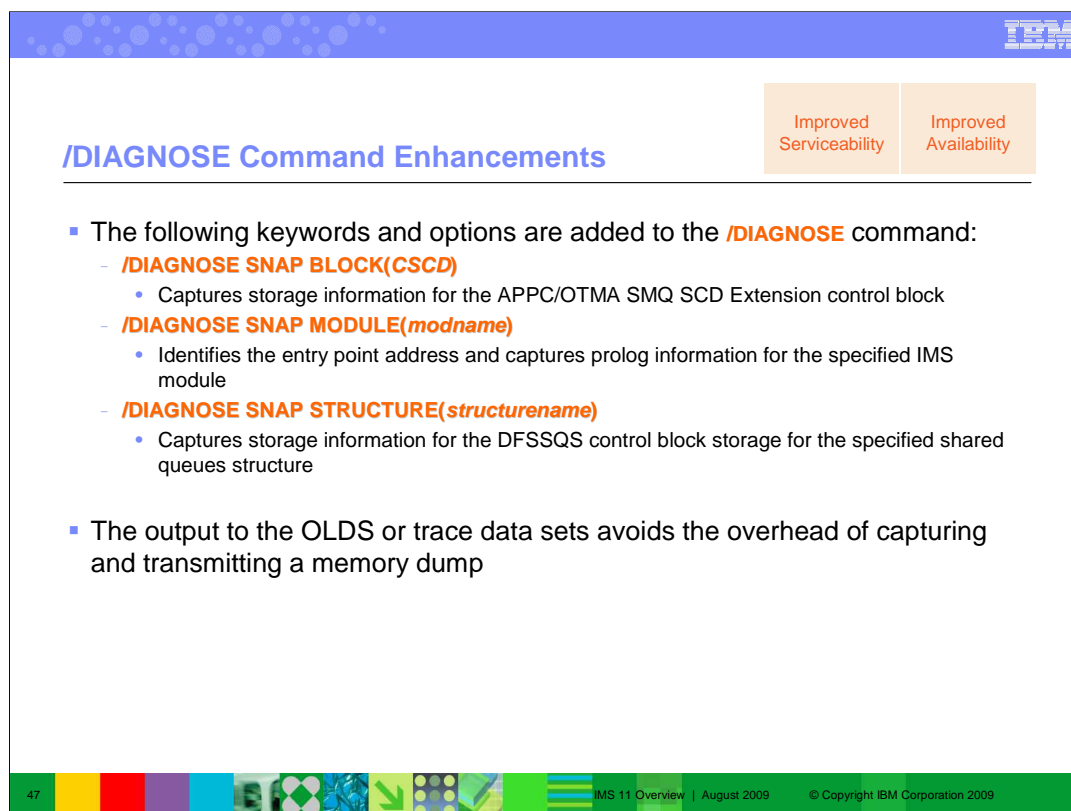
Improved Manageability Improved Usability

- **Knowledge Based Log Analysis (KBLA)** has the following small **usability** enhancements:
 - The ability to **scroll ISPF panels** to view data hidden due to exceeding 24 lines on a panel, or in 'split screen' mode is provided
 - The 'Define KBLA Environment' panel and JCL is modified to allow the user to **allocate data sets** which **span multiple volumes**
 - Should avoid **'out of space'** abends during log processing

46 IMS 11 Overview | August 2009 © Copyright IBM Corporation 2009

When you upgrade the KBLA utilities to support IMS Version 11, the following functions are added:

- Data entry panel scrolling
- Multi-volume KBLA output data set allocation



/DIAGNOSE Command Enhancements

Improved Serviceability

Improved Availability

- The following keywords and options are added to the **/DIAGNOSE** command:
 - **/DIAGNOSE SNAP BLOCK(CSCD)**
 - Captures storage information for the APPC/OTMA SMQ SCD Extension control block
 - **/DIAGNOSE SNAP MODULE(modname)**
 - Identifies the entry point address and captures prolog information for the specified IMS module
 - **/DIAGNOSE SNAP STRUCTURE(structurename)**
 - Captures storage information for the DFSSQS control block storage for the specified shared queues structure
- The output to the OLDS or trace data sets avoids the overhead of capturing and transmitting a memory dump

The /DIAGNOSE command takes a snap shot of IMS system resources at any time without impacting availability. The output produced can be quickly transmitted to IBM Software Support, thus avoiding the overhead of capturing and transferring a memory dump. The following keywords and options are new to the /DIAGNOSE command:

- BLOCK(CSCD)** Captures storage information for the APPC/OTMA SMQ SCD Extension control block.
- MODULE(modname)** Identifies the entry point address and captures prolog information for the specified IMS module. The prolog information contains the current maintenance level for a module and can help you to determine if any maintenance is missing.
- STRUCTURE(structurename)** Captures storage information for the DFSSQS control block storage for the specified shared queues structure.



IBM Software Group



Transaction Manager Enhancements

Information Management software

© Copyright IBM Corporation 2009

Improved Manageability
Improved Usability

Type-2 Query TM Commands

- The **type-2 QUERY** commands are extended for **TM Resources**:
 - **QUERY LTERM** - used to query logical terminal (LTERM) information
 - **QUERY NODE** - used to query VTAM node or terminal information
 - **QUERY USER** - used to query ETO user or ISC subpool information
 - **QUERY USERID** - used to query user ID information
- **Filtering and wildcard support makes it easier to manage your IMSplex**
 - Output of several type-1 commands is consolidated into a single type-2 command

```
File Action Manage resources SPOC View Options Help
IMS11          IMS Single Point of Control
Command ==>>  QRY NODE NAME(NODE21) SHOW(GLOBAL,CONV,LTERM,STATUS)

Response for: QRY NODE NAME(NODE21) SHOW(GLOBAL,CONV,LTERM,STATUS)

Plex . . . . . Route . . . . . Wait . . . . .

Node   MbrName CC  Gbl  Lterm   ConvID ConvTran ConvStat Status
NODE21 IMS1   0  Y    LTERM21A 1  TRAN1A  CONVHELD CONVACT,STATIC,RM,RMACTIVE,RMOWNED
NODE21 IMS1   0  Y    LTERM21B 2  TRAN1A  CONVHELD
NODE21 IMS1   0  Y    3  TRAN1A  CONVHELD
NODE21 IMS1   0  Y    1  TRAN1A  CONVHELD
NODE21 IMS1   0  Y    2  TRAN1A  CONVHELD
NODE21 IMS1   0  Y    3  TRAN1A  CONVACTV
```

49
IMS 11 Overview | August 2009
© Copyright IBM Corporation 2009

This enhancement provides type-2 command support for data communications resources. It provides QUERY command support for NODE, LTERM, USER and USERID. This enhancement improves ease of use because the QUERY commands are enhanced commands which allow wildcards and filters that provide options not available to type-1 commands. The new TM type-2 QUERY commands are an alternative, not a replacement, for existing type-1 commands. The new commands enable you to manage your IMSplex environment more efficiently.

The new TM type-2 QUERY commands and their usage are:

QUERY LTERM - used to query logical terminal (LTERM) information

QUERY NODE - used to query VTAM node or terminal information

QUERY USER - used to query ETO user or ISC subpool information

QUERY USERID - used to query user ID information

The diagnostic information that can be generated by using a single TM type-2 command is comparable to the diagnostic information that can be generated by consolidating the output of several existing type-1 commands. In general, equivalent type-1 and type-2 commands operate in the same manner. Issuing type-2 commands requires a Common Service Layer with a minimum of an OM and SCI.

Query Node example:

There are two IMS systems in the IMSplex: IMS1 and IMS2.

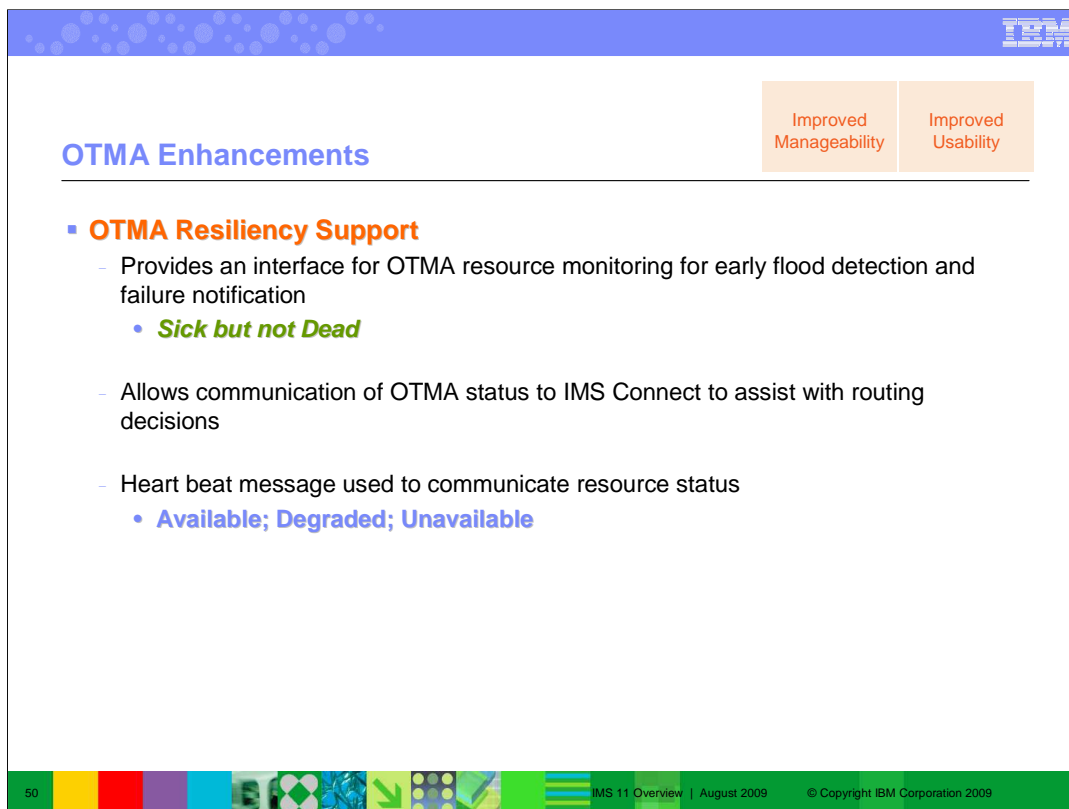
The Resource Manager (RM) is maintaining status in the resource structure (STM=YES), and Shared queues are enabled.

IMS1 is the command master, and because SHOW(GLOBAL) was specified, IMS1 is the only system that processes the command.

Any other IMS ignores the command (RC=4, RSN=x1000). NODE21 exists in the resource structure. IMS1 displays a global line which shows that the node is active in the IMSplex (RM, RMACTIVE, and RMOWNED status), and has a conversation active. There are two lterms assigned to the node, and are displayed on separate output lines. There are three conversations associated with the node, and are displayed on separate output lines.

IMS1 displays an additional global line for each assigned LTERM, so 2 lines are displayed for the 2 lterms LTERM21A and LTERM21B.

IMS1 displays an additional global line for each active conversation for the node. 3 lines are displayed, because the node has 2 held conversations and 1 active conversation.



OTMA Enhancements

- **OTMA Resiliency Support**
 - Provides an interface for OTMA resource monitoring for early flood detection and failure notification
 - *Sick but not Dead*
 - Allows communication of OTMA status to IMS Connect to assist with routing decisions
 - Heart beat message used to communicate resource status
 - **Available; Degraded; Unavailable**

50 | IMS 11 Overview | August 2009 | © Copyright IBM Corporation 2009

OTMA Resiliency Support - When an OTMA becomes unable to effectively process the work submitted by an OTMA client (that is, sick, but not dead), the client application is usually unaware of the situation and continues sending work to the IMS. This situation could lead to flooded transactions and can result in an IMS outage that disrupts all client applications.

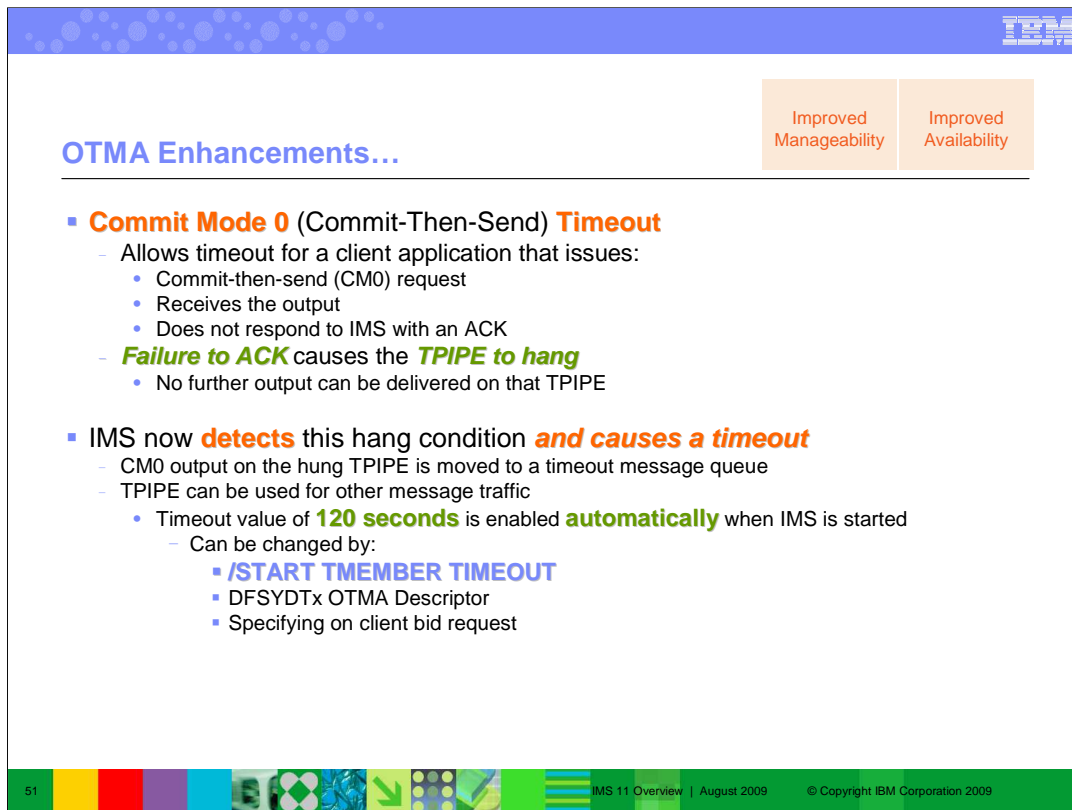
OTMA Resiliency Support provides a client-server protocol that allows OTMA to detect the sick-but-not-dead situations related to OTMA processing, identify the resources involved, take server actions if needed, and send out a protocol message to the client. This protocol message would then be processed by an OTMA client, such as IMS Connect or MQSeries, where the transaction can be redirected to another IMS. If there is no other suitable IMS to reroute the transaction to, the OTMA client could mark the IMS system "temporary unavailable, try again later".

IMS Connect utilizes this OTMA function by processing these protocol messages, updating its data store entry, and recording new data store events for warning and severe status so that IMS Connect vendor applications and user exits can access the information and redirect the transaction requests to a different IMS if needed.

Example 1: IMS Connect submits a lot of transactions to OTMA. However, OTMA in IMS has experienced a slow-down issue and failed to process the transaction in a timely manner. Once OTMA detects that the message flood level for the IMS Connect reaches 80% of the threshold, the OTMA resource monitor will send an action message to inform IMS Connect. In addition, OTMA periodically sends out a "heart beat" message with the resource info. to IMS Connect. In this case, the "heart beat" message will indicate a warning condition with the flood status if the flood condition still exists for the IMS Connect.

Example 2: Both IMS Connect and MQSeries submit transactions to IMS. If the total number of un-completed send-then-commit (CM1) transactions reaches the warning level, OTMA resource monitor will send out an action message to inform them. Also, a heart-beat message which includes any flood control status will be sent out later.

Example 3: Due to the potential I/O hung or IMS TCB hung conditions, OTMA transactions may not be able to be enqueued and the OTMA transaction message block (YTIB) may not even be created. However, an OTMA client such as IMS Connect may not know this condition and continues to submit transactions to IMS OTMA via the XCF interface. OTMA has been enhanced to detect the number of unprocessed OTMA XCF messages and send out an action message if a high number of XCF messages exist in OTMA.



OTMA Enhancements...

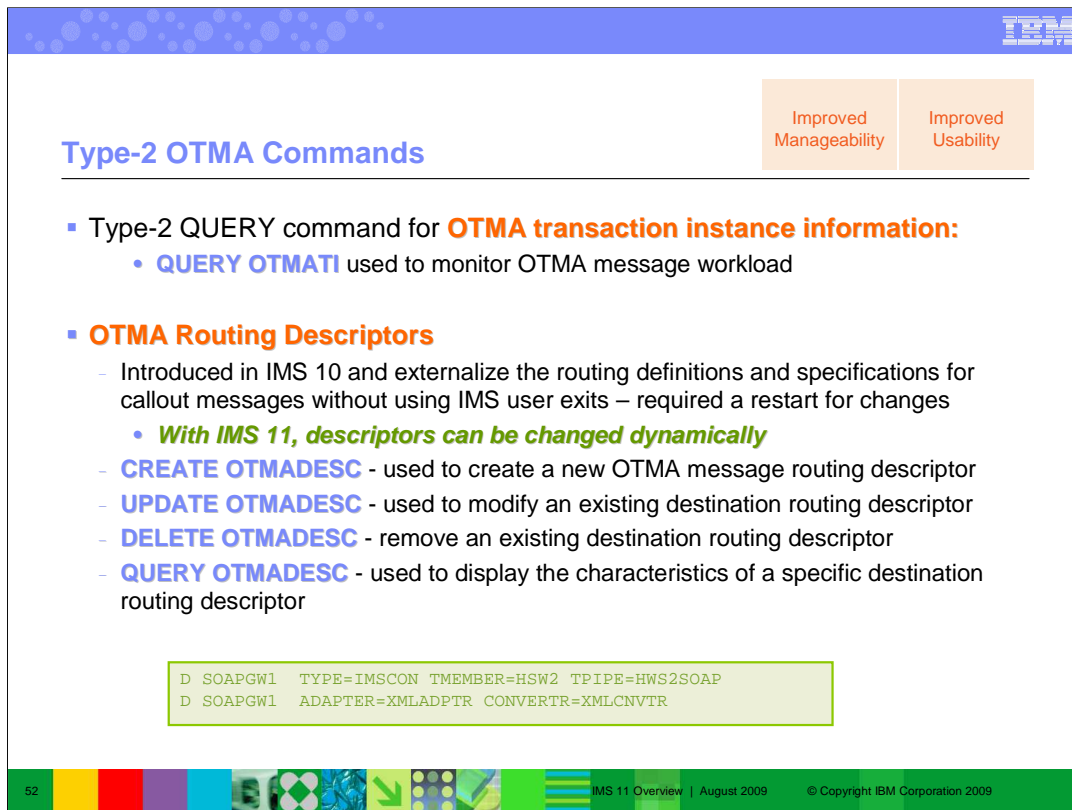
Improved Manageability Improved Availability

- **Commit Mode 0 (Commit-Then-Send) Timeout**
 - Allows timeout for a client application that issues:
 - Commit-then-send (CM0) request
 - Receives the output
 - Does not respond to IMS with an ACK
 - **Failure to ACK** causes the **TPIPE to hang**
 - No further output can be delivered on that TPIPE
- IMS now **detects** this hang condition **and causes a timeout**
 - CM0 output on the hung TPIPE is moved to a timeout message queue
 - TPIPE can be used for other message traffic
 - Timeout value of **120 seconds** is enabled **automatically** when IMS is started
 - Can be changed by:
 - **/START TMEMBER TIMEOUT**
 - DFSYDTx OTMA Descriptor
 - Specifying on client bid request

51 IMS 11 Overview | August 2009 © Copyright IBM Corporation 2009

If an ACK or NAK (positive or negative acknowledgement) is not received, commit mode 0 TPIPEs could hang and become unusable. Lots of output messages could then be queued to the TPIPE. The time-out function is needed for CM0 TPIPEs so that the TPIPE can be usable. This is similar to the existing CM1 Timeout function.

If a client application receives commit-then-send (CM0) output via an OTMA transaction pipe (TPIPE) and fails to respond with an acknowledgement (ACK), the TPIPE in IMS will hang and the subsequent output on the TPIPE cannot be delivered. The IMS V11 OTMA CM0 ACK timeout enhancement will detect this hang condition and take the time-out action so that the CM0 output on the hung TPIPE can be moved to a time-out message queue and the rest of the output on the TPIPE can continue to flow.



Type-2 OTMA Commands

Improved Manageability Improved Usability

- Type-2 QUERY command for **OTMA transaction instance information**:
 - **QUERY OTMATI** used to monitor OTMA message workload
- **OTMA Routing Descriptors**
 - Introduced in IMS 10 and externalize the routing definitions and specifications for callout messages without using IMS user exits – required a restart for changes
 - **With IMS 11, descriptors can be changed dynamically**
 - **CREATE OTMADESC** - used to create a new OTMA message routing descriptor
 - **UPDATE OTMADESC** - used to modify an existing destination routing descriptor
 - **DELETE OTMADESC** - remove an existing destination routing descriptor
 - **QUERY OTMADESC** - used to display the characteristics of a specific destination routing descriptor

```
D SOAPGW1  TYPE=IMSCON TMEMBER=HSW2 TPIPE=HWS2SOAP
D SOAPGW1  ADAPTER=XMLADPTR  CONVERTR=XMLCNVTR
```

52 IMS 11 Overview | August 2009 © Copyright IBM Corporation 2009

The OTMA type-2 QUERY OTMATI command for monitoring workload enables you to view a summary of the number of messages in the OTMA send-then-commit (CM1) message queue. You can retrieve this information by issuing the command QUERY OTMATI and include parameters to filter the result information. If the SHOW parameter is used, the individual characteristics of each transaction instance block are displayed. You can view the length of time that a transaction instance has been in existence and the correlation ID of the input message. This information can help you determine if there are problems in processing OTMA input messages.

Prior to IMS Version 11, if you wanted to make changes to the destination routing descriptors you had to wait for a scheduled outage or for IMS to be restarted. The OTMA type-2 commands for destination routing descriptors allow you to change the descriptors dynamically without interruption to a running IMS instance. You can dynamically add, update, delete and query destination routing descriptors while IMS is actively running by issuing these commands: CREATE OTMADESC, UPDATE OTMADESC, DELETE OTMADESC, and QUERY OTMADESC. In IMS Version 10, the definition of destination routing descriptors in the DFSYDTx PROCLIB member data set had to be entered in a specific order. In IMS Version 11, you no longer have to be concerned with order. You can input the entries in any order.

Transaction Expiration Enhancements


Improved
Manageability

Improved
Usability

- Allows you to set **transaction expiration values (in seconds)**
 - IMS can **avoid processing** transactions **where the response is no longer needed**

- **Expiration** can be specified **by transaction**:
 - In the message prefix (for OTMA messages)
 - By using the following type-2 commands
 - CREATE TRAN
 - CREATE TRANDESC
 - UPDATE TRAN
 - UPDATE TRANDESC
 - By specifying EXPRTIME on the TRANSACT macro
 - By including an expiration time for the transaction it creates in the Destination Creation exit (DFSINSX0)

- If an expiration value is set it is checked:
 - By IMS when an application issues its first GU call
 - For OTMA, when the message is received from XCF and again before the message is enqueued to IMS
 - If the time specified in the expiration value has passed, **IMS discards the transaction**
 - Abend 0243 will be issued at application Get Unique (GU) time if the transaction has expired



IMS Connect is
enhanced
to take advantage of
Transaction
Expiration

53

IMS 11 Overview | August 2009

© Copyright IBM Corporation 2009

IMS Version 11 has the ability to interrogate an expiration time associated with transactions and discard (not process) the transactions if the time specified has passed. By discarding transactions whose expiration time has passed, processing costs and CPU cycles are not spent for the unwanted transactions. The expiration time is specified in the TRANSACT macro, type-2 commands, or the OTMA message prefix. OTMA, IMS Connect, some type-2 commands, and the RDDS Extraction utility are enhanced to support the transaction expiration function. (The Resource Definition Data Set (RDDS) Extraction utility is enhanced to recognize the new transaction expiration attribute so that transactions with the attribute specified can be imported from or exported to the RDDS.)

Shared Queues Affinity Routing Enhancement

Improved Manageability Improved Availability

- Exit **DFSMSCEO** is enhanced as follows:
 - To allow routing of **Synchronous APPC/OTMA Shared Message Queue** transactions to a **back end IMS** when:
 - Resource Recovery Service (RRS) or
 - APPC/OTMA Shared Queues support (AOS=y) are not active
 - *This resolves the APPC/OTMA affinity restriction for some customers*
 - The IMSID affinity field in DFSMSCEP user parameter list is increased from 4 to 8 bytes
 - To support the XRF RSENAME
 - DFSMSCEB and DFSMSCEP control block sizes are increased to provide more room for service and future enhancements
- **QUERY TRAN** command has a new status display value (STATUS=LCLAFFIN) if the transaction is registered for affinity
 - *Makes it easier to manage affinity routing*

DFSMSCEO sample exit has examples of back end affinity routing

54 IMS 11 Overview | August 2009 © Copyright IBM Corporation 2009

The capability and usability of the TM and MSC Message Routine and Control User exit routine (DFSMSCEO) is enhanced for IMS Version 11. The affinity support in DFSMSCEO is enhanced in the following ways:

- The DFSMSCEO user exit can route APPC/OTMA synchronous messages to a back-end IMS when the Resource Recovery Service (RRS) or the APPC/OTMA Shared Queues support (AOS) are not active. This resolves the APPC/OTMA affinity restriction for some customers.
- The IMSID affinity field in DFSMSCEP user parameter list is increased from four to eight bytes to support XRF.
- A new status display value is added to the QUERY TRAN command that displays if the transaction is registered for affinity (STATUS=LCLAFFIN). This makes affinity routing easier to manage.
- The DFSMSCEB and DFSMSCEP control block sizes are increased to provide more room for service and future enhancements to the DFSMSCEO user exit.



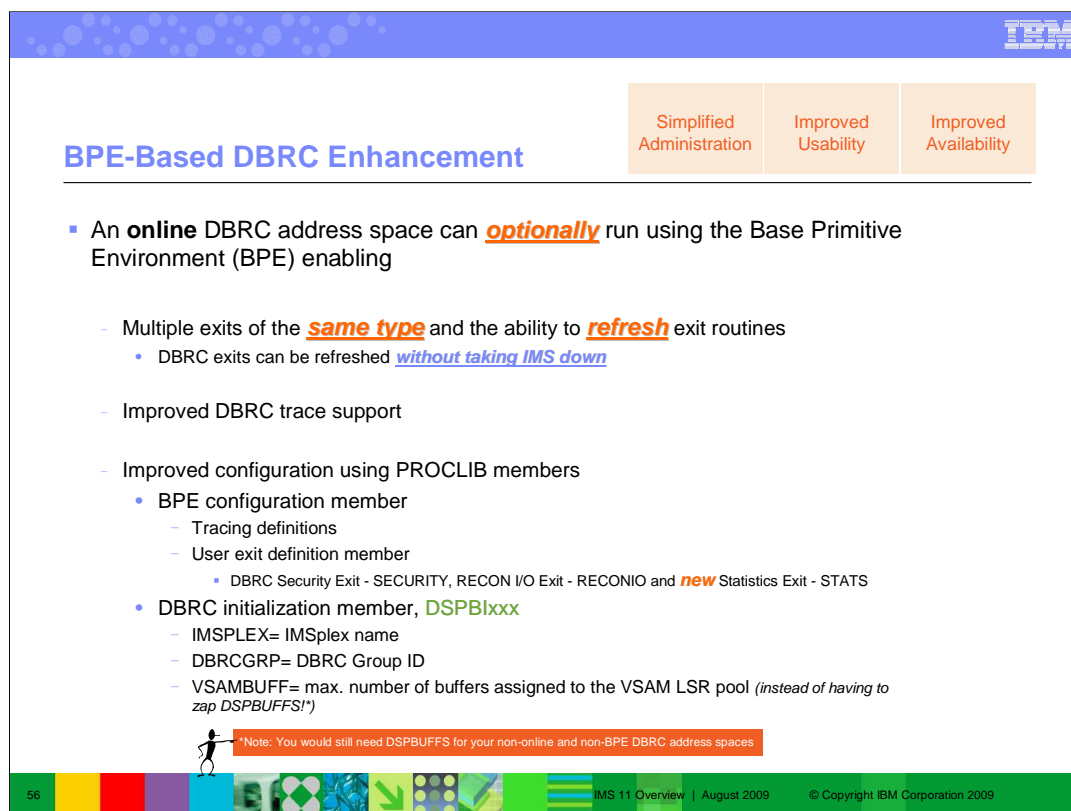
IBM Software Group



DBRC Enhancements

Information Management software

© Copyright IBM Corporation 2009



BPE-Based DBRC Enhancement

- Simplified Administration
Improved Usability
Improved Availability
- An **online** DBRC address space can **optionally** run using the Base Primitive Environment (BPE) enabling
 - Multiple exits of the **same type** and the ability to **refresh** exit routines
 - DBRC exits can be refreshed **without taking IMS down**
 - Improved DBRC trace support
 - Improved configuration using PROCLIB members
 - BPE configuration member
 - Tracing definitions
 - User exit definition member
 - DBRC Security Exit - SECURITY, RECON I/O Exit - RECONIO and **new** Statistics Exit - STATS
 - DBRC initialization member, **DSPBlxxx**
 - IMSPLEX= IMSplex name
 - DBRCGRP= DBRC Group ID
 - VSAMBUFF= max. number of buffers assigned to the VSAM LSR pool (*instead of having to zap DSPBUFFS!*)

*Note: You would still need DSPBUFFS for your non-online and non-BPE DBRC address spaces

In order to take advantage of several of the services the Base Primitive Environment provides, we are adding an option to start DBRC on a BPE environment. With BPE, we can provide improved user exit management and trace support. This support is totally optional. You either start DBRC with your current procedures, or you can use a new startup JCL member that uses BPE. This allows you to convert your JCL over time.

Unlike the current DBRC start-up JCL, the new startup JCL and procedures will have a PROCLIB DD statement which allows online DBRC to have proclib members to support certain functional capabilities. For example, a BPE configuration member can specify trace options as well as user exit members. The user exit member defines which user exits should be used by this DBRC. BPE provides a much better user exit interface than currently provided through DBRC. For one thing, you have the ability to update user exits without shutting IMS/DBRC down. Also, you can have more than one exit routine called.

SECURITY type exit replaces DSPDCAX0. The parameter list is mapped by DSPDCABK.

RECONIO type exit replaces DSPCEXT0. The parameter list is mapped by DSPRIOX.

STATS type exit is new. The parameter lists are mapped by DSPBSTX, DSPBST1 and DSPBST2.

The STATS exit includes a lot of information on specific types of work done by DBRC, such as the number of individual DBRC requests that are issued, number of VSAM requests made by DBRC and average time it takes to process a DBRC request.

We now allow you to define a DBRC Security Exit and RECON I/O Exit through the BPE EXITDEF proclib members. If specified, they would be used instead of the current DSPDCAX0 and DSPCEXT0. The new exits would have a slightly different interface as the current exits. Basically, we plan on using an interface that is similar to other CSL user exits, and put the current parameters passed to the exits at the end of the new parameters in the same order. We plan on providing a new Statistics Exit that will provide such things as average request times, average wait times (for reserve, and other things). We also plan on taking advantage of BPE tracing services. For example, BPE external trace could be used for the online region instead of GTF tracing.

A DBRC initialization member can be used to specify things like the DBRCGRP and IMSPLEX name. We are also adding the ability to specify VSAM LSR buffer usage for the online DBRC instead of requiring you to zap DSPBUFFS. Note: you would still need DSPBUFFS for your non-online DBRC image and non-BPE.

RECON Security Override

Simplified Administration Improved Usability

- For **non-production** copies of the RECON
 - You can optionally override the **inherited** security level
 - Allowing access for testing or diagnostic purposes
 - New optional sub parameter added to the CMDAUTH keyword on INIT.RECON and CHANGE.RECON commands

CMDAUTH(SAF|EXIT|BOTH|NONE,safh|q{,rcnqual})

- Where: *rcnqual* must be a substring of the COPY1 RECON DSN

The RECON Header & DBRC API RECON Status Block are increased to support this option

57 IMS 11 Overview | August 2009 © Copyright IBM Corporation 2009

Non-production copies of the RECON data set inherit the security level of the production RECON from which they were copied. This security characteristic poses a problem when someone who does not have the original authority level needs to access the RECON data set, such as for testing or debugging purposes. To solve this access problem:

- The CMDAUTH keyword on the INIT.RECON and CHANGE.RECON commands is enhanced in IMS Version 11 with the option to override DBRC security for non-production copies of the RECON data set.
- The RECON header record (DSPRCNRC) is enhanced with the new RCNQUAL field, RCNCMDRNQ, which tells DBRC if the RECON data set is a copy and whether command authorization should be enforced or not.
- The new apqrc_CmdRNQ field that contains the RCNQUAL value is added to the DBRC API RECON status block (DSPAPQRC).

Unconditional Deletion of PRILOG Information

Simplified Administration

Improved Usability

- A new **CLEANUP.RECON** command
 - To **delete obsolete or expired** recovery-related information from the RECON
 - In cases where PRILOG record compression was unable to delete inactive entries
 - Can be issued from DSPURX00 or through the DBRC Command API


```

          CLEANUP.RECON {RETPRD(time_interval) | TIME(time_stamp)}
                    {DBRANGE(firstdb,lastdb)} {DBONLY}
                    {LASTIC} {LISTDL|NOLISTDL}
          
```
 - Parameters on the command enable you to specify:
 - Either a **retention period** or an **absolute time** to which recovery-related information and log information is to be deleted
 - The **databases** that are associated with the inactive information
 - Whether to just delete **database recovery-related information** and not log information
 - That information that pertains to the **last available image copy** for a database can be deleted
 - Whether data about the deleted information should be included in the **SYSPRINT** listing

Obviously use this with caution!
Should not be necessary if databases are image copied regularly

58

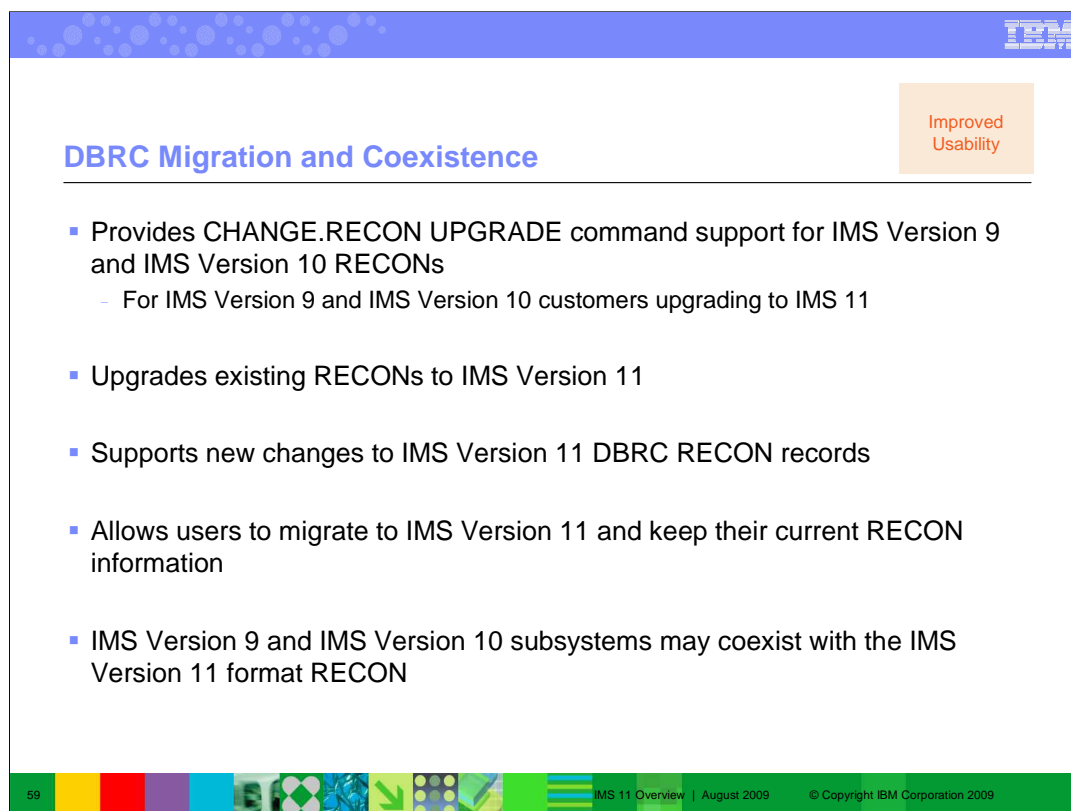
IMS 11 Overview | August 2009

© Copyright IBM Corporation 2009

You can use a new DBRC command, CLEANUP.RECON, to delete old or expired recovery-related information from the RECON data set. The RECON data set contains information that is needed to recover databases. The PRILOG record in the RECON data set can grow to be very large. PRILOG record compression, which is the automatic deletion of inactive data set entries in the PRILOG record, sometimes cannot take place for various reasons. In these cases, the information must be deleted manually.

With the appropriate authorization, you can use the CLEANUP.RECON command to remove obsolete information from the RECON data set. The CLEANUP.RECON command can be issued from within the Database Recovery Control utility (DSPURX00) or embedded in a DBRC Command API request. The parameters on the CLEANUP.RECON command enable you to specify:

- Either a retention period or an absolute time to which recovery-related information and log information is to be deleted. It is recommended to use small deletion periods when using this cleanup command. The smaller the deletion periods, the less time the RECON will be held.
- The databases that are associated with the inactive information
- Whether to just delete database recovery-related information and not log information
- That information that pertains to the last available image copy for a database can be deleted
- Whether data about the deleted information should be included in the SYSPRINT listing.



The slide features a blue header with a pattern of white dots. In the top right corner, there is a small logo and the text 'IBM Business Partner'. Below the header, the title 'DBRC Migration and Coexistence' is displayed in blue. To the right of the title is an orange box with the text 'Improved Usability'. The main content is a bulleted list of features. At the bottom of the slide, there is a green footer bar with a decorative pattern of colorful squares and icons on the left, and the text 'IMS 11 Overview | August 2009 © Copyright IBM Corporation 2009' on the right.

DBRC Migration and Coexistence

Improved Usability

- Provides CHANGE.RECON UPGRADE command support for IMS Version 9 and IMS Version 10 RECONS
 - For IMS Version 9 and IMS Version 10 customers upgrading to IMS 11
- Upgrades existing RECONS to IMS Version 11
- Supports new changes to IMS Version 11 DBRC RECON records
- Allows users to migrate to IMS Version 11 and keep their current RECON information
- IMS Version 9 and IMS Version 10 subsystems may coexist with the IMS Version 11 format RECON


59 IMS 11 Overview | August 2009 © Copyright IBM Corporation 2009

Migration to IMS Version 11 and coexistence support is provided for the DBRC RECONS. An IMS Version 9.1 RECON or an IMS Version 10 RECON may be upgraded to IMS Version 11. RECONS using Parallel RECON Access may also be upgraded. IMS Version 9.1 and IMS Version 10.1 may coexist with an IMS Version 11 RECON provided the respective coexistence Small Programming Enhancements are applied to the lower releases.

The RECON header, Change Accumulation execution record, and database records have been increased in size.



IBM Software Group

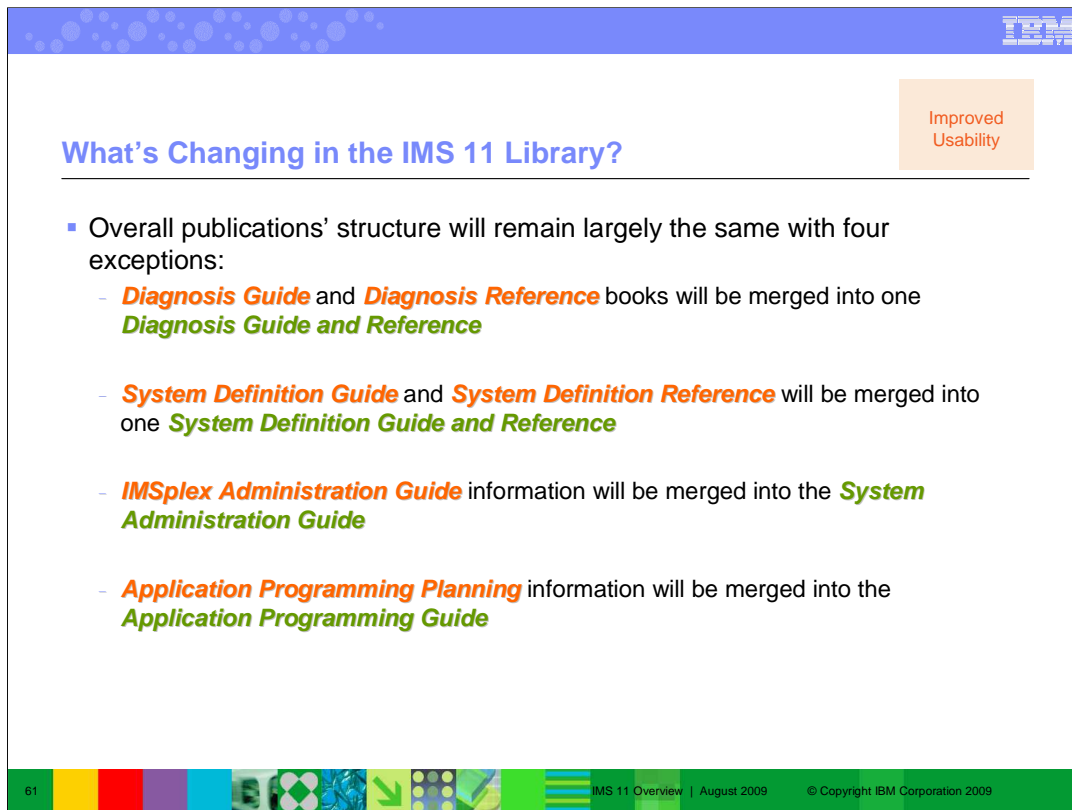


Miscellaneous

Information Management software



© Copyright IBM Corporation 2009



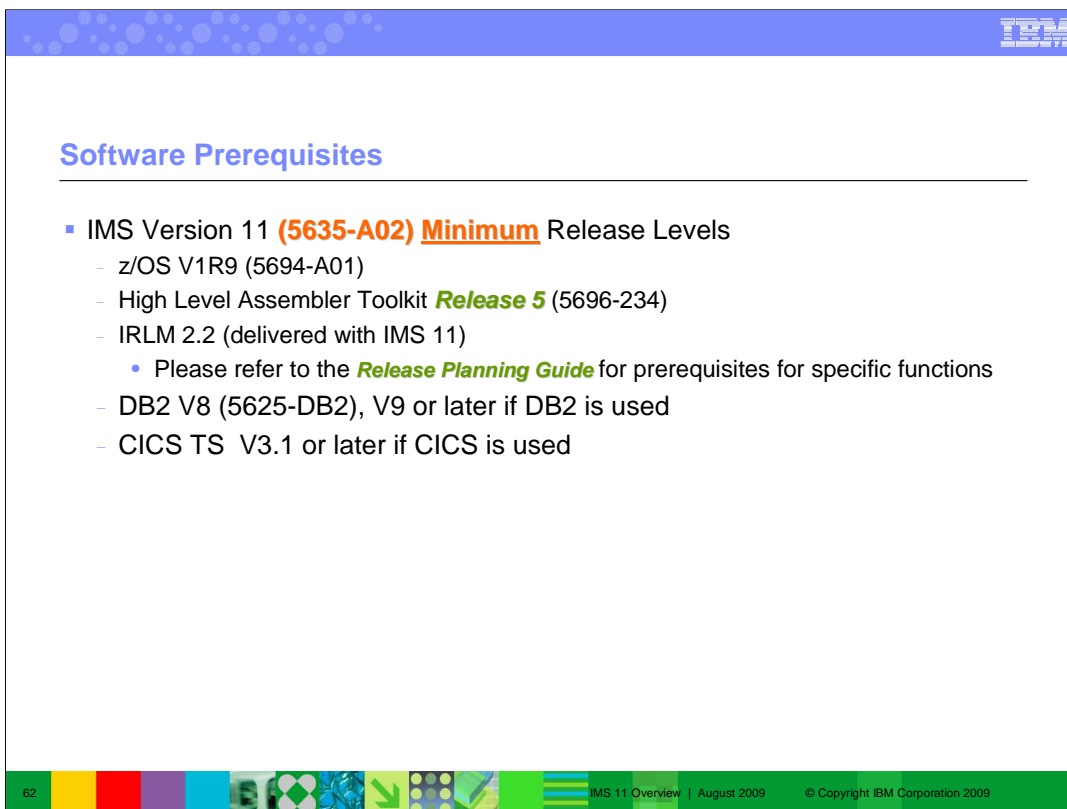
What's Changing in the IMS 11 Library?

Improved Usability

- Overall publications' structure will remain largely the same with four exceptions:
 - Diagnosis Guide* and *Diagnosis Reference* books will be merged into one *Diagnosis Guide and Reference*
 - System Definition Guide* and *System Definition Reference* will be merged into one *System Definition Guide and Reference*
 - IMSplex Administration Guide* information will be merged into the *System Administration Guide*
 - Application Programming Planning* information will be merged into the *Application Programming Guide*

61 | IMS 11 Overview | August 2009 | © Copyright IBM Corporation 2009

The structure of the library changed quite a lot with IMS Version 10. In Version 11 we will improve the content of information deliverables, but the overall library structure will remain largely the same, with the exception to the noted areas above.



Software Prerequisites

- IMS Version 11 **(5635-A02) Minimum** Release Levels
 - z/OS V1R9 (5694-A01)
 - High Level Assembler Toolkit **Release 5** (5696-234)
 - IRLM 2.2 (delivered with IMS 11)
 - Please refer to the **Release Planning Guide** for prerequisites for specific functions
 - DB2 V8 (5625-DB2), V9 or later if DB2 is used
 - CICS TS V3.1 or later if CICS is used

62 IMS 11 Overview | August 2009 © Copyright IBM Corporation 2009





IMS 11 operates under z/OS V1R9 (5694-A01) configurations, or subsequent versions, releases and modification levels unless otherwise stated, and requires the following minimum version, or release or modification levels:


- z/OS V1R9 (5694-A01) with DFSMS (a base element of z/OS V1R9).
- RACF (included in a separately orderable SecureWay Security Server feature of z/OS V1R9), or equivalent, if security is used.
- IBM High-level Assembler Toolkit (5696-234), a separately orderable feature of z/OS.
- Note, CICS 3.1 no longer supports VS COBOL (see CICS 3.1 materials). CICS 2.3 is planned to be out of service shortly after IMS 11 GA.

IBM, the IBM logo, and other marks are trademarks of International Business Machines Corporation. © Copyright International Business Machines Corporation 2009.

Migration and Coexistence

- Migration/coexistence is supported from IMS 9 & 10 to IMS Version 11

	IMS Version 11 Function	IMS Version 9 Coexistence APAR	IMS Version 10 Coexistence APAR	IMS Version 11 Function	IMS Version 9 Coexistence APAR	IMS Version 10 Coexistence APAR
	DBRC RECON Data sets	PK61582	PK61583	ODBA	PK66020	PK66022
	IMS Connect	PK24912 PK29938 PK00895 PK87088	None	OTMA	PK47172	None
	IMSplex	PK23402 PK32970 PK27280 PK30189	None	System Management Enhancements	PK30189	None
	Global Online Change	PK23402 PK32970	None			

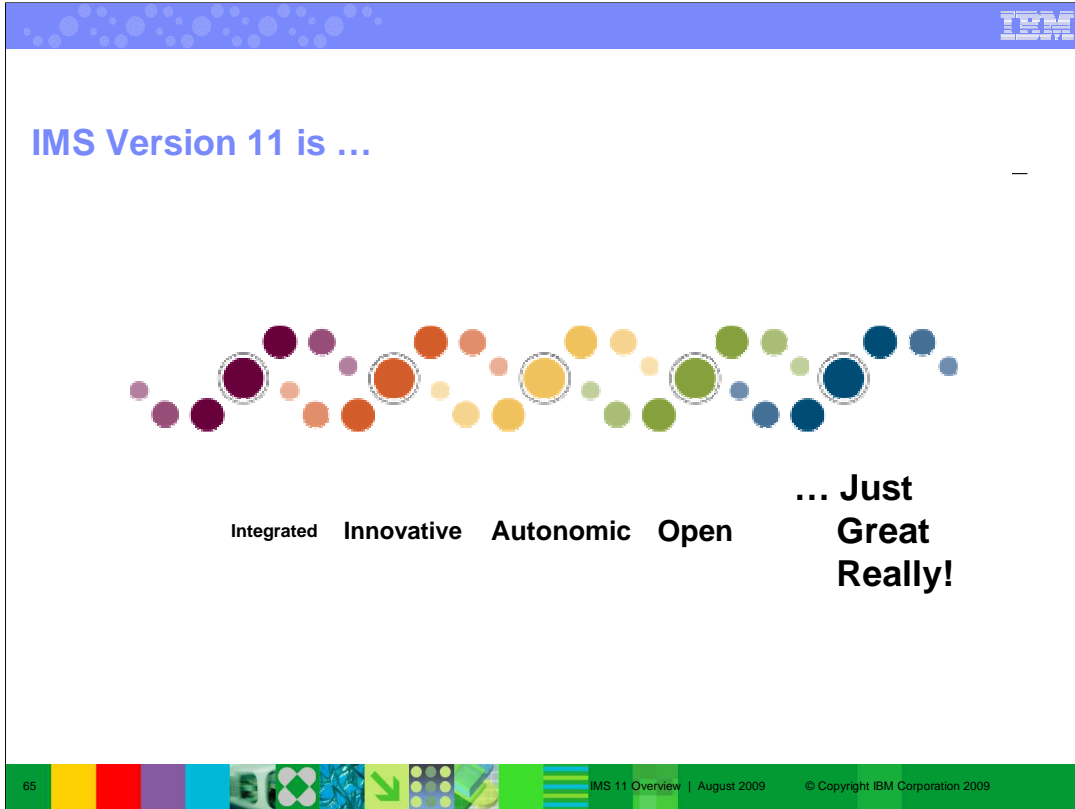
63

IMS 11 Overview | August 2009
© Copyright IBM Corporation 2009




For more information

- IMS 11 Release Planning Guide, GC19-2442
 - Available from the Information Management Software for z/OS® Solutions Information Center at <http://publib.boulder.ibm.com/infocenter/imzic>
- IMS 11 Fact Sheet, GC19-2451 (available October, 2008)
- IMS 11 Announcement Letters
 - EMEA – [ZP08-0416](#)
 - US – [208-258](#)
- IMS Family Web site:
 - ibm.com/ims
- IMS Version 9 will be Withdrawn from Marketing on 7th September 2009
 - See Announcement Letter [ZP09-0212](#) issued 2nd June 2009





IMS Version 11 is ...



Integrated Innovative Autonomic Open ... **Just Great Really!**

65 IMS 11 Overview | August 2009 © Copyright IBM Corporation 2009

The slide features a blue header bar with a pattern of white dots. Below the header, the text "IMS Version 11 is ..." is displayed in blue. A decorative graphic of various colored circles (purple, orange, yellow, green, blue) is arranged in a horizontal line. Below this graphic, the words "Integrated", "Innovative", "Autonomic", and "Open" are listed in black, followed by "..." and the phrase "Just Great Really!" in a larger, bold black font. At the bottom of the slide, there is a green footer bar containing the number "65", a series of small icons, the text "IMS 11 Overview | August 2009", and the copyright notice "© Copyright IBM Corporation 2009".