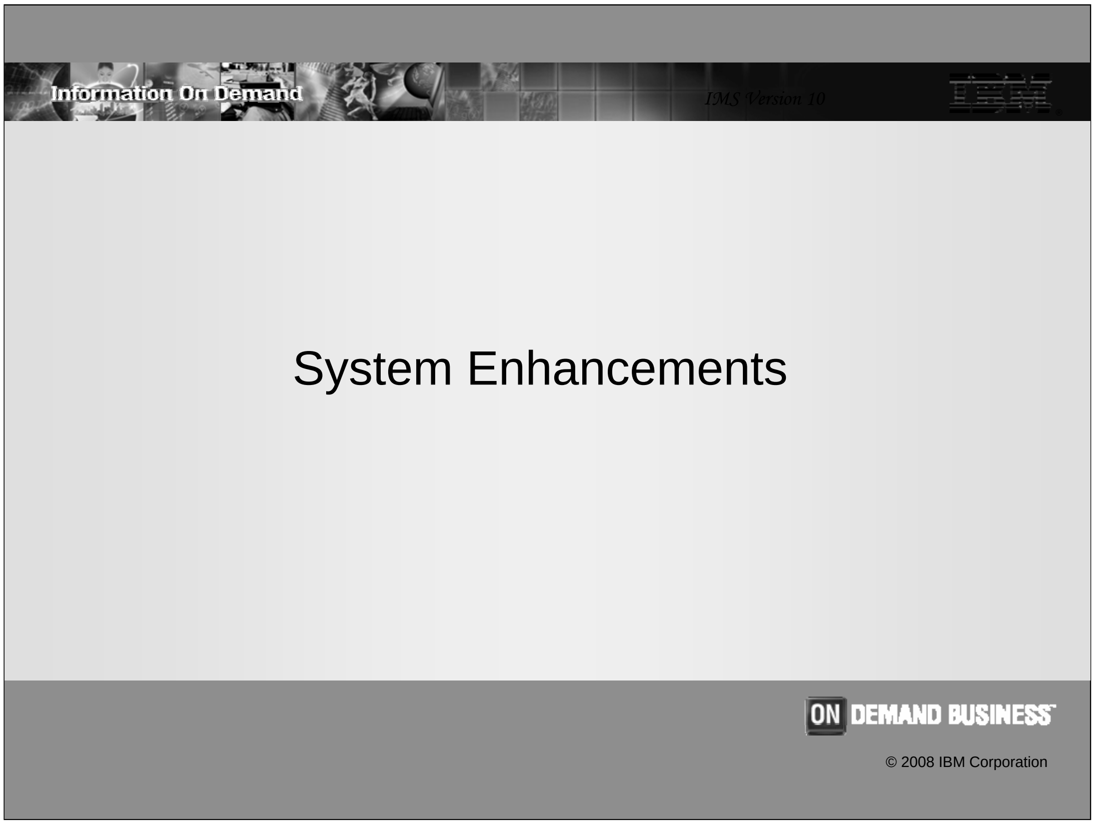

Information On Demand

IMS Version 10

System Enhancements

ON DEMAND BUSINESS™

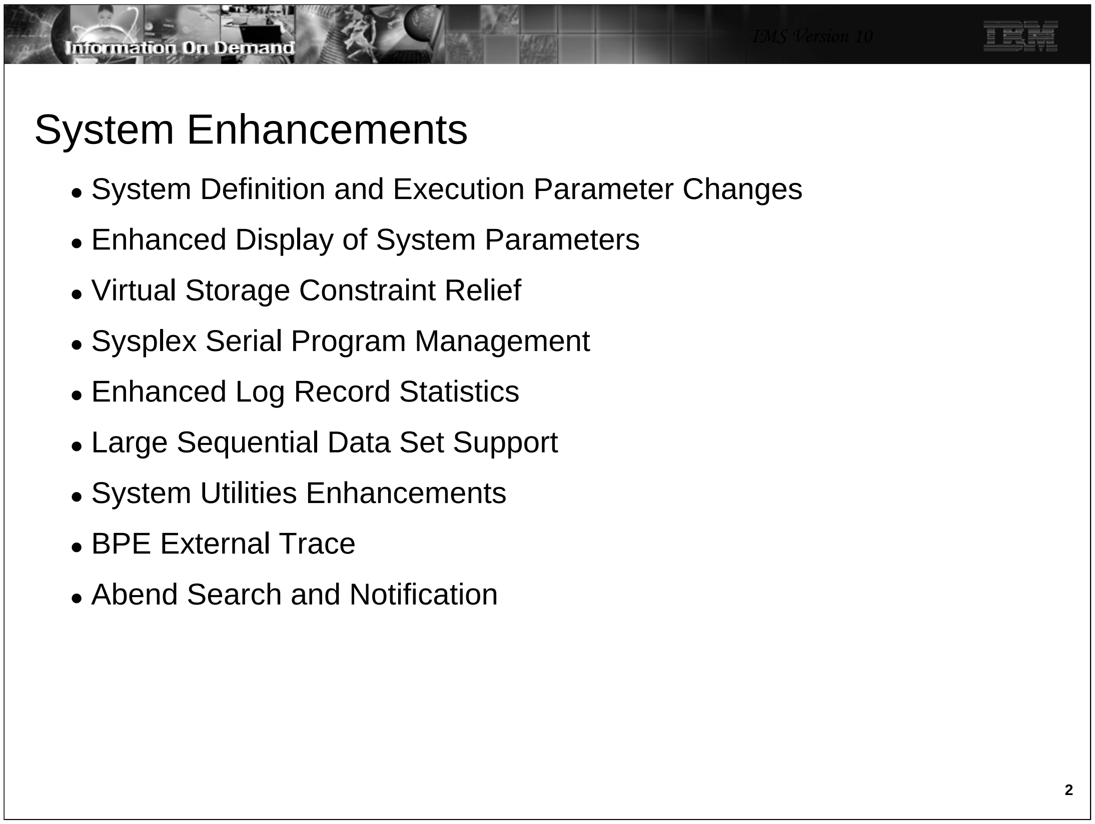

Information On Demand

IMS Version 10

IBM

System Enhancements

- System Definition and Execution Parameter Changes
- Enhanced Display of System Parameters
- Virtual Storage Constraint Relief
- Sysplex Serial Program Management
- Enhanced Log Record Statistics
- Large Sequential Data Set Support
- System Utilities Enhancements
- BPE External Trace
- Abend Search and Notification

2

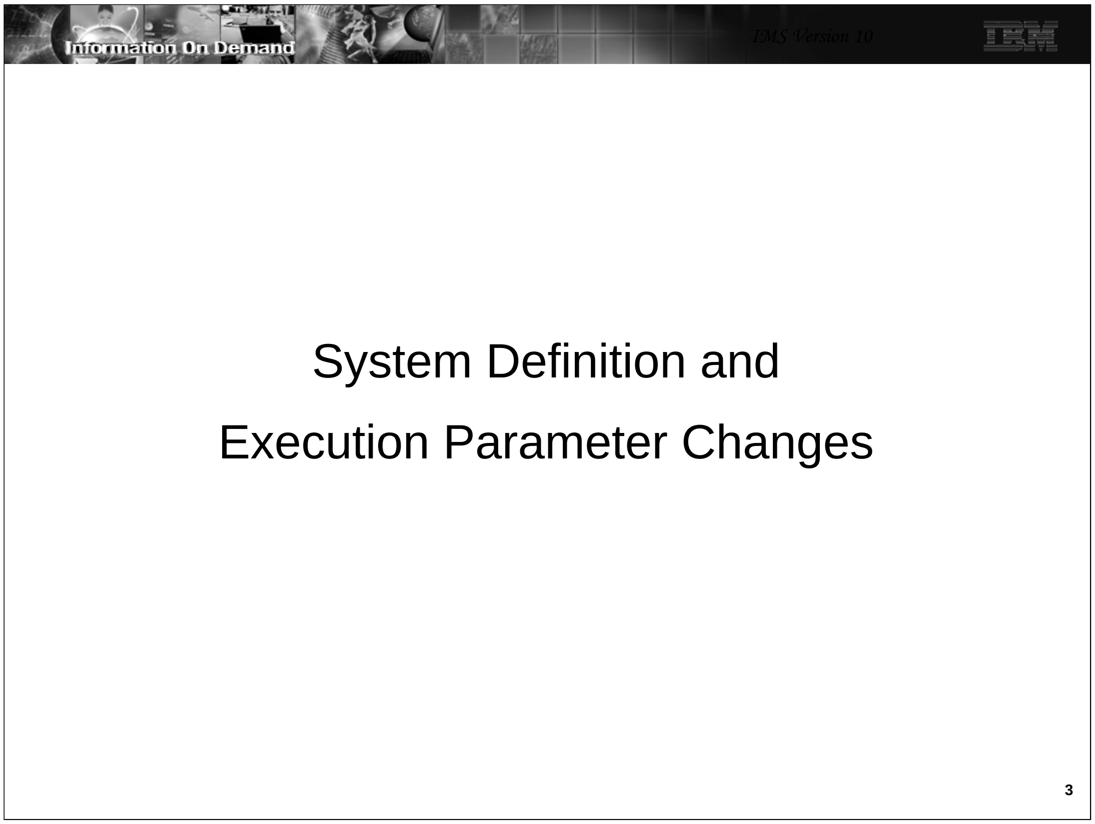
Information On Demand
IMS Version 10

System Definition and Execution Parameter Changes

3

This section covers changes to IMS system definition and execution parameters that affect functions available in previous releases. Definition and execution changes for new functions are discussed where the new functions are explained.

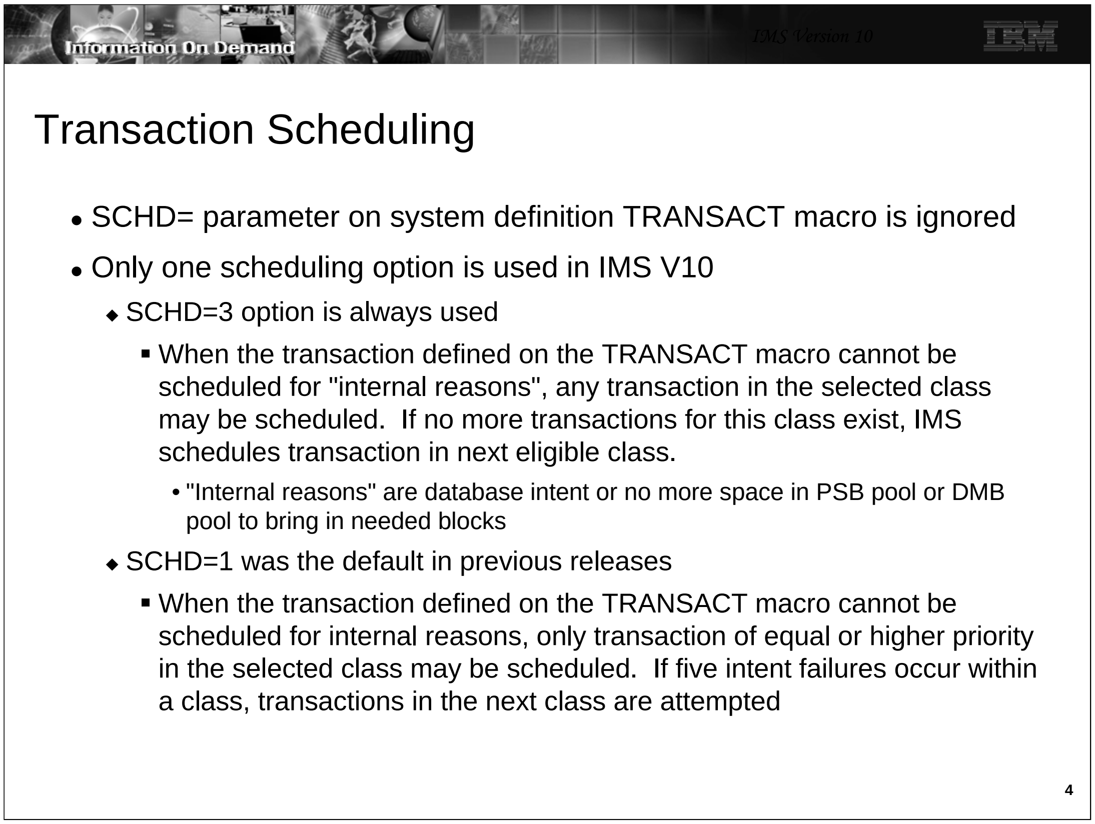
Information On Demand
IMS Version 10

Transaction Scheduling

- SCHD= parameter on system definition TRANSACT macro is ignored
- Only one scheduling option is used in IMS V10
  - SCHD=3 option is always used
    - When the transaction defined on the TRANSACT macro cannot be scheduled for "internal reasons", any transaction in the selected class may be scheduled. If no more transactions for this class exist, IMS schedules transaction in next eligible class.
      - "Internal reasons" are database intent or no more space in PSB pool or DMB pool to bring in needed blocks
  - SCHD=1 was the default in previous releases
    - When the transaction defined on the TRANSACT macro cannot be scheduled for internal reasons, only transaction of equal or higher priority in the selected class may be scheduled. If five intent failures occur within a class, transactions in the next class are attempted

4

In previous release of IMS, the SCHD= parameter on the system definition TRANSACT macro specified the scheduling option used when the transaction defined on the TRANSACT macro could not be scheduled for internal reasons (database intent or no more space in PSB pool or DMB pool to bring in needed blocks). IMS V10 always uses the SCHD=3 option. The default in previous releases was SCHD=1.

SCHD=1 specified that only transactions of equal or higher priority in the selected class would be scheduled. Five consecutive intent conflicts are allowed within a class before IMS starts scheduling the next eligible class.

SCHD=3 specified that any transaction in the selected class could be scheduled. IMS starts scheduling the next eligible class after attempting to schedule all the transactions in the current class

Two other options were available in previous releases.

SCHD=2 specified that only higher-priority transactions in the selected class could be scheduled.

SCHD=4 specified that IMS should skip to the next class and attempt to schedule the highest-priority transaction in that class.

Scheduling failures for the "internal reasons" should be rare. Intent conflicts only occur when PROCOPTs with the E (exclusive) option are used. The PSB and DMB pools should always be large enough to hold the currently required PSBs and DMBs.

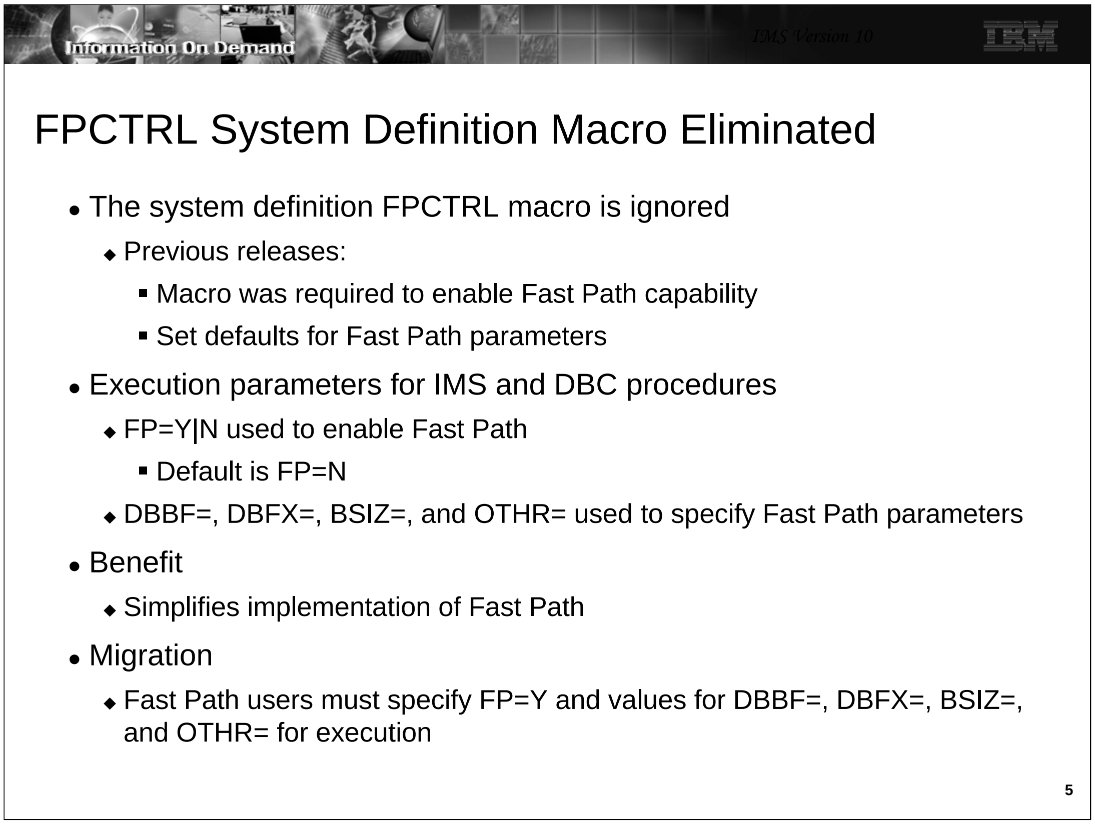

Information On Demand

IMS Version 10

FPCTRL System Definition Macro Eliminated

- The system definition FPCTRL macro is ignored
  - Previous releases:
    - Macro was required to enable Fast Path capability
    - Set defaults for Fast Path parameters
- Execution parameters for IMS and DBC procedures
  - FP=Y|N used to enable Fast Path
    - Default is FP=N
  - DBBF=, DBFX=, BSIZ=, and OTHR= used to specify Fast Path parameters
- Benefit
  - Simplifies implementation of Fast Path
- Migration
  - Fast Path users must specify FP=Y and values for DBBF=, DBFX=, BSIZ=, and OTHR= for execution

5

IMS V10 has simplified the implementation of Fast Path. It does not have to specified at system definition time. The FPCTRL system definition macro is ignored. In previous releases it was required to enable Fast Path capabilities. It also was used to specify the default values for some Fast Path parameters. Fast Path capabilities are always generated for DB/DC and DBCTL systems. They are enabled by a parameter at execution time.

You must specify FP=Y at execution time for an online system to enable Fast Path capabilities in IMS V10. The default for the parameter is FP=N. Since defaults for Fast Path parameters cannot be specified at system definition time, they should be specified at execution time. These parameters are DBBF=, DBFX=, BSIZ=, and OTHR=. The default values for these parameters are the same as the default values for the FPCTRL macro in previous releases. These are:

DBBF=10, DBFX=4, BSIZ=2048, and OTHR=2

These values are highly unlikely to be appropriate for most Fast Path users.

When an FPCTRL macro is included in the system definition input, the following message is issued:

```
G1010 THE FPCTRL MACRO IS NO LONGER SUPPORTED.
      THE KEYWORD PARAMETERS WILL BE IGNORED.
      PLEASE USE THE EQUIVALENT EXECUTION
      PARAMETER TO SPECIFY VALUES THAT WERE
      CODED IN THIS MACRO.
```

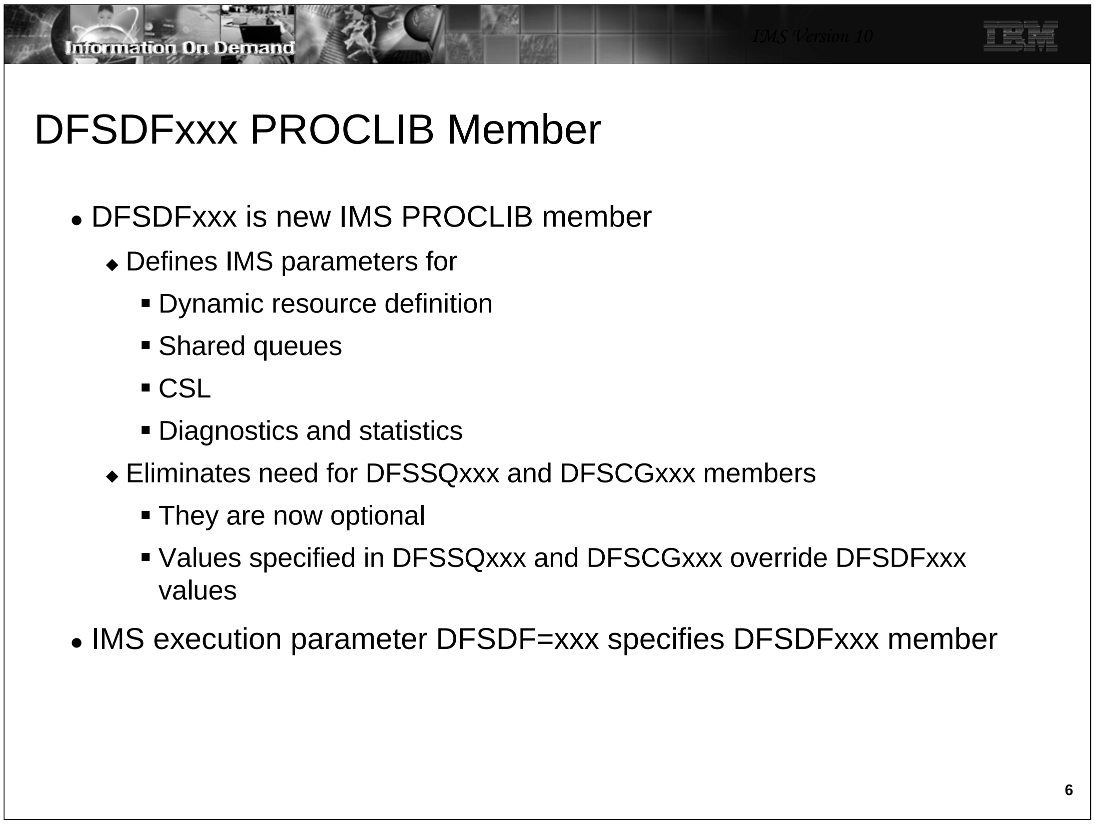

Information On Demand
IMS Version 10

DFSDFxxx PROCLIB Member

- DFSDFxxx is new IMS PROCLIB member
  - Defines IMS parameters for
    - Dynamic resource definition
    - Shared queues
    - CSL
    - Diagnostics and statistics
  - Eliminates need for DFSSQxxx and DFSCGxxx members
    - They are now optional
    - Values specified in DFSSQxxx and DFSCGxxx override DFSDFxxx values
- IMS execution parameter DFSDF=xxx specifies DFSDFxxx member

6

IMS V10 has a new PROCLIB member which is used to consolidate execution definitions which otherwise might be in several PROCLIB members. The DFSDFxxx member may be used to define parameters for several components of IMS.

DFSDFxxx eliminates the need to use a DFSSQxxx member for shared queues parameters or a DFSCGxxx member for CSL parameters. Instead, these parameters may be specified in the DFSDFxxx member. If a parameter is specified in both members, the values specified in DFSSQxxx and DFSCGxxx override those specified in DFSDFxxx.

DFSDFxxx is required to specify some IMS parameters. These are explained later.

The DFSDFxxx member used by IMS is specified in the DFSDF= execution parameter.

.

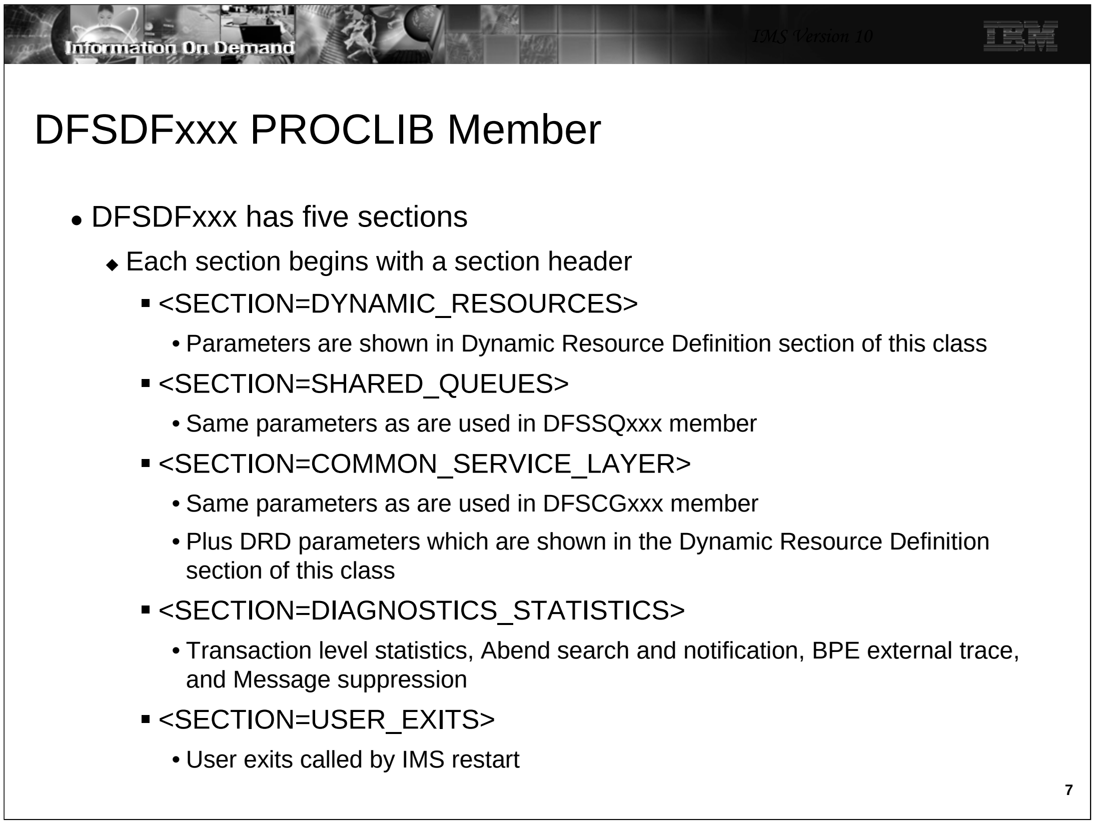

Information On Demand
IMS Version 10

DFSDFxxx PROCLIB Member

- DFSDFxxx has five sections
  - Each section begins with a section header
    - <SECTION=DYNAMIC_RESOURCES>
      - Parameters are shown in Dynamic Resource Definition section of this class
    - <SECTION=SHARED_QUEUES>
      - Same parameters as are used in DFSSQxxx member
    - <SECTION=COMMON_SERVICE_LAYER>
      - Same parameters as are used in DFSCGxxx member
      - Plus DRD parameters which are shown in the Dynamic Resource Definition section of this class
    - <SECTION=DIAGNOSTICS_STATISTICS>
      - Transaction level statistics, Abend search and notification, BPE external trace, and Message suppression
    - <SECTION=USER_EXITS>
      - User exits called by IMS restart

7

The DFSDFxxx member has up to five sections of definitions. Each section begins with a section identification statement. These statements are shown on the slide.

Dynamic resource definition parameters are defined only in DFSDFxxx. They are in the DYNAMIC_RESOURCES and COMMON-SERVICE_LAYER sections.

Shared queues parameters may be defined in the SHARED_QUEUES section.

CSL parameters may be defined in the COMMON_SERVICE_LAYER section.

Some diagnostic and statistics parameters are defined in DFSDFxxx. These are defined in the DIAGNOSTICS_STATISTICS section.

IMS Version 10 includes an IMS restart user exit capability. These exits are specified in the USER_EXITS section. They are called at the beginning and end of IMS restarts. Restart exits are called during all types of IMS restart. An exit routine is passed a function code and a code that indicates the type of restart that is being done. Multiple routines can be defined. The routines are called in the order that they are listed in the EXITDEF parameter. If an exit is defined multiple times, it is called multiple times. The form of the specification is:
EXITDEF=(TYPE=RESTART,EXITS=(*exitname*))

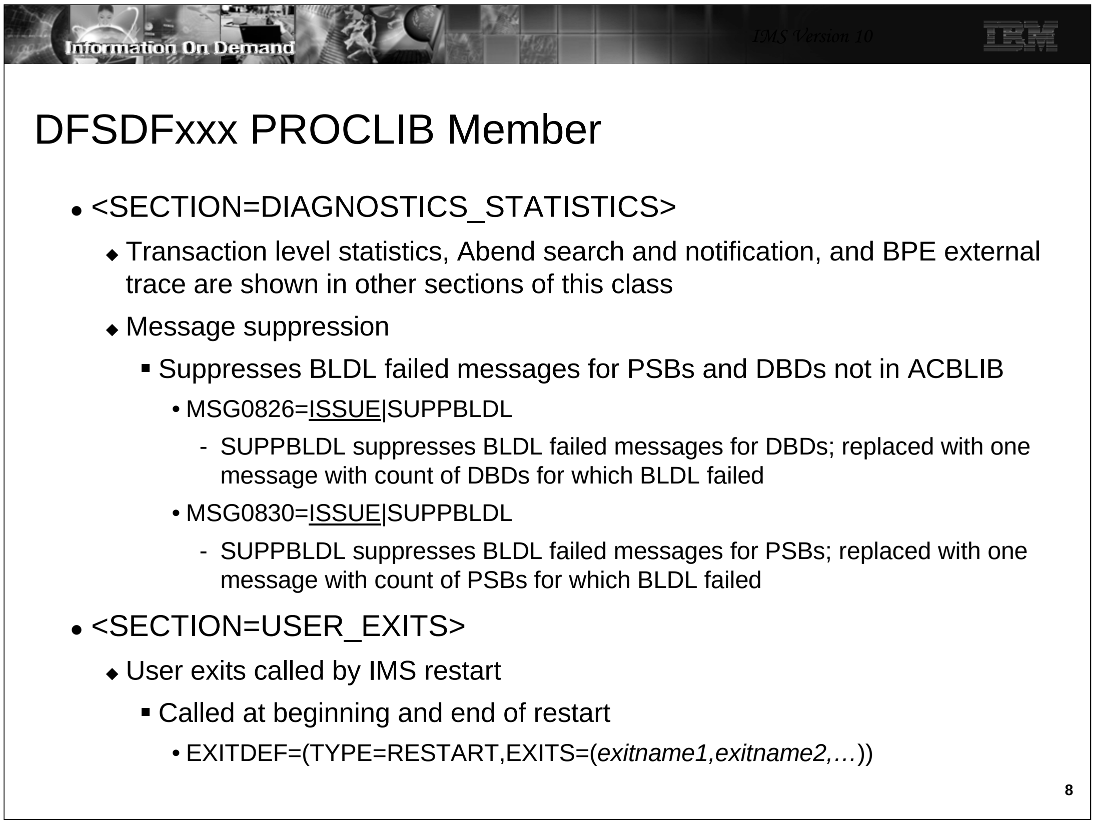

DFSDFxxx PROCLIB Member

- <SECTION=DIAGNOSTICS_STATISTICS>
  - Transaction level statistics, Abend search and notification, and BPE external trace are shown in other sections of this class
  - Message suppression
    - Suppresses BLDL failed messages for PSBs and DBDs not in ACBLIB
      - MSG0826=<u>ISSUE</u>|SUPPBLDL
        - SUPPBLDL suppresses BLDL failed messages for DBDs; replaced with one message with count of DBDs for which BLDL failed
      - MSG0830=<u>ISSUE</u>|SUPPBLDL
        - SUPPBLDL suppresses BLDL failed messages for PSBs; replaced with one message with count of PSBs for which BLDL failed
- <SECTION=USER_EXITS>
  - User exits called by IMS restart
    - Called at beginning and end of restart
      - EXITDEF=(TYPE=RESTART,EXITS=(*exitname1,exitname2,…*))

8

The DIAGNOSTICS_STATISTICS section includes a capability to suppress some DFS826I and DFS830I messages. There messages are issued by IMS restart to indicate that DBDs and PSBs which are defined in the system do not have members in ACBLIB. These messages are suppressed with the following statements:

MSG0826=SUPPBLDL and MSG0830=SUPPBLDL.

If the DBD messages are not suppressed, there is one message for each DBD for which the BLDL failed. Its form is: DFS826I BLDL FAILED FOR FOLLOWING DBD This is followed by a message which includes a count of the number of DBD BLDL failures. Its form is: DFS826I xxx DBD ERRORS SENT TO JOB LOG
If the DBD messages are suppressed, the following message is issued: DFS826I *xxx* DBD ERRORS SUPPRESSED This includes the count of the messages suppressed.

If the PSB messages are not suppressed, there is one message for each PSB for which the BLDL failed. Its form is: DFS830I BLDL FAILED FOR FOLLOWING PSB This is followed by a message which includes a count of the number of DBD BLDL failures. Its form is: DFS830I xxx PSB ERRORS SENT TO JOB LOG
If the PSB messages are suppressed, the following message is issued: DFS830I *xxx* PSB ERRORS SUPPRESSED This includes the count of the messages suppressed.

IMS Version 10 includes an IMS restart user exit capability. These exits are specified in the USER_EXITS section. They are called at the beginning and end of IMS restarts. Restart exits are called during all types of IMS restart. An exit routine is passed a function code and a code that indicates the type of restart that is being done. Multiple routines can be defined. The routines are called in the order that they are listed in the EXITDEF parameter. If an exit is defined multiple times, it is called multiple times. The form of the specification is:
EXITDEF=(TYPE=RESTART,EXITS=(*exitname1.exitname2,…*))

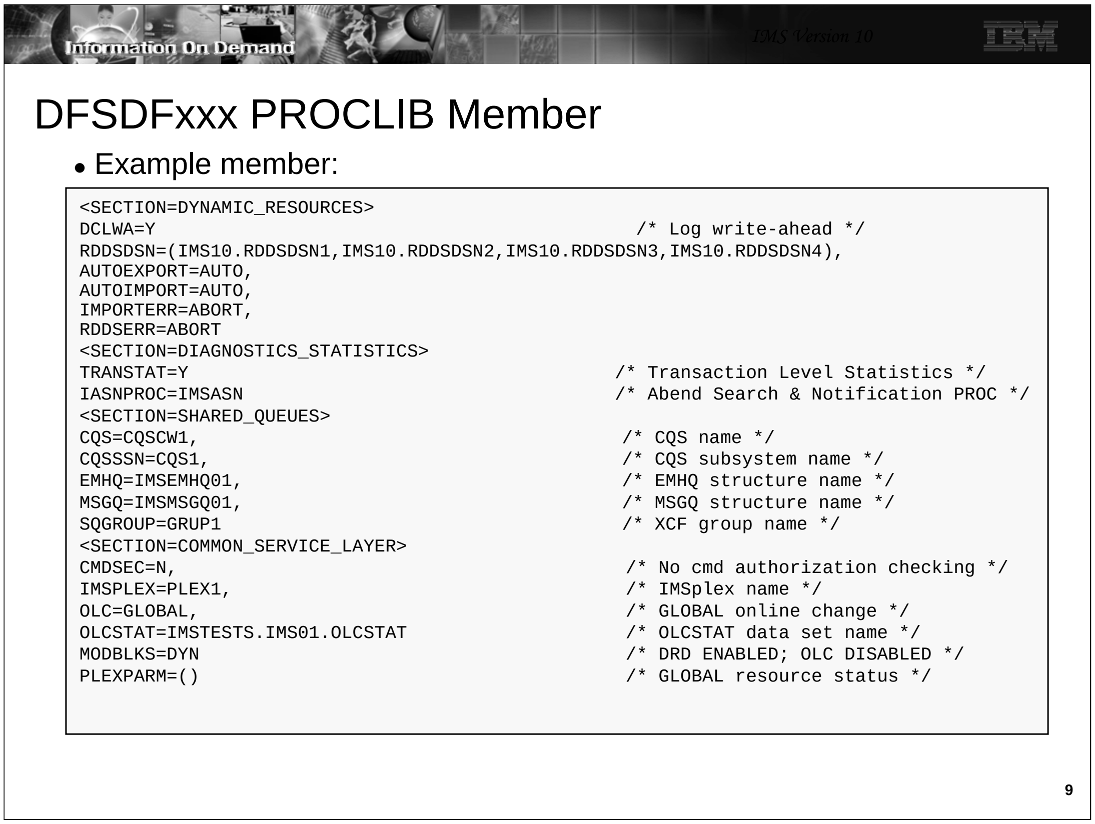
Information On Demand
IMS Version 10

DFSDFxxx PROCLIB Member

- Example member:

```
<SECTION=DYNAMIC_RESOURCES>
DCLWA=Y                                            /* Log write-ahead */
RDDSDSN=(IMS10.RDDSDSN1,IMS10.RDDSDSN2,IMS10.RDDSDSN3,IMS10.RDDSDSN4),
AUTOEXPORT=AUTO,
AUTOIMPORT=AUTO,
IMPORTERR=ABORT,
RDDSERR=ABORT
<SECTION=DIAGNOSTICS_STATISTICS>
TRANSTAT=Y                                       /* Transaction Level Statistics */
IASNPROC=IMSASN                                  /* Abend Search & Notification PROC */
<SECTION=SHARED_QUEUES>
CQS=CQSCW1,                                       /* CQS name */
CQSSSN=CQS1,                                      /* CQS subsystem name */
EMHQ=IMSEMHQ01,                                   /* EMHQ structure name */
MSGQ=IMSMSGQ01,                                   /* MSGQ structure name */
SQGROUP=GRUP1                                     /* XCF group name */
<SECTION=COMMON_SERVICE_LAYER>
CMDSEC=N,                                         /* No cmd authorization checking */
IMSPLEX=PLEX1,                                    /* IMSplex name */
OLC=GLOBAL,                                       /* GLOBAL online change */
OLCSTAT=IMSTESTS.IMS01.OLCSTAT                    /* OLCSTAT data set name */
MODBLKS=DYN                                       /* DRD ENABLED; OLC DISABLED */
PLEXPARM=()                                       /* GLOBAL resource status */
```

9

This is an example of a DFSDFxxx member. It includes parameters for DRD, traces, transaction level statistics, abend search and notification, shared queues, and CSL.

The meanings of the parameters other that those for shared queues and CSL are explained elsewhere in this class. Shared queues and CSL parameters have the same meaning as they had in the DFSSQxxx and DFSCGxxx members for previous releases.

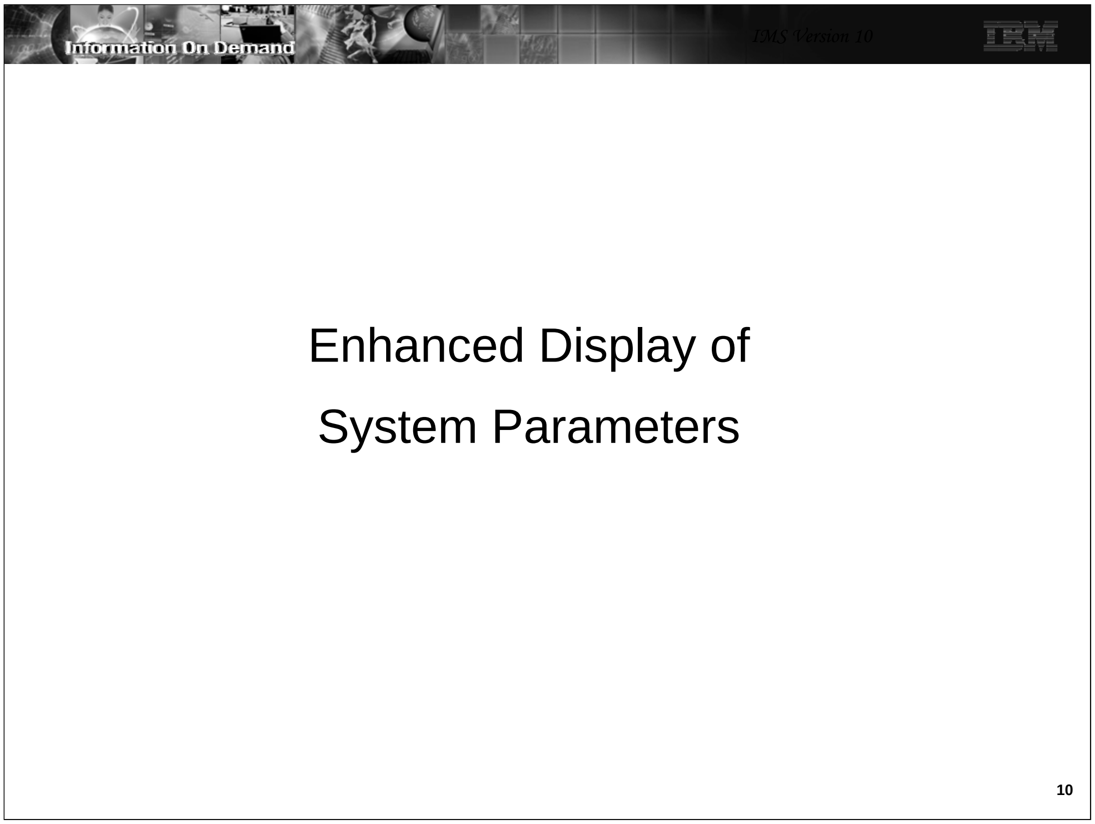

Enhanced Display of System Parameters

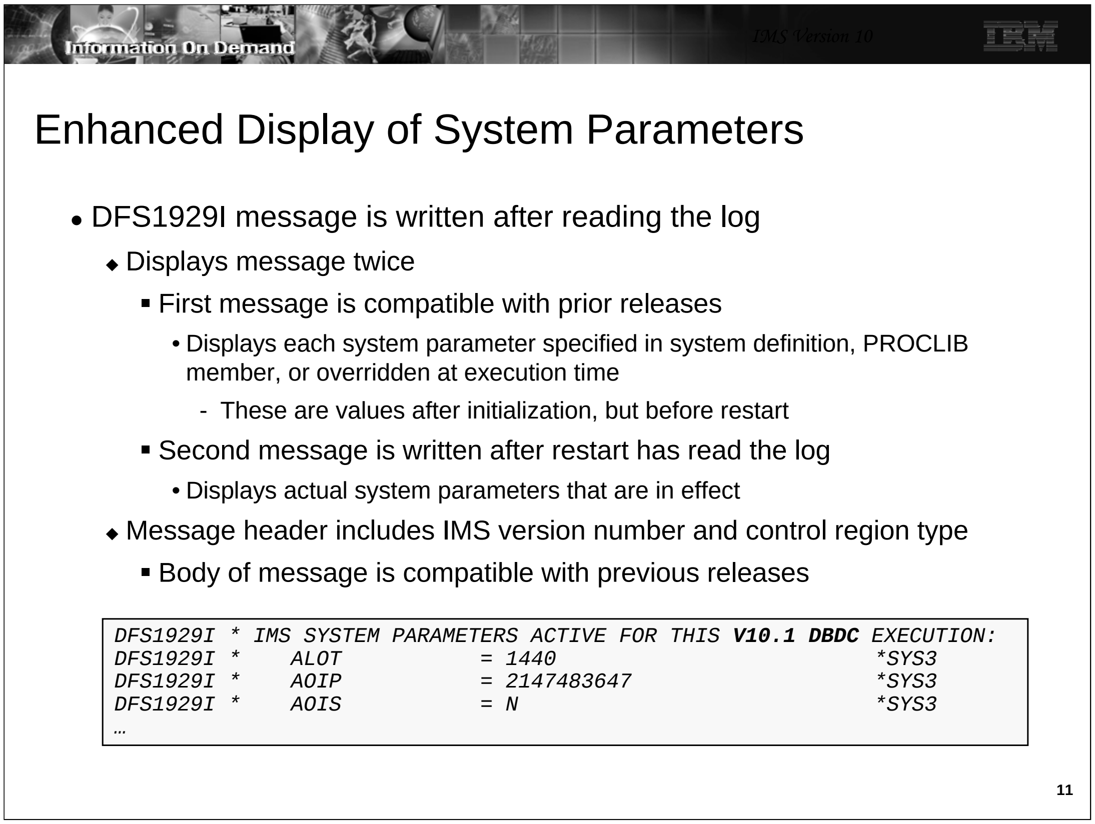
Information On Demand
IMS Version 10
IBM

Enhanced Display of System Parameters

- DFS1929I message is written after reading the log
  - Displays message twice
    - First message is compatible with prior releases
      - Displays each system parameter specified in system definition, PROCLIB member, or overridden at execution time
        - These are values after initialization, but before restart
    - Second message is written after restart has read the log
      - Displays actual system parameters that are in effect
  - Message header includes IMS version number and control region type
    - Body of message is compatible with previous releases

```
DFS1929I * IMS SYSTEM PARAMETERS ACTIVE FOR THIS V10.1 DBDC EXECUTION:
DFS1929I *    ALOT           = 1440                        *SYS3
DFS1929I *    AOIP           = 2147483647                  *SYS3
DFS1929I *    AOIS           = N                           *SYS3
…
```

11

IMS issues message DFS1929I to list the IMS system parameters. Since the message is issued early in the initialization process, the values displayed reflect those provided during system generation or overridden through startup parameters. Every parameter is shown. If a value has not been provided by the user then the default, if one exists, is displayed. The values shown, however, might not reflect the actual values that are used during online execution. The discrepancy between the actual and requested values may be minor, but in some situations it can be large. IMS V10 continues to issue the DFS1929I message early in initialization, for compatibility, but issues it once again after restart is complete and the log has been read so that the values displayed reflect the actual values in effect for the IMS execution.

The DFS1929I message has been enhanced to include the IMS Control Region type along with the IMS Version number. The remainder of the message is compatible with the format of previous releases. The first form of the message that is printed during IMS initialization contains all the IMS parameters with user-specified or default values.

The second DFS1929I message which is produced after the log has been read has been enhanced to list:

►Only the parameters that are applicable to the control region type.

►Only the parameters associated with active components. For example, Fast Path parameters will be shown only if activated. Prior releases always displayed the parameters regardless of whether or not Fast Path was enabled. The new IMS V10 parameter FP=Y | N will always be displayed.

►Actual values that will be used during execution. For example, the CPLOG value displayed in the first DFS1929I message during initialization might not be valid after a warm or emergency restart because its value was changed by a command. As a result, the value in the first DFS1929I message does not help in diagnosing problems. The problem encompasses other parameters such as APPC, APPCSE, OTMA, OTMASE, etc., which can also be changed by command and are recoverable during an IMS restart. The actual values, therefore, are provided in the second DFS1929I message after restart is complete.

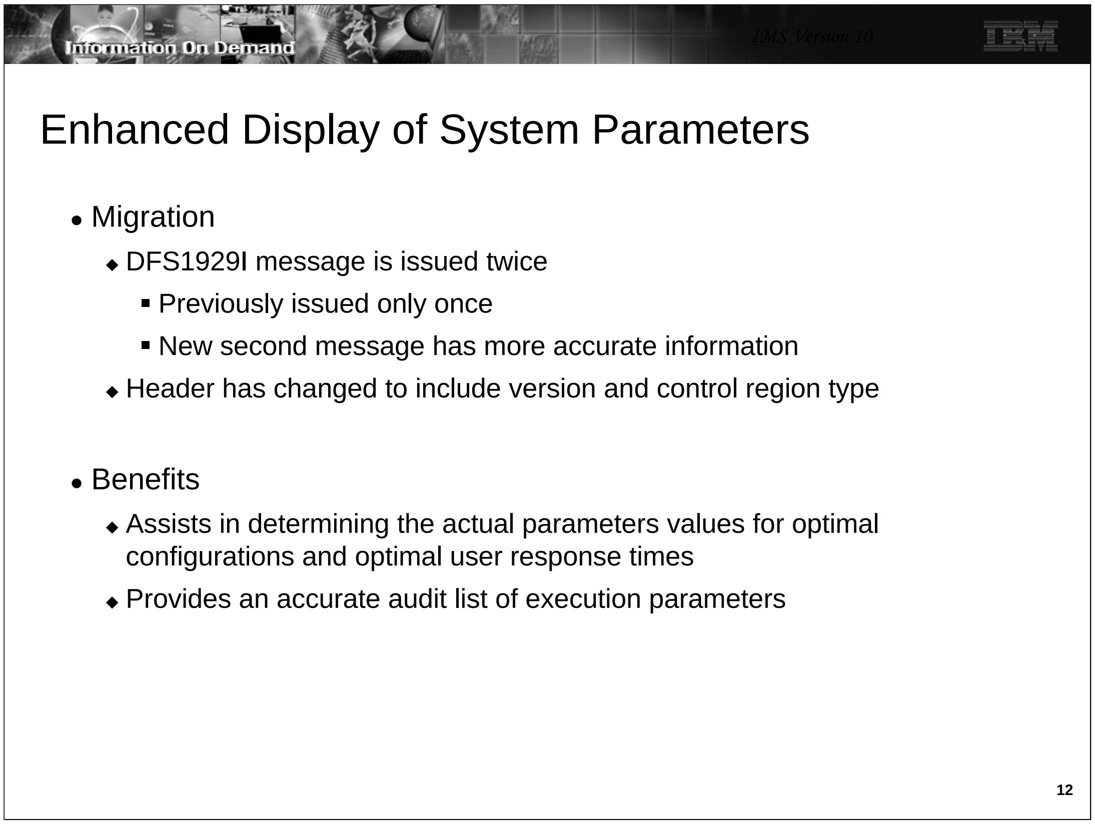
Information On Demand
IMS Version 10

Enhanced Display of System Parameters

- Migration
  - DFS1929I message is issued twice
    - Previously issued only once
    - New second message has more accurate information
  - Header has changed to include version and control region type

- Benefits
  - Assists in determining the actual parameters values for optimal configurations and optimal user response times
  - Provides an accurate audit list of execution parameters

When migrating to IMS V10, user-written routines or procedures that look for the DFS1929I message should recognize the header format change to include the IMS version and Control Region type. Additionally, these routines should be modified to recognize that there are two versions of the message.

Virtual Storage Constraint Relief

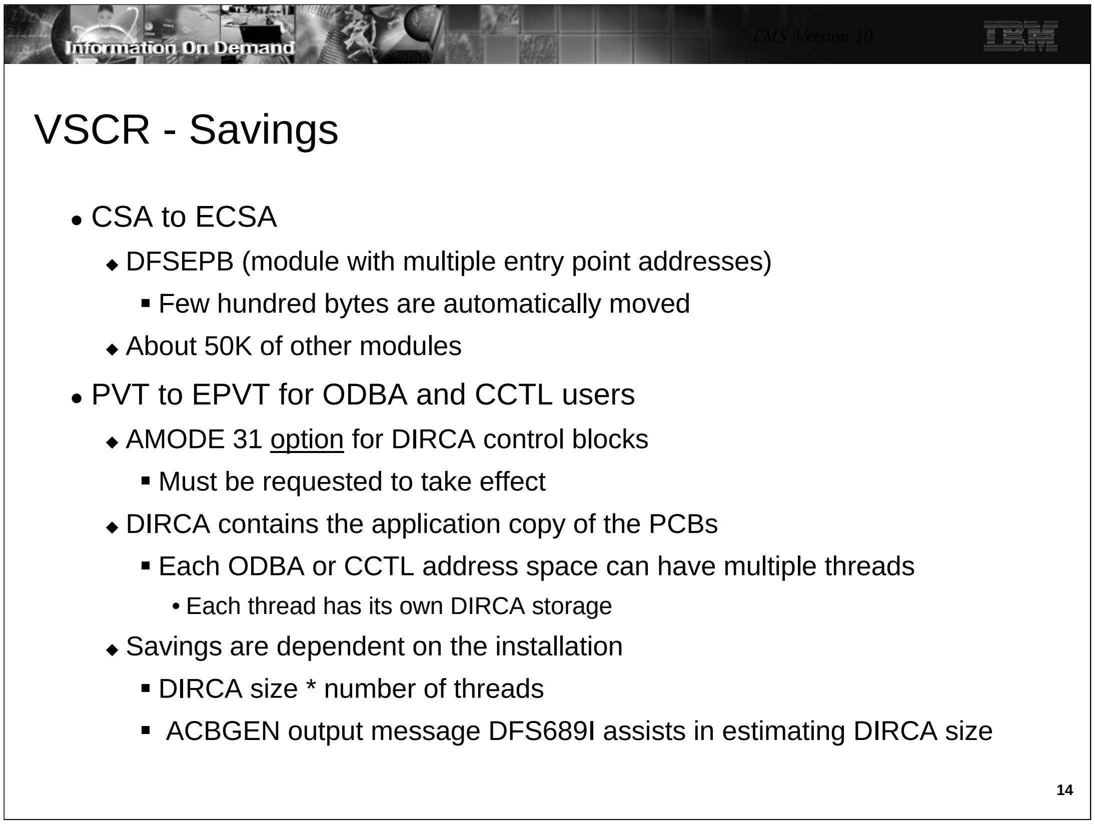
Information On Demand
IMS Version 10
IBM

VSCR - Savings

- CSA to ECSA
  - DFSEPB (module with multiple entry point addresses)
    - Few hundred bytes are automatically moved
  - About 50K of other modules
- PVT to EPVT for ODBA and CCTL users
  - AMODE 31 option for DIRCA control blocks
    - Must be requested to take effect
  - DIRCA contains the application copy of the PCBs
    - Each ODBA or CCTL address space can have multiple threads
      - Each thread has its own DIRCA storage
  - Savings are dependent on the installation
    - DIRCA size * number of threads
    - ACBGEN output message DFS689I assists in estimating DIRCA size

The DFSEPB (Entry Point Block) was introduced in IMS V9 to keep track of entry points of CSECTs of large composite modules that needed to be divided into smaller modules. The movement of this storage from CSA to ECSA saves a few hundred bytes in IMS V10. In addition, a few other modules have been moved from CSA to ECSA. These modules total about 50K bytes.

A more substantial impact is that of potential DIRCA (Dependent Region Interregion Communication Area) storage movement above the line for ODBA and CCTL address spaces. The savings are installation dependent and must be requested for the environment to take advantage of the capability. Each thread from the ODBA/CCTL address space has its own DIRCA and each DIRCA contains the application's copy of the PCBs. This enhancement provides a benefit for those environments, such as CICS, that need to use larger PSBs and more concurrent threads without impacting storage usage below the 16MB line.

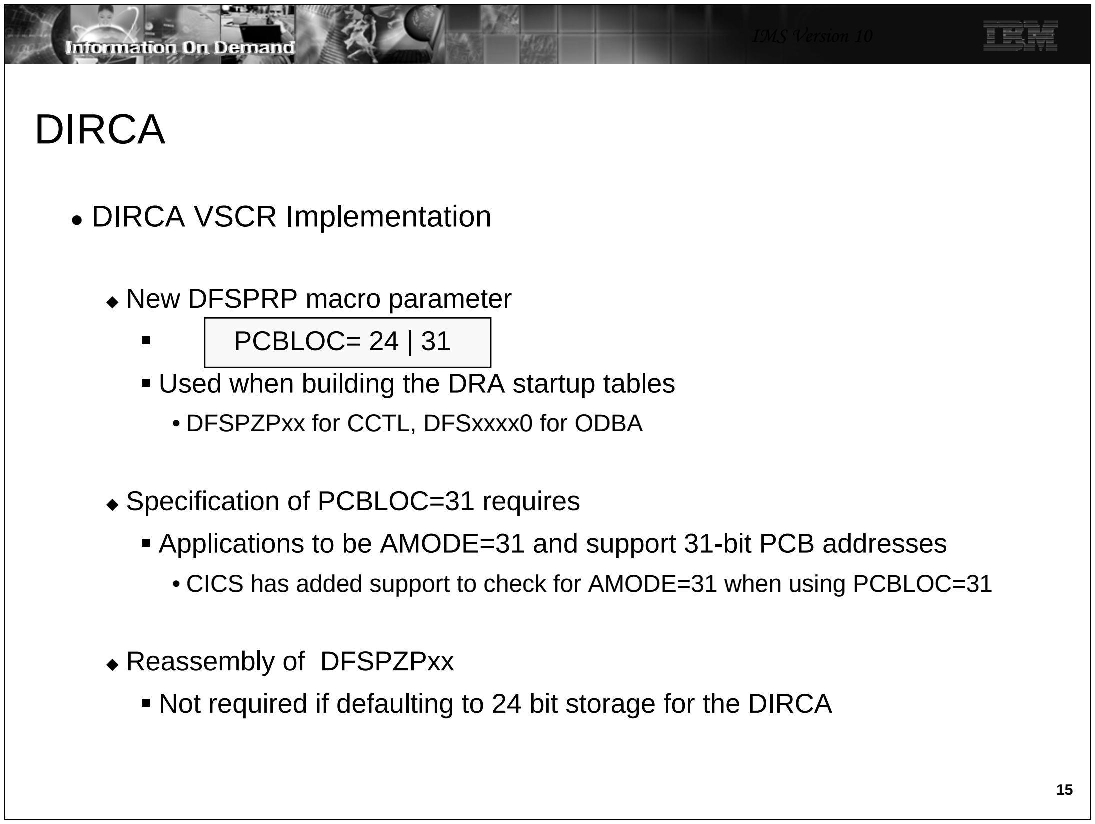
Information On Demand
IMS Version 10
IBM

DIRCA

- DIRCA VSCR Implementation
  - New DFSPRP macro parameter
    - PCBLOC= 24 | 31
    - Used when building the DRA startup tables
      - DFSPZPxx for CCTL, DFSxxxx0 for ODBA
  - Specification of PCBLOC=31 requires
    - Applications to be AMODE=31 and support 31-bit PCB addresses
      - CICS has added support to check for AMODE=31 when using PCBLOC=31
  - Reassembly of DFSPZPxx
    - Not required if defaulting to 24 bit storage for the DIRCA

15

The DFSPRP macro provides the definition for the DRA parameters when building the DRA startup table. A new parameter PCBLOC provides the option of defining whether the DIRCA is to be built using 24-bit or 31-bit storage. Note that a specification of 31 also requires that the applications actually support AMODE 31 and 31-bit PCB addresses.

CICS has added support so that it checks for the use of AMODE=31 by applications when PCBLOC=31 is specified. The application will abend with a new CICS abend code ADCF during schedule of the PSB if it is AMODE=24 and PCBLOC=31 is in effect. This support is added by APARS PK54099 for CICS TS 2.2 and CICS TS 2.3, and PK54100 for CICS TS 3.1 and CICS TS 3.2.

When a CCTL connector starts up, the DRA INIT call can be used to override the PCBLOC specification in the DRA startup table. A new parameter, PPLLPSO, allows specification of 31 or 24. Other values are ignored and the startup defaults to the value specified in the startup table.

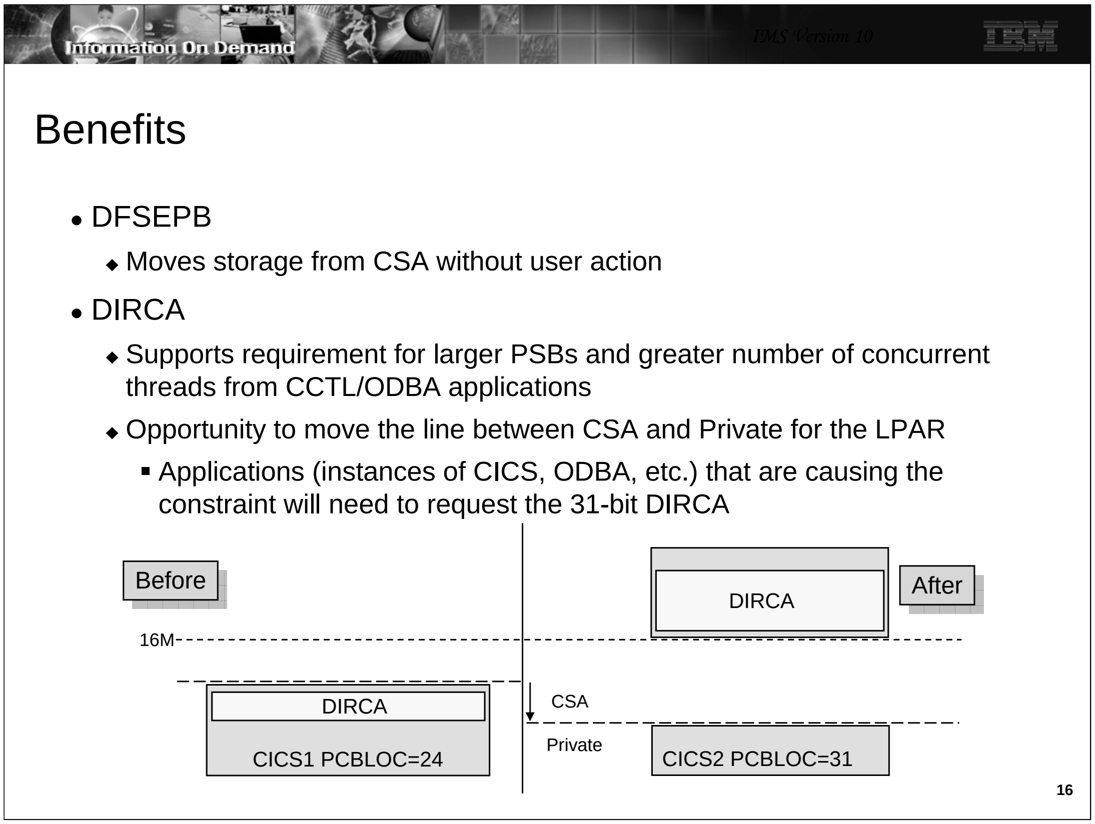

Information On Demand

IMS Version 10

Benefits

- DFSEPB
  - Moves storage from CSA without user action
- DIRCA
  - Supports requirement for larger PSBs and greater number of concurrent threads from CCTL/ODBA applications
  - Opportunity to move the line between CSA and Private for the LPAR
    - Applications (instances of CICS, ODBA, etc.) that are causing the constraint will need to request the 31-bit DIRCA

16

The DFSEPB module is automatically moved above the line from CSA to ECSA.

The DIRCA is moved above the 16MB line from private to extended private. This provides two benefits. First, PSBs may be larger. Second, the CICS or ODBA address space private area does not have to be as large, so the line between private and CSA may be lower. This provides more room in CSA.

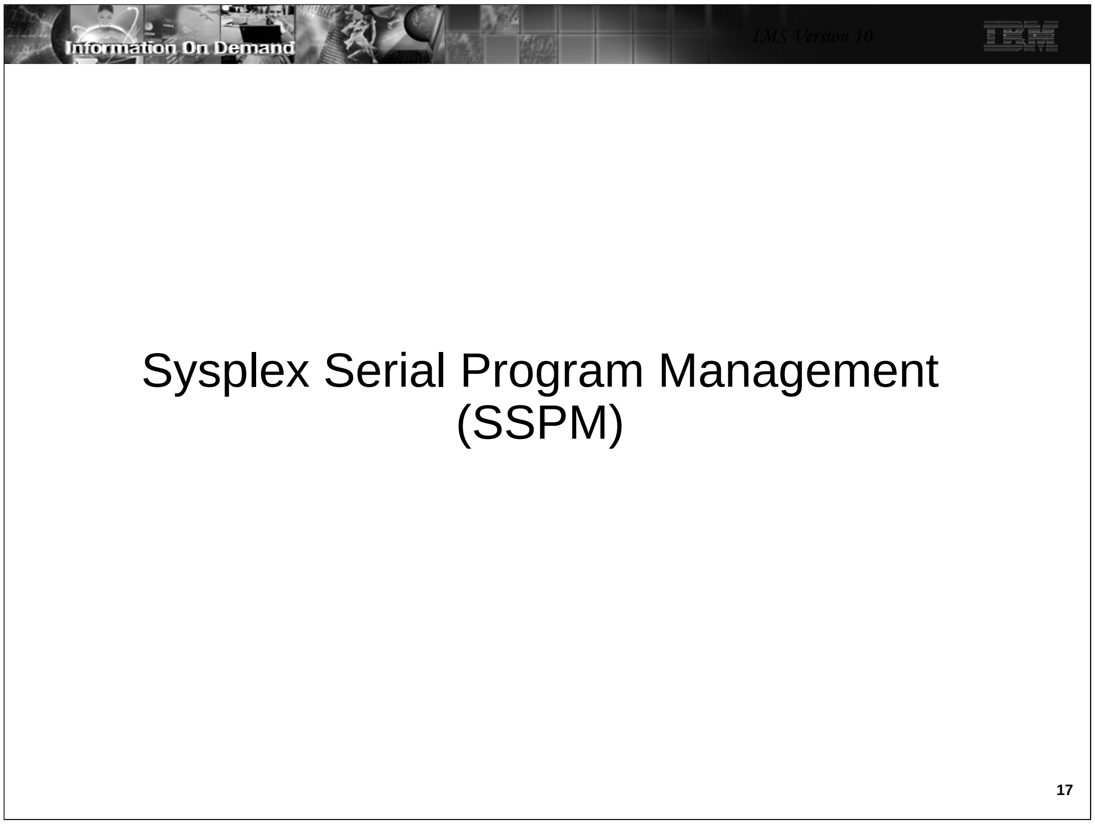

Information On Demand
IMS Version 10

Sysplex Serial Program Management (SSPM)

17

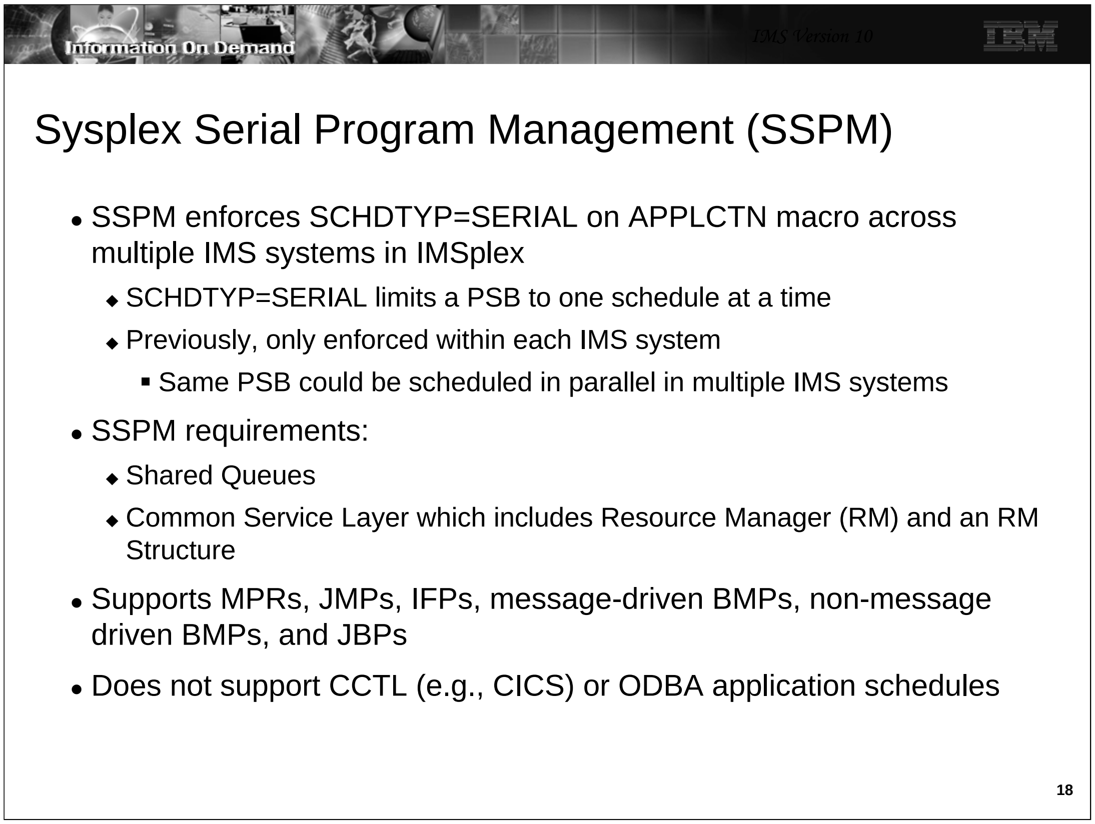
Information On Demand
IMS Version 10
IBM

Sysplex Serial Program Management (SSPM)

- SSPM enforces SCHDTYP=SERIAL on APPLCTN macro across multiple IMS systems in IMSplex
  - SCHDTYP=SERIAL limits a PSB to one schedule at a time
  - Previously, only enforced within each IMS system
    - Same PSB could be scheduled in parallel in multiple IMS systems
- SSPM requirements:
  - Shared Queues
  - Common Service Layer which includes Resource Manager (RM) and an RM Structure
- Supports MPRs, JMPs, IFPs, message-driven BMPs, non-message driven BMPs, and JBPs
- Does not support CCTL (e.g., CICS) or ODBA application schedules

IMS Version 10 has enhanced the support for serialization of programs that are defined as SCHDTYP=SERIAL to be across all the members of an IMSplex, rather than just within a single IMS when using Shared Queues. This serialization is only done in systems using shared queues and a Resource Manager with a Resource Manager structure.

IMS V10 will ensure that SCHDTYP=SERIAL programs will be scheduled in only one IMS dependent region (MPR, JMP, IFP, message-driven BMP, non-message driven BMP, JBP) across an entire IMSplex with Shared Queues at a time.

CICS and ODBA are not supported for SSPM. SCHDTYP=SERIAL may be specified for an APPLCTN used by CICS or ODBA. As in previous releases, this will serialize the scheduling of the PSB by CICS, ODBA, and IMS TM within an IMS system. On the other hand, the serialization across multiple IMS systems will only be enforced for IMS TM regions including MPPs, IFPs, JMPs, BMPs, and JBPs.

SCHDTYP=SERIAL is the default if not specified on the APPLCTN macro.

IMS will use the IMS Resource Manager (RM) component of the IMS Common Services Layer (CSL) to provide this serialization.

SCHDTYP=SERIAL is specified on the APPLCTN macro or on the CREATE PGM SCHDTYPE(SERIAL) command when using DRD. This should not be confused with the SERIAL=YES parameter on the TRANSACT macro. SCHDTYP=SERIAL for APPLCTN limits the scheduling of a PSB. SERIAL=YES for TRANSACT limits the scheduling of a transaction. The meaning and enforcement of SERIAL=YES on the TRANSACT macro is not changed in IMS Version 10.

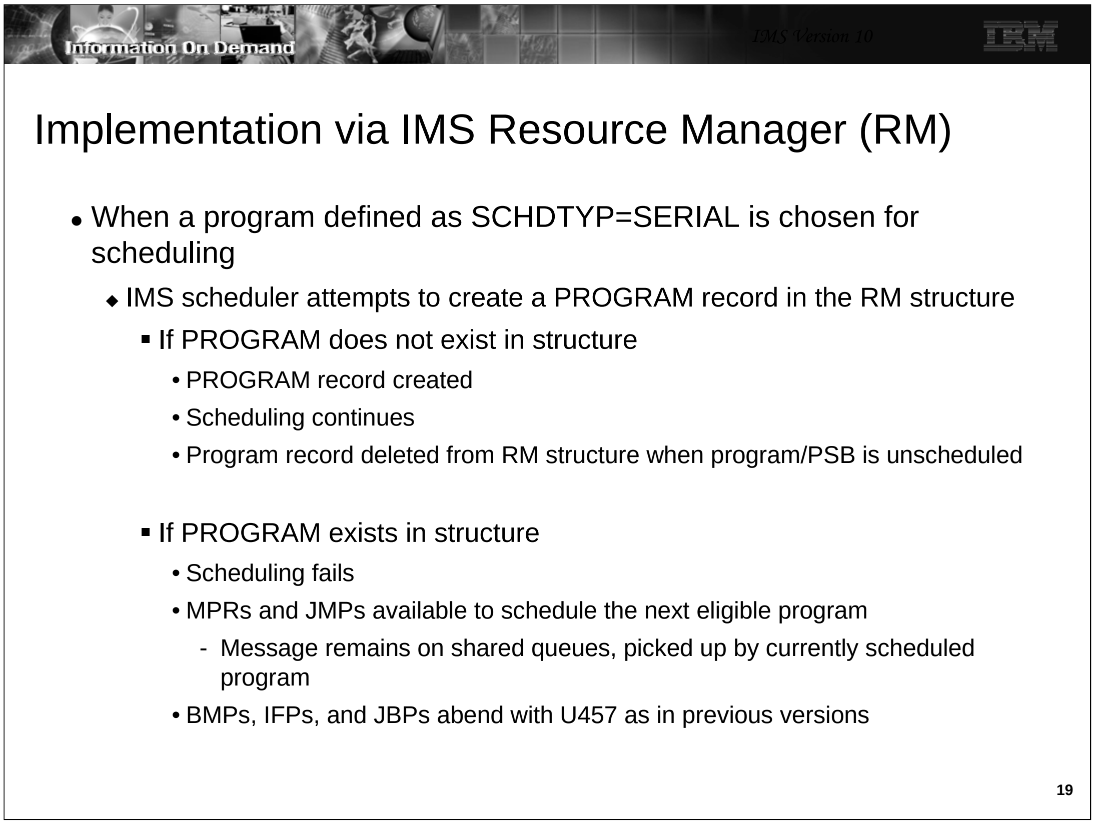

Implementation via IMS Resource Manager (RM)

- When a program defined as SCHDTYP=SERIAL is chosen for scheduling
  - IMS scheduler attempts to create a PROGRAM record in the RM structure
    - If PROGRAM does not exist in structure
      - PROGRAM record created
      - Scheduling continues
      - Program record deleted from RM structure when program/PSB is unscheduled
    - If PROGRAM exists in structure
      - Scheduling fails
      - MPRs and JMPs available to schedule the next eligible program
        - Message remains on shared queues, picked up by currently scheduled program
      - BMPs, IFPs, and JBPs abend with U457 as in previous versions

IMS Version 10 will determine if a SCHDTYP=SERIAL program is already scheduled by interacting with the IMS Resource Manager and its resource structure. IMS uses RM to attempt to create the program resource in the resource structure.

If the program resource does not exist in the resource structure, that means no other copy of this program is scheduled at this time. Therefore, this scheduling can continue. A program record for this program will be created in the resource structure. For MPRs, JMPs, and IFPs, the program resource remains active in the resource structure until the PSB is unscheduled. For BMPs and JBPs, the program resource will be deleted when the batch message processing program terminates.

If the serial program is already scheduled on another IMS, the program resource already exists in RM and the attempt to create the program resource in RM fails. Therefore, this attempt to schedule what would be the second copy of the serial program will fail. For MPRs and JMPs, the dependent region will be available to schedule the next eligible program. The message that got the scheduling failure will remain on the shared queues (was never taken off) and will hopefully be picked up by the currently scheduled program. BMPs, JMPs, and IFPs will get a U457 abend when scheduling finds the program is already scheduled as they would have before in previous versions within a single IMS versus in IMS V10 this abend would be on another IMS in the IMSplex.

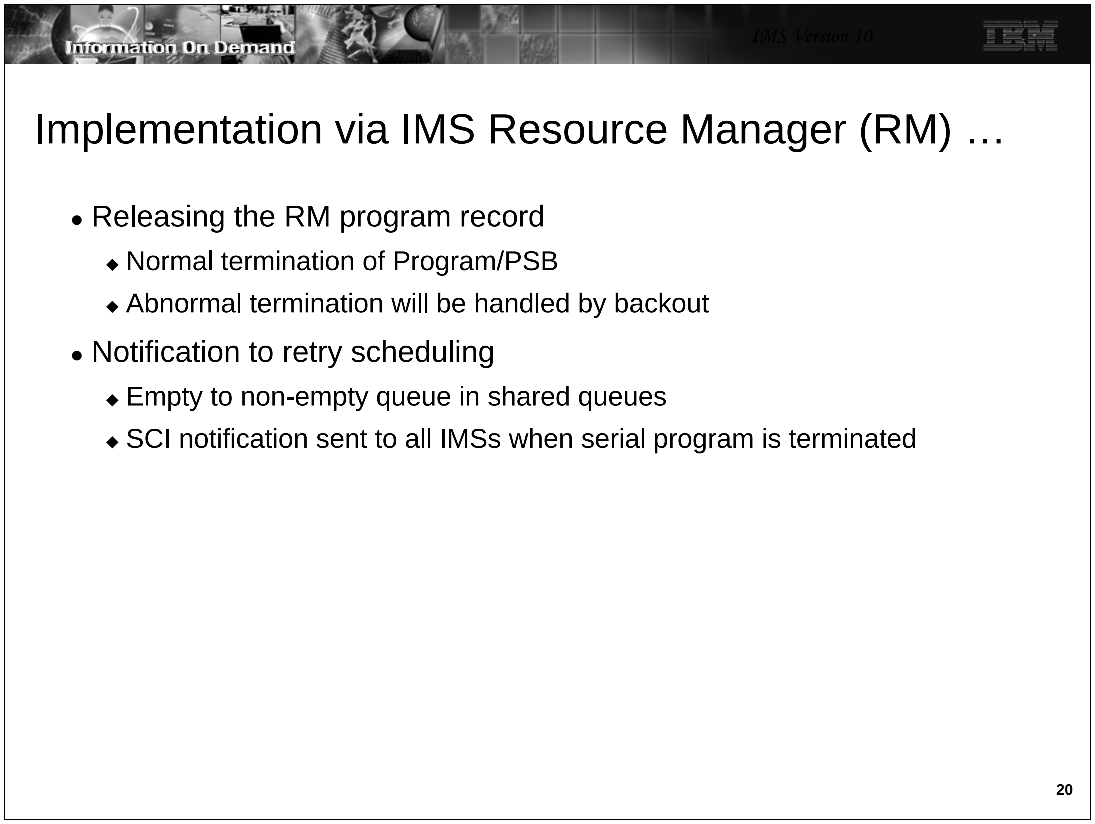

Implementation via IMS Resource Manager (RM) …

- Releasing the RM program record
  - Normal termination of Program/PSB
  - Abnormal termination will be handled by backout
- Notification to retry scheduling
  - Empty to non-empty queue in shared queues
  - SCI notification sent to all IMSs when serial program is terminated

The program record in the RM structure is deleted when the program terminates. It the program ABENDs, backout will delete the program record.

After a scheduling attempt fails due to this serialization, the IMS system will not attempt to schedule the PSB again until it receives a notification. This can occur in two ways. It may occur due to the queue going from an empty to a non-empty state. It may occur due to an SCI notification which is sent when the serial program is terminated.

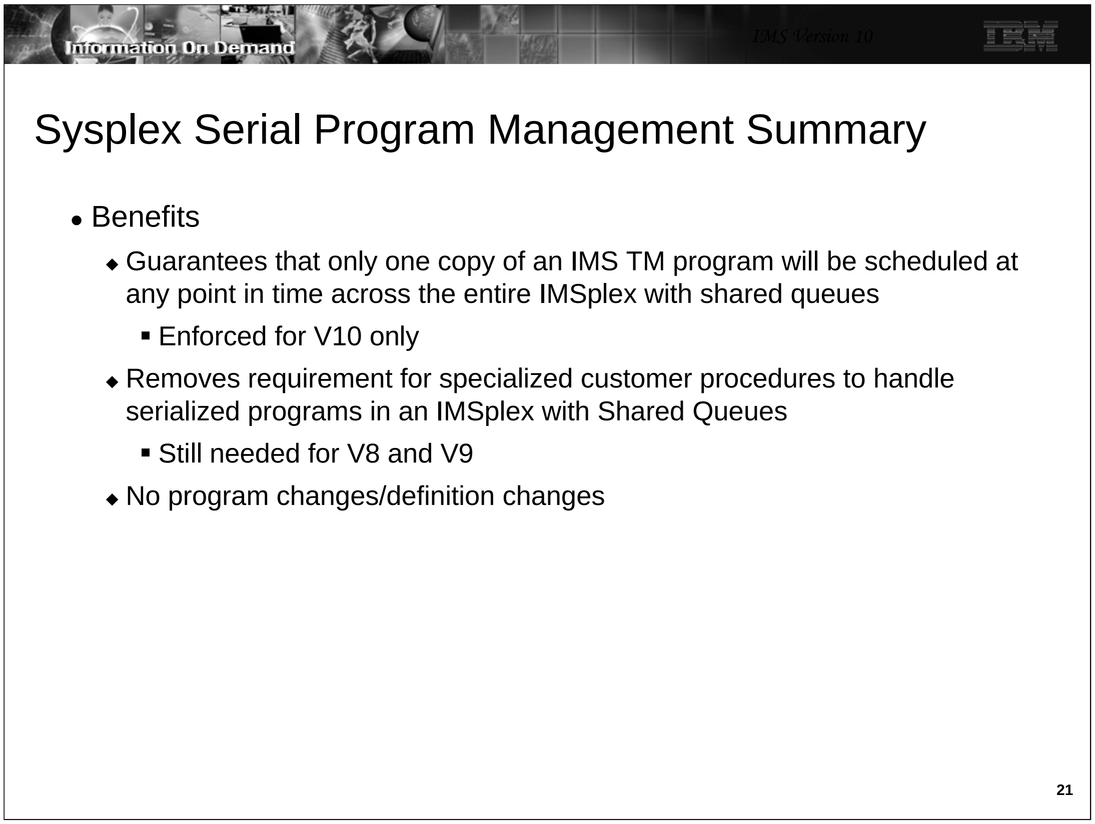

Sysplex Serial Program Management Summary

- Benefits
  - Guarantees that only one copy of an IMS TM program will be scheduled at any point in time across the entire IMSplex with shared queues
    - Enforced for V10 only
  - Removes requirement for specialized customer procedures to handle serialized programs in an IMSplex with Shared Queues
    - Still needed for V8 and V9
  - No program changes/definition changes

Enhanced Log Record Statistics

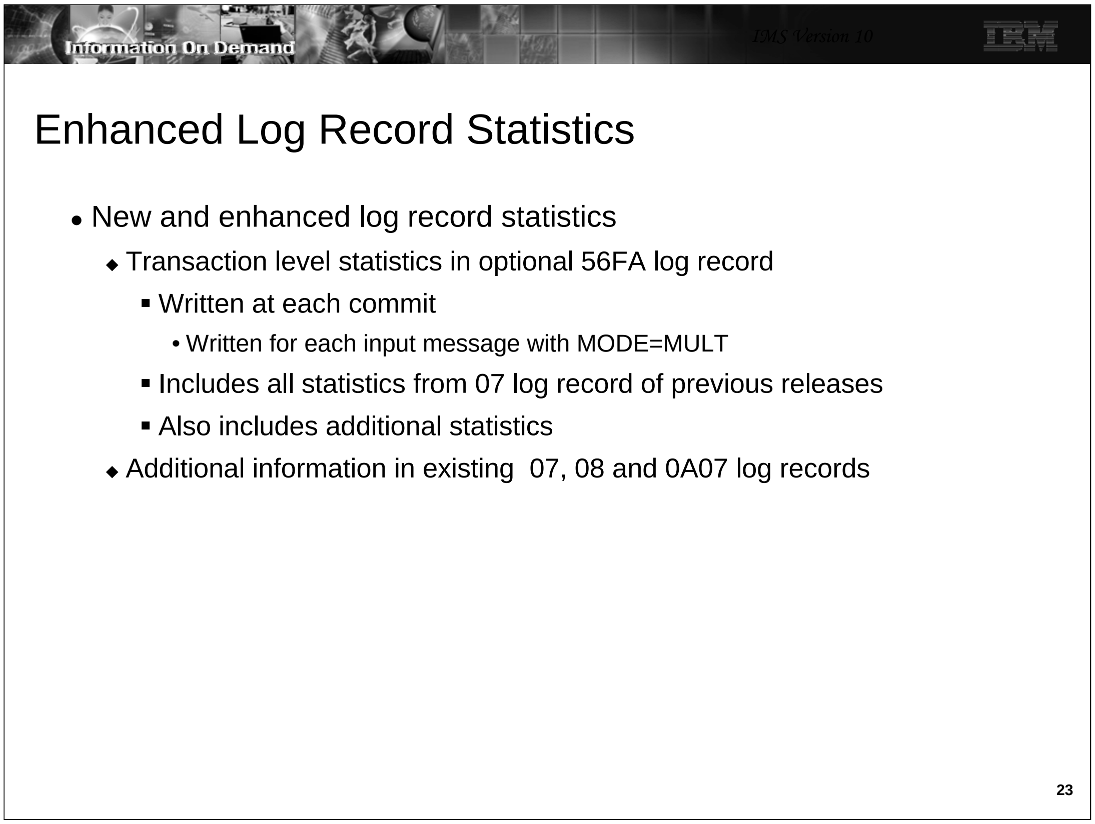
Information On Demand
IMS Version 10

Enhanced Log Record Statistics

- New and enhanced log record statistics
  - Transaction level statistics in optional 56FA log record
    - Written at each commit
      - Written for each input message with MODE=MULT
    - Includes all statistics from 07 log record of previous releases
    - Also includes additional statistics
  - Additional information in existing 07, 08 and 0A07 log records

Enhanced statistics are written in several IMS log records. This includes a new 56FA log record which records statistics for individual transactions, not just for executions of programs or schedulings of PSBs. When chosen it is written for each commit for transactions defined as MODE=SNGL and for non-message driven BMPs and JBPs. For transactions defined as MODE=MULT the log record is written for each message. This record includes all of the statistics which were written to the 07 log records in previous releases. This includes CPU time and call counts. In addition, there are also new statistics which are shown on the nest page.

Additional information has been added to the application start accounting log record (08), the application termination accounting log record (07), and the CPI-C driven transaction termination log record (0A07).

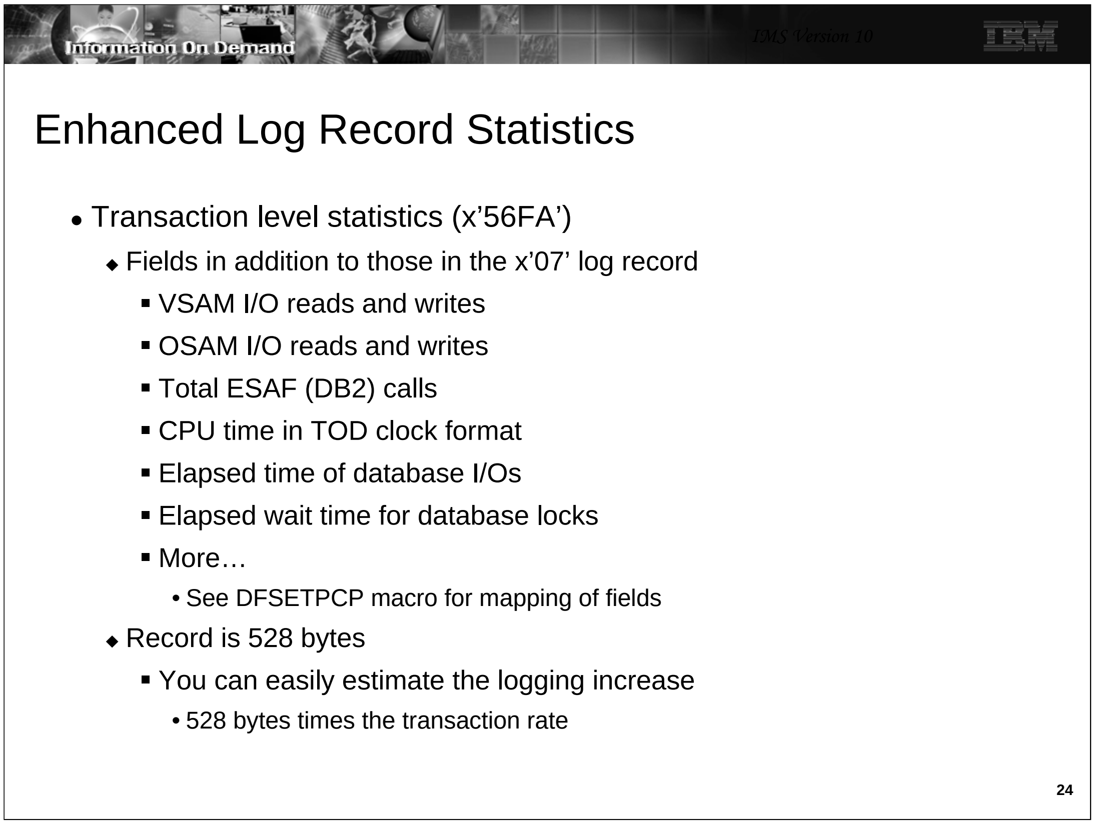

Information On Demand
IMS Version 10
IBM

Enhanced Log Record Statistics

- Transaction level statistics (x’56FA’)
  - Fields in addition to those in the x’07’ log record
    - VSAM I/O reads and writes
    - OSAM I/O reads and writes
    - Total ESAF (DB2) calls
    - CPU time in TOD clock format
    - Elapsed time of database I/Os
    - Elapsed wait time for database locks
    - More…
      - See DFSETPCP macro for mapping of fields
  - Record is 528 bytes
    - You can easily estimate the logging increase
      - 528 bytes times the transaction rate

24

The additional fields in the x’56FA’ log record include read counts and write counts for both VSAM and OSAM database data sets. There are counts of External Subsystem Attach Facility calls. There are typically calls to DB2, but they also may be for MQ or another subsystem. The CPU time for the transaction is written in time of day (TOD) format. The statistics include the total elapsed time for database I/Os and the total elapsed time for waits for locks. These two elapsed times were previously available in the x’07’ log records for DBCTL threads. It was not available for IMS TM except in IMS Monitor trace records. The x’56FA’ log record is mapped by the DFSETPCP macro.

The x’56FA’ log record is approximately 528 bytes. It is written for every transaction and for every commit for non-message driven programs for those transactions and programs with the option. If you choose to use this option, you can easily estimate its effect on logging volumes. The increase in logging may be noticeable, but most IMS systems will easily handle this increase.

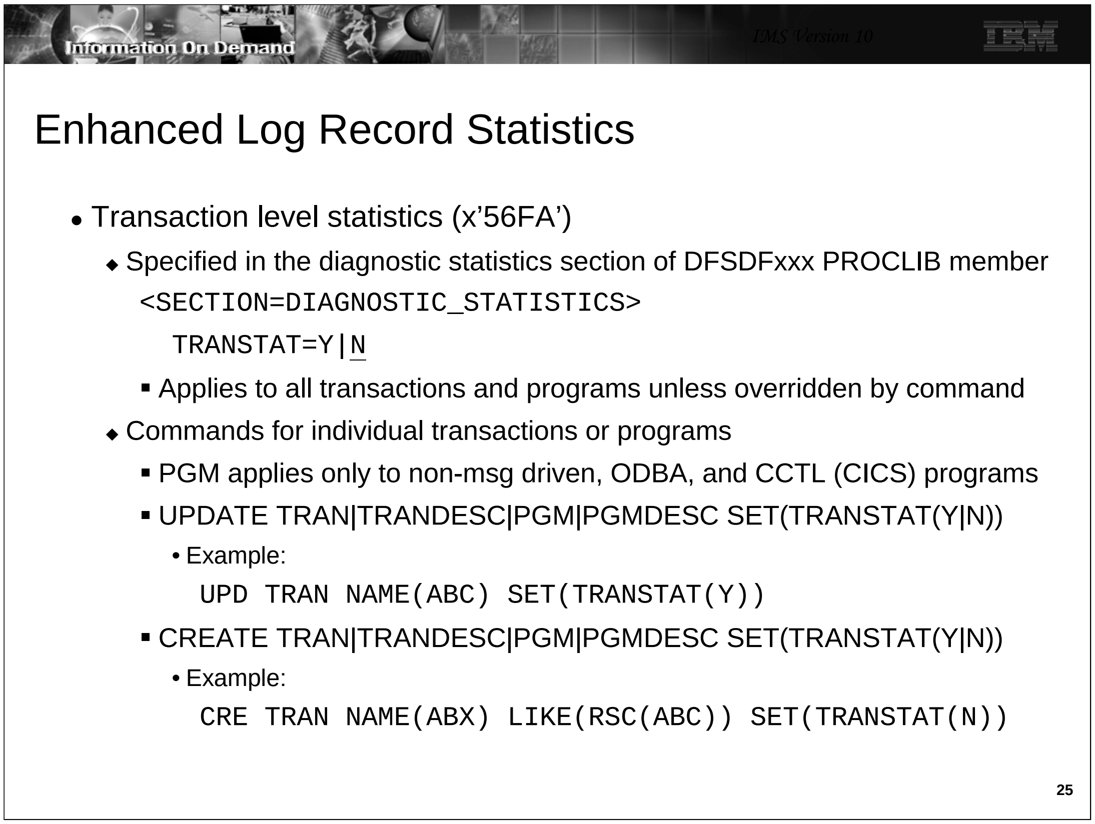

Information On Demand
IMS Version 10
IBM

Enhanced Log Record Statistics

- Transaction level statistics (x'56FA')
  - Specified in the diagnostic statistics section of DFSDFxxx PROCLIB member

```
<SECTION=DIAGNOSTIC_STATISTICS>
  TRANSTAT=Y|N
```

    - Applies to all transactions and programs unless overridden by command
  - Commands for individual transactions or programs
    - PGM applies only to non-msg driven, ODBA, and CCTL (CICS) programs
    - UPDATE TRAN|TRANDESC|PGM|PGMDESC SET(TRANSTAT(Y|N))
      - Example:

```
UPD TRAN NAME(ABC) SET(TRANSTAT(Y))
```

    - CREATE TRAN|TRANDESC|PGM|PGMDESC SET(TRANSTAT(Y|N))
      - Example:

```
CRE TRAN NAME(ABX) LIKE(RSC(ABC)) SET(TRANSTAT(N))
```

25

Transaction level statistics log records are optional. The TRANSTAT specification in the diagnostic statistics section of the DFSDFxxx PROCLIB member determines the default setting for all transactions and non-message driven programs. TRANSTAT=N is the default. TRANSTAT=Y causes the log records to be written. The setting in this member may be overridden for individual transactions and non-message driven programs by using commands.

There are two commands that may be used to set transaction level statistics for transactions and non-message driven programs.

The UPDATE command can be used to change the current setting of transaction level statistics for any of the entities.

The CREATE command includes the SET(TRANSTAT(Y|N) parameter for setting transaction level statistics on or off when creating a new transaction, transaction descriptor, program, or program descriptor. It defaults to the setting in the descriptor or the transaction or program specified in the LIKE parameter.

There is also a QUERY command that may be used to show the current setting for a transaction (TRAN), transaction descriptor (TRANDESC), program (PGM), or program descriptor (PGMDESC).

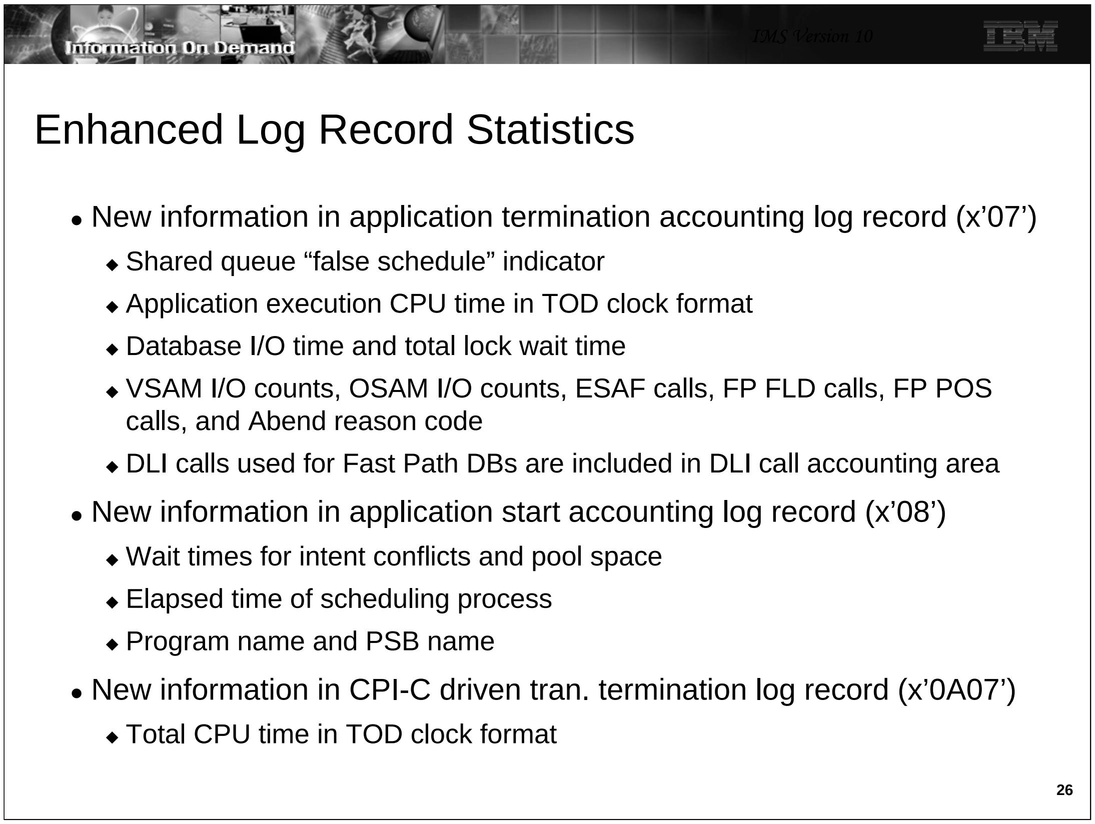

Information On Demand
IMS Version 10
IBM

Enhanced Log Record Statistics

- New information in application termination accounting log record (x’07’)
  - Shared queue “false schedule” indicator
  - Application execution CPU time in TOD clock format
  - Database I/O time and total lock wait time
  - VSAM I/O counts, OSAM I/O counts, ESAF calls, FP FLD calls, FP POS calls, and Abend reason code
  - DLI calls used for Fast Path DBs are included in DLI call accounting area
- New information in application start accounting log record (x’08’)
  - Wait times for intent conflicts and pool space
  - Elapsed time of scheduling process
  - Program name and PSB name
- New information in CPI-C driven tran. termination log record (x’0A07’)
  - Total CPU time in TOD clock format

26

The application termination accounting log record (x’07’) has been enhanced. The DLRFALSE bit in the DLRFLAG2 byte indicates that the scheduling of the transaction was false in a shared queues environment. The DLREXTIM field contains the CPU time for the program in time of day (TOD) format. The DLRTMEIO and DLRTMEPL fields are now populated in all environments. Previously, only DBCTL threads caused these fields to be populated. DLRTMEIO contains the total I/O time for database calls. DLRTMEPL contains the total wait time for database locks. Additional fields have been added. These include counts for VSAM and OSAM database data set reads and writes, ESAF calls, such as those to DB2, and additional Fast Path call counts for FLD calls and POS calls. When an abend occurs the abend reason code is included when available. The DLI call counts for GU, GN. ISRT, etc. now include both full function and Fast Path database calls. The log record is mapped by the DFSLOG07 macro. The offsets to some fields have changed from previous releases.

The application start accounting log record (x’08’) is written for each scheduling of a program. The statistics in the log record have been enhanced. Statistics about the scheduling of the program were written for DBCTL threads in previous releases. IMS V10 writes these statistics for all programs, not just DBCTL threads. These statistics show the time that scheduling waited due to intent conflicts as could occur when a PROCOPT includes E, the time that scheduling waited due to a lack of pool space, such as in the PSB pool, and the elapsed time of the scheduling process. Two new fields have been added. They are LINTPGM which contains the program name and LINTPSB which contains the PSB name. The program name and PSB name were previously available in the LINTSY1 and LINTSY2 fields, however, there were complicated rules governing the contents of these two fields. LINTPGM and LINTPSB greatly simplify the discovery of the program name and PSB name. LINTSY1 and LINTSY2 remain in IMS V10 and contain the same contents as in previous releases. The log record is mapped by the DFSLOG08 macro.

The x’0A07’ log record is written when a CPI-C driven transaction terminates. It includes accounting data similar to that in the x’07’ log records. IMS V10 has added the L0AEXT field to the log record. It contains the CPU time in time of day (TOD) format. Previously, this time was only available in timer units in the L0AETIM field in the log record. The log record is mapped by the DFSLOG0A macro.

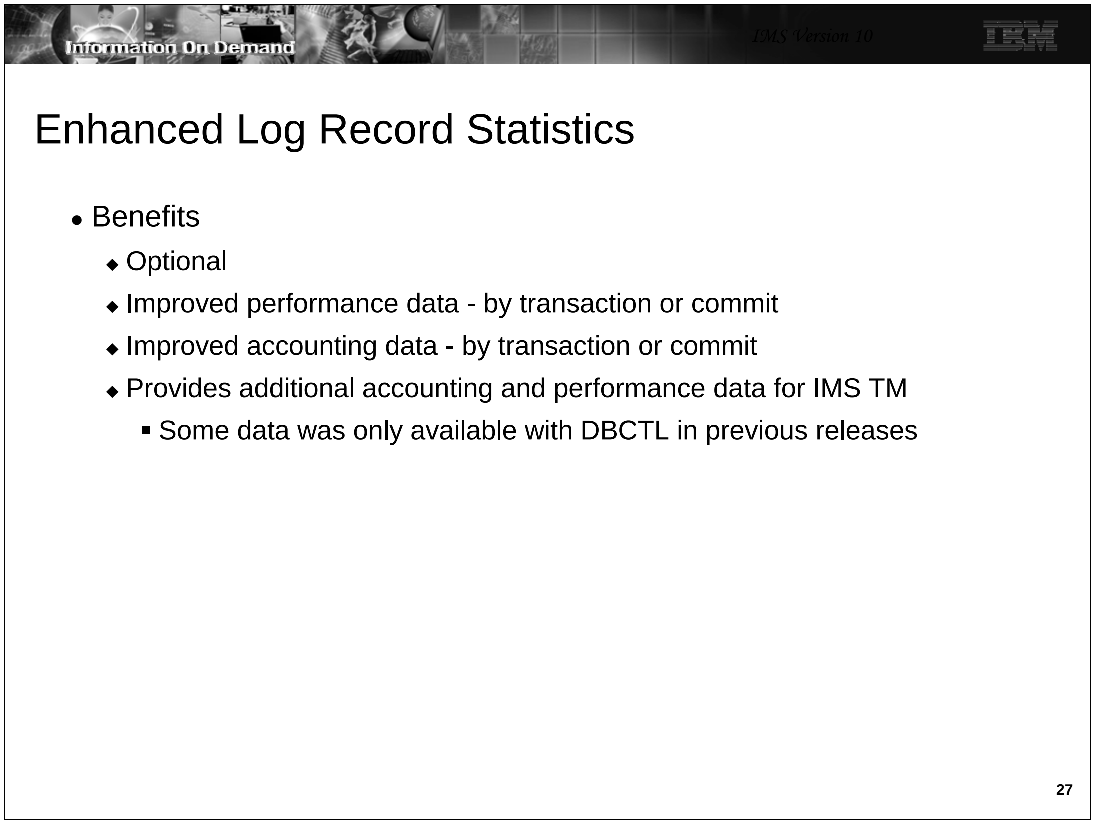

Enhanced Log Record Statistics

- Benefits
  - Optional
  - Improved performance data - by transaction or commit
  - Improved accounting data - by transaction or commit
  - Provides additional accounting and performance data for IMS TM
    - Some data was only available with DBCTL in previous releases

IMS V10 provides enhanced performance and accounting information in its log records. This can provide for granular information for accounting and charge out purposes. It also provides more detailed information for understanding the performance of individual transactions and programs. Some of the information in the log records was previously available only when the IMS Monitor trace was on for IMS TM environments. It was then processed by the IMS Monitor report program or the IMS Performance Analysis tool. IMS V10 provides new options for obtaining this data.

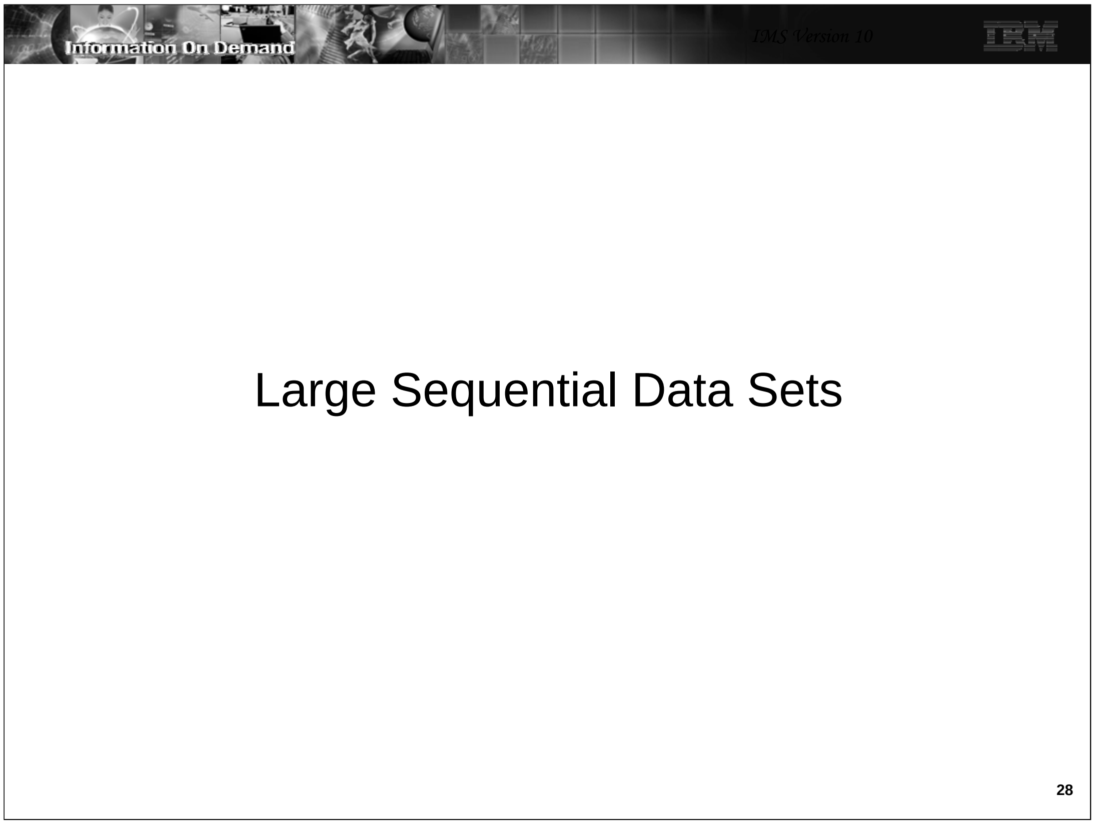

Large Sequential Data Sets

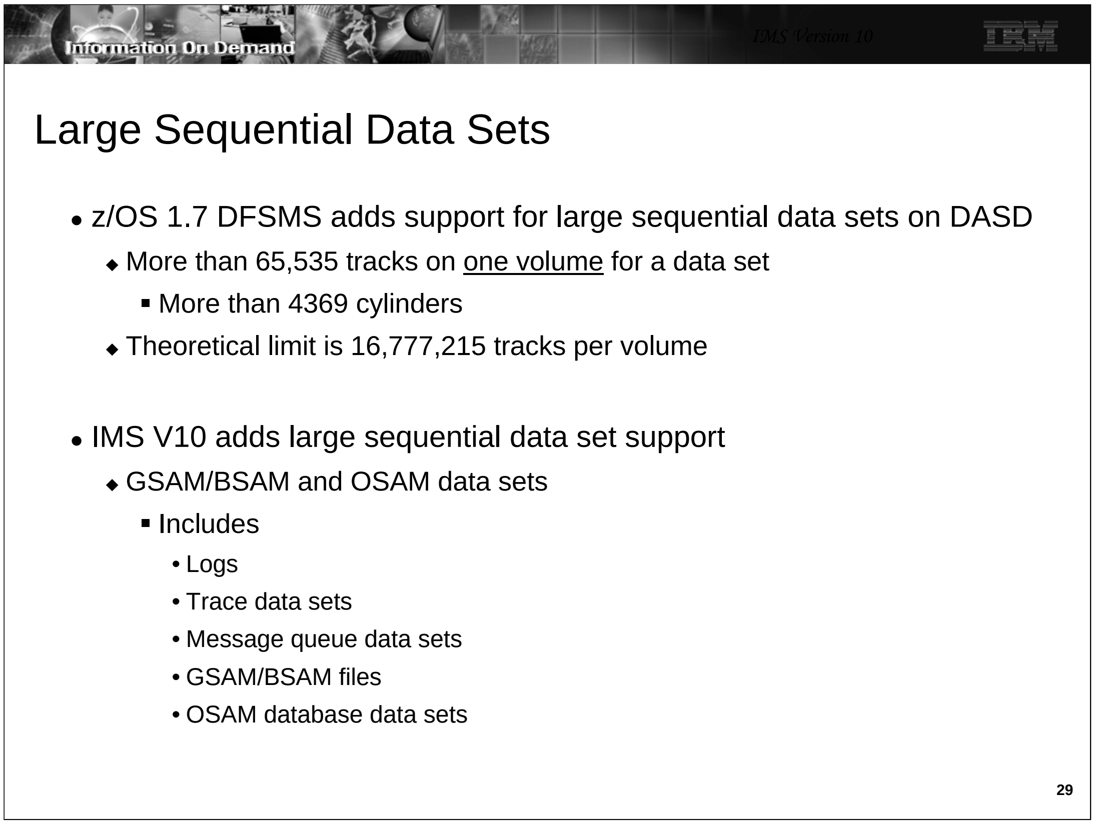
Information On Demand
IMS Version 10
IBM

Large Sequential Data Sets

- z/OS 1.7 DFSMS adds support for large sequential data sets on DASD
  - More than 65,535 tracks on <u>one volume</u> for a data set
    - More than 4369 cylinders
  - Theoretical limit is 16,777,215 tracks per volume

- IMS V10 adds large sequential data set support
  - GSAM/BSAM and OSAM data sets
    - Includes
      - Logs
      - Trace data sets
      - Message queue data sets
      - GSAM/BSAM files
      - OSAM database data sets

29

z/OS 1.7 added support for large sequential data sets on DASD. These data sets may occupy more than 65,535 tracks on a single volume. Without this support, physical sequential data sets (DSORG=PS) were restricted to 65,535 tracks on any volume. Multi-volume data sets could have up to 65,535 tracks on each volume. The architectural (theoretical) limit with large data set support is 16,777,215 tracks per volume. This is far beyond the maximum number of tracks per volume that storage subsystems support.

IMS Version 10 adds support for large sequential data sets. This applies to GSAM/BSAM and OSAM data sets. These include logs (OLDS and SLDS), trace data sets, message queue data sets, GSAM files, and OSAM database data sets.

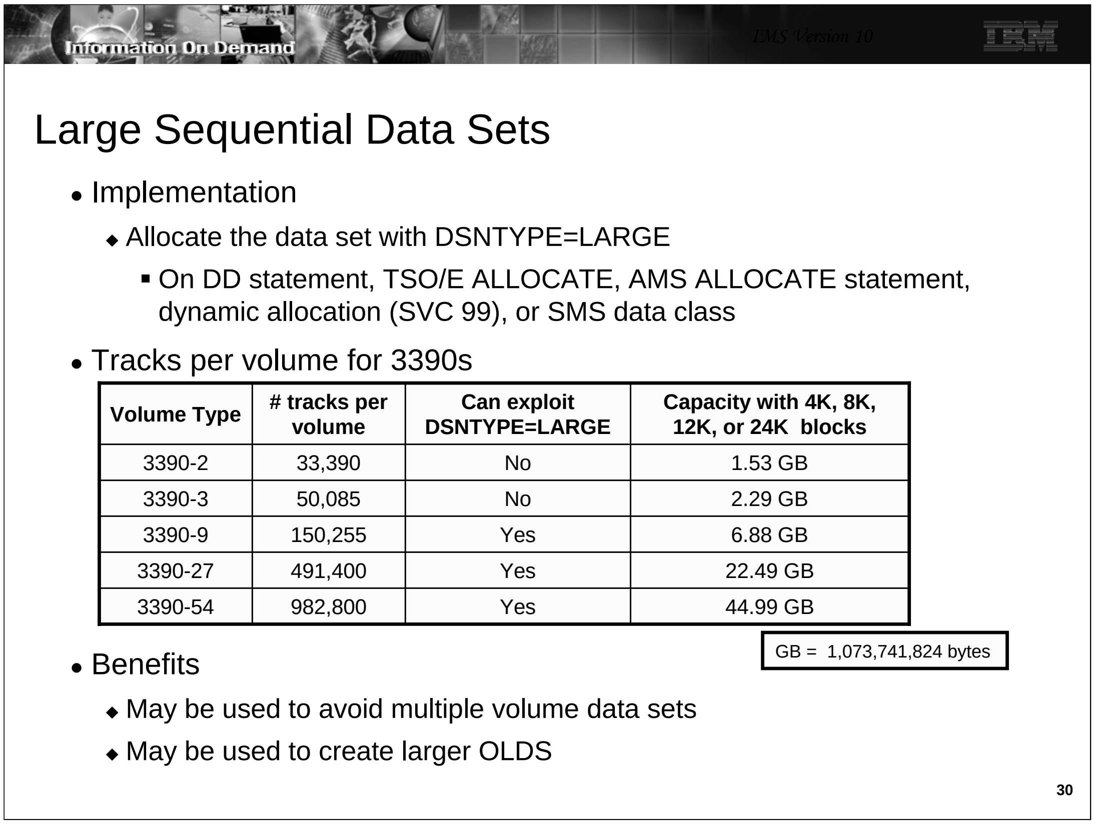
Information On Demand
IMS Version 10

Large Sequential Data Sets

- Implementation
  - Allocate the data set with DSNTYPE=LARGE
    - On DD statement, TSO/E ALLOCATE, AMS ALLOCATE statement, dynamic allocation (SVC 99), or SMS data class
- Tracks per volume for 3390s

Volume Type	# tracks per volume	Can exploit DSNTYPE=LARGE	Capacity with 4K, 8K, 12K, or 24K blocks
3390-2	33,390	No	1.53 GB
3390-3	50,085	No	2.29 GB
3390-9	150,255	Yes	6.88 GB
3390-27	491,400	Yes	22.49 GB
3390-54	982,800	Yes	44.99 GB

GB = 1,073,741,824 bytes

- Benefits
  - May be used to avoid multiple volume data sets
  - May be used to create larger OLDS

30

A data set may use large data set support only if if is allocated with DSNTYPE=LARGE specified. That is, large data set support only applies to data sets that are created with this specification. DSNTYPE=LARGE may be specified on the DD statement, a TSO/E ALLOCATE statement, or an AMS ALLOCATE statement. For dynamic allocation (SVC 99) a DALDSNT text unit may be used to specify LARGE.

The following is an example of the use of a DD statement with DSNTYPE=LARGE specified.

```
//ABCDEF  DD  DSN=IMS.ABC.XYZ,DSNTYPE=LARGE,
//            UNIT=SYSDA,SPACE=(CYL,(4500,100)),
//            DISP=(NEW,CATLG),VOL=SER=LRGVOL1
```

The table is provided to show which volume types may contain large format data sets. 3390-1, 3390-2, and 3390-3 have less than 65,535 tracks per volume. Large format data sets may be defined on them, however, they cannot take advantage of the larger number of tracks per volume. The table also shows the approximate capacity of a volume for certain block sizes. This applies to 4K, 8K, 12K and 24K. Other block sizes may yield more or less capacity due to the number of blocks which will fit on a track. For example, a 3390-9 will hold a 7.45 GB data set with a block size of 26K.

Large format data sets may be used for any log, trace data set, message queue data set, GSAM/BSAM data set or OSAM database data set. They may be used to avoid the requirement for multiple volumes for very large data sets. Since OLDS data sets must reside on only one volume, large format data sets allow OLDS to be larger than previously possible.

OSAM database data sets are limited to a maximum of 8 gigabytes. The large format data set support does not change this limit.

System Utilities Enhancements

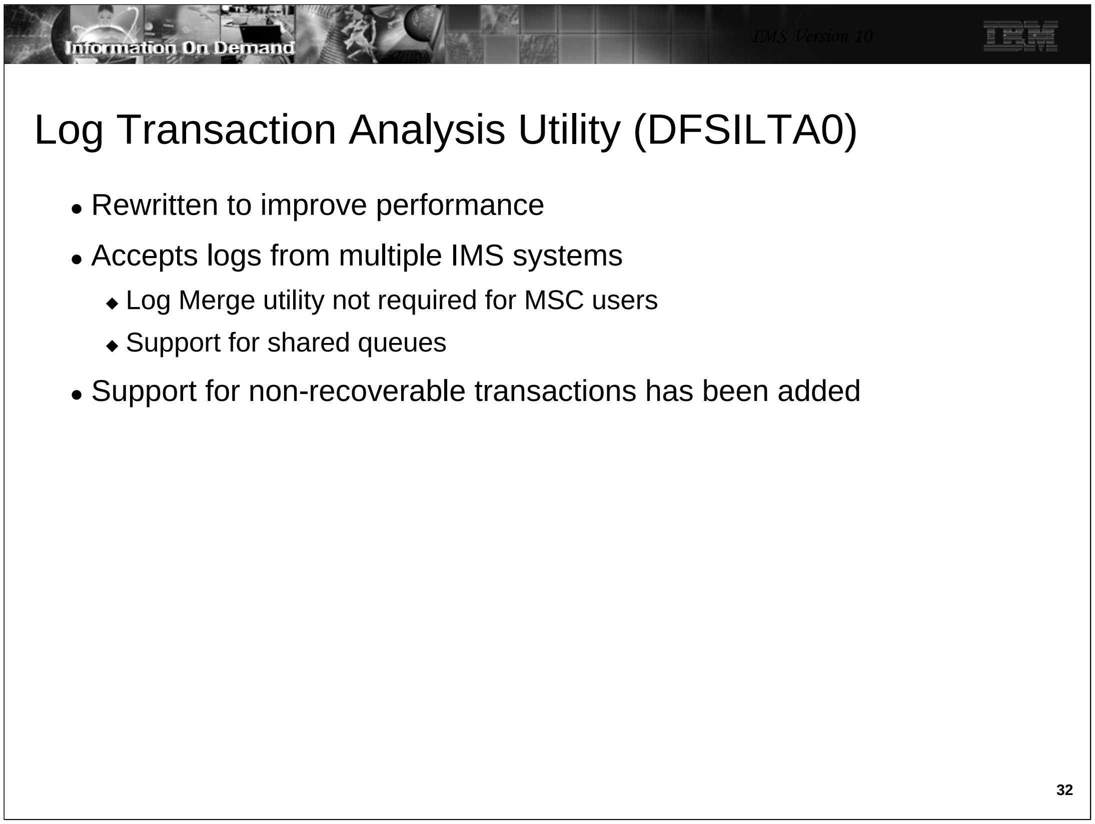
Information On Demand
IMS Version 10
IBM

Log Transaction Analysis Utility (DFSILTA0)

- Rewritten to improve performance
- Accepts logs from multiple IMS systems
  - Log Merge utility not required for MSC users
  - Support for shared queues
- Support for non-recoverable transactions has been added

32

The Log Transaction Analysis utility (DFSILTA0) has several improvements.

In previous releases the utility could produce reports for systems connected by MSC. This required that the user invoke the Log Merge utility to merge the logs from multiple systems and use its output as input to the Log Transaction Analysis utility. The requirement to run Log Merge has been eliminated. Logs from multiple IMS systems may be input to a single execution of the utility. In previous releases the Log Transaction Analysis utility did not support shared queues. This restriction has been eliminated. One execution of the utility can process logs from multiple systems using shared queues.

When logs from multiple systems are input to the utility multiple LOGINxxx DD statements are used. Any valid JCL characters may be used to replace “xxx” in the DD names. The last character in “xxx” is used as the System ID in reports. When multiple log data sets from one IMS system are inputs, they are concatenated under a single DD name.

In previous releases the utility did not report non-recoverable transactions, including APPC messages. Support for reporting non-recoverable transactions has been added in IMS V10.

The Log Transaction Analysis utility has an option to write the output report to an output log data set, to a printer DD, or to both. In previous releases this was controlled by the OUT= parm on the EXEX statement. Both outputs were produced unless the OUT= parameter eliminated them. Possible values were OUT=NOREPORT, OUT=NOLOG, and OUT=(NOREPORT,NOLOG). IMS V10 has eliminated this parm. If you do not want either output just do not include its DD statement. PRINTER is used for the report and LOGOUT is used for the output log data set. The need for the output log data set has been reduced in IMS V10. In previous releases it was used to create a log for input to the Statistical Analysis utility. This is available in IMS V10 but not required.

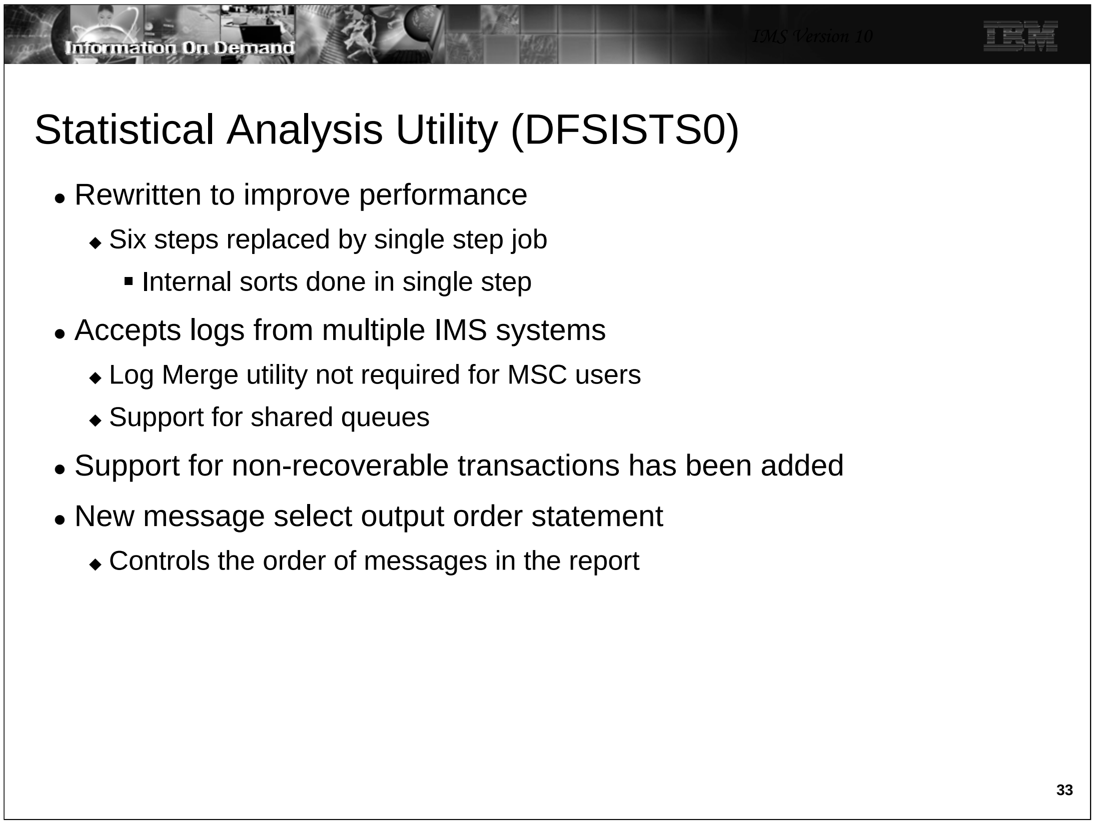

Information On Demand
IMS Version 10

Statistical Analysis Utility (DFSISTS0)

- Rewritten to improve performance
  - Six steps replaced by single step job
    - Internal sorts done in single step
- Accepts logs from multiple IMS systems
  - Log Merge utility not required for MSC users
  - Support for shared queues
- Support for non-recoverable transactions has been added
- New message select output order statement
  - Controls the order of messages in the report

The Statistical Analysis utility (DFSISTS0) has several improvements.

In previous releases the Statistical Analysis utility reports were created by a six step job which included two sorts. The reports are now produced by a single step. This steps invokes internal sorts.

Support for shared queues and MSC is comparable to that in the Log Transaction Analysis utility for IMS V10. This includes the support for multiple LOGINxxx DD statements. When logs from multiple systems are input to the utility multiple LOGINxxx DD statements are used. Any valid JCL characters may be used to replace “xxx” in the DD names. When multiple log data sets from one IMS system are inputs, they are concatenated under a single DD name. The capability to specify multiple input logs eliminates the requirement to use an output log from the Log Transaction Analysis utility when you want a report from multiple IMS systems. In previous releases this was the way that log records from multiple IMS systems were given to the Statistical Analysis utility.

In previous releases the utility did not report non-recoverable transactions, including APPC messages. Support for reporting non-recoverable transactions has been added in IMS V10.

The message select output order statement is a new utility control statement specified in the SYSIN DD data set. The format of the message select output order statement is:

```
ORDER=TIME|SOURCE|TRAN
```

This statement determines the order in which the message select function lists or copies messages. If TIME is specified, the messages are produced in input time order. If SOURCE is specified, the messages are ordered by the originating LTERM names. If TRAN is specified

`TIME` - ordered by original input sequence

`SOURCE` - ordered by originating LTERM

`TRAN` - ordered by initial tran. (message switches will use destination LTERM)

If SOURCE is specified, the messages are produced in order of the message source, either an LTERM or a

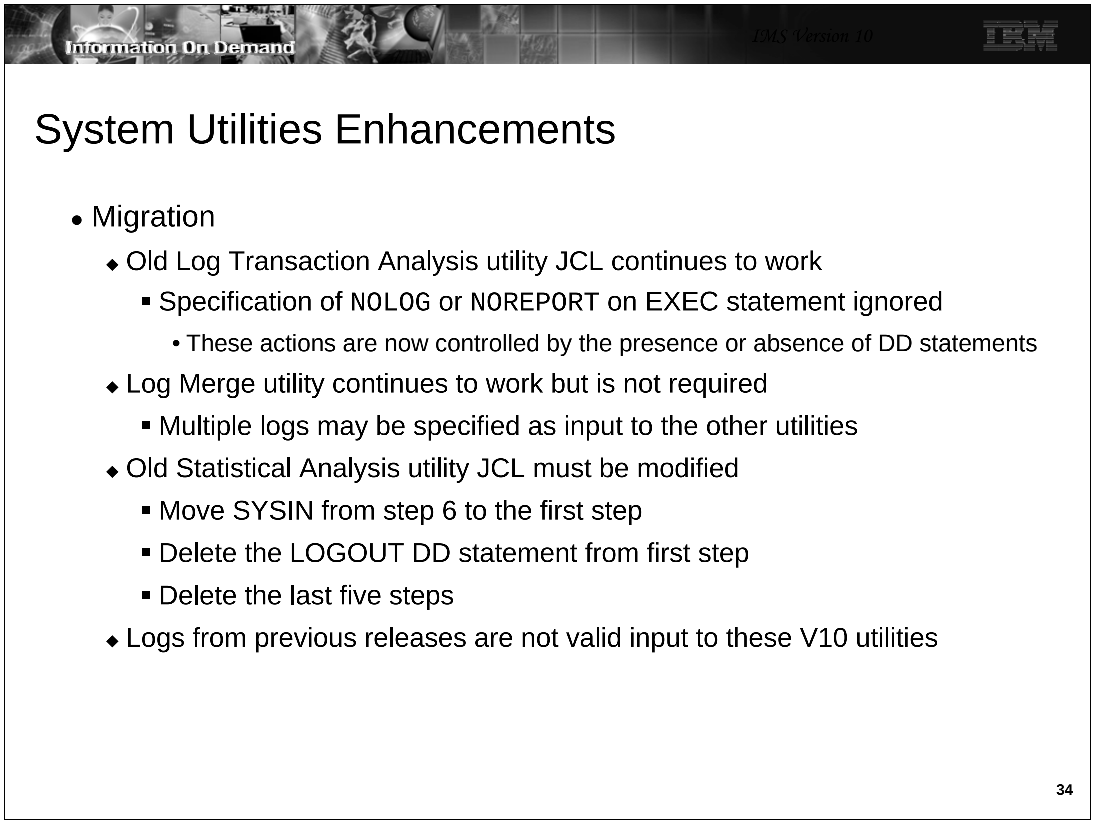
Information On Demand
IMS Version 10

System Utilities Enhancements

- Migration
  - Old Log Transaction Analysis utility JCL continues to work
    - Specification of `NOLOG` or `NOREPORT` on EXEC statement ignored
      - These actions are now controlled by the presence or absence of DD statements
  - Log Merge utility continues to work but is not required
    - Multiple logs may be specified as input to the other utilities
  - Old Statistical Analysis utility JCL must be modified
    - Move SYSIN from step 6 to the first step
    - Delete the LOGOUT DD statement from first step
    - Delete the last five steps
  - Logs from previous releases are not valid input to these V10 utilities

The JCL used in previous releases for the Log Transaction Utility can be used with IMS V10. There is one slight incompatibility. Previously you could specify an OUT= parameter on the EXEC statement. OUT=NOLOG specified that an output report should not be written to the LOGOUT DD data set. OUT=NOREPORT specified that an output report should not be written to the PRINTER DD data set. In IMS V10 if you do not want either of these data sets do not include the DD statement for it. IMS V10 ignores the OUT= parameter on the EXEC statement.

The Log Merge utility was used in previous releases to merge logs from different systems. The output of Log Merge was used as input to the Log Transaction Analysis or Statistical Analysis utilities. The Log Merge utility can be used with IMS V10, but it is not required. Both analysis utilities can merge multiple logs as part of their processing.

The JCL used for the Statistical Analysis utilities in previous releases must be changed for IMS V10. The actions needed to change the JCL are listed on the slide. The control statements were specified in the SYSIN data set of the sixth step. The SYSIN data set is now part of the first (and only) step. The LOGOUT DD in the old first step was used to pass data to the following steps. Since only one step is used in IMS V10, the LOGOUT DD is no longer used. The last five steps are no longer used, so their JCL should be eliminated.

Logs produced by IMS V8 and IMS V9 are not valid input to the IMS V10 Log Transaction Analysis utility and the IMS V10 Statistical Analysis utility.

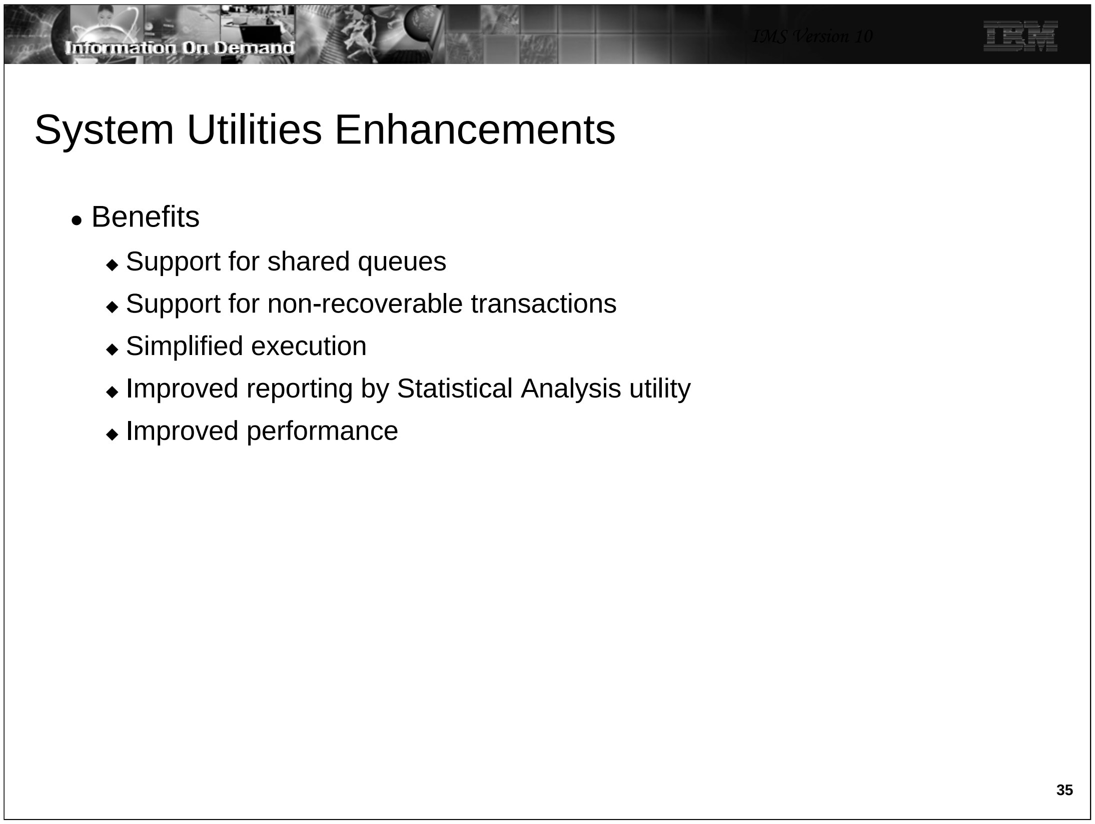

System Utilities Enhancements

- Benefits
  - Support for shared queues
  - Support for non-recoverable transactions
  - Simplified execution
  - Improved reporting by Statistical Analysis utility
  - Improved performance

The Log Transaction Analysis utility (DFSILTA0) and the Statistical Analysis utility (DFSISTS0) enhancements provide several benefits. They now provide information for shared queues systems and for non-recoverable transactions in any systems. Their execution has been simplified in two ways. You can now feed logs from multiple systems into both utilities without first merging the logs with the Log Merge utility. The Statistical Analysis utility is now just a single step execution instead of multiple programs and sorts executed in multiple steps. The Statistical Analysis utility has new reporting options that provide you a choice in how the report lines are ordered. The simplification of execution also provides improved performance since entire data sets are not written in each step and passed to following steps.

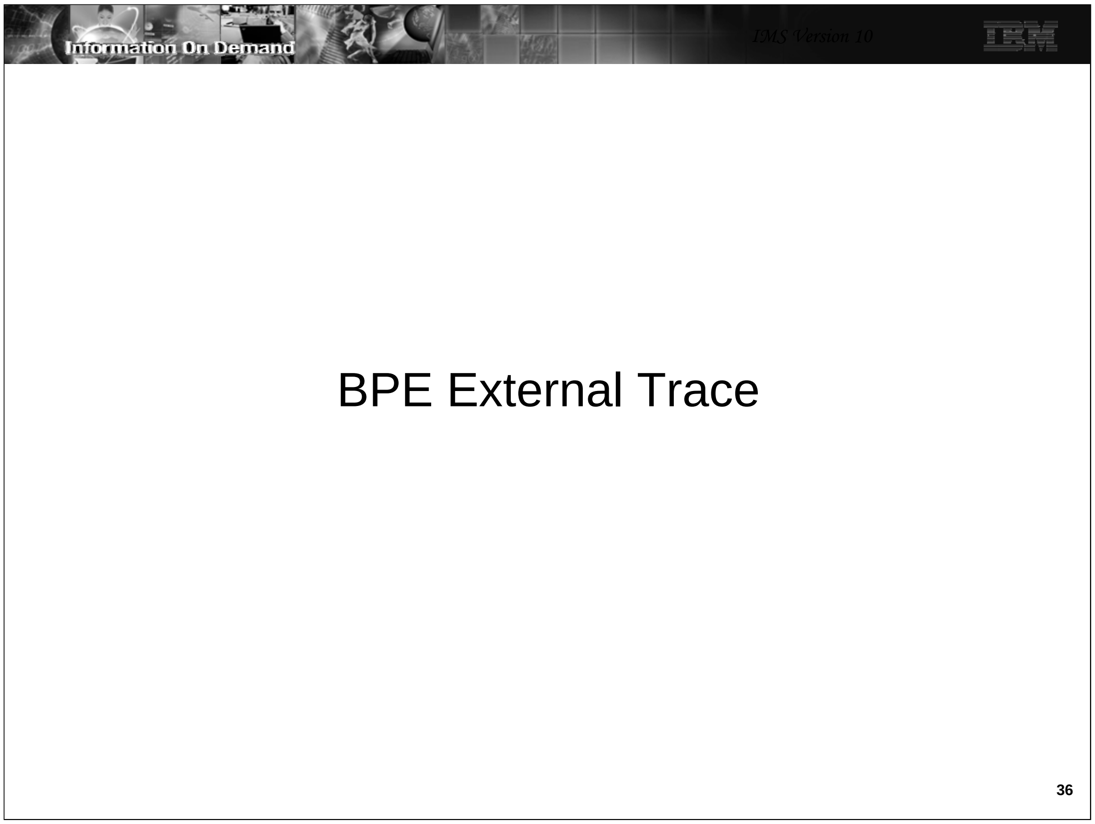

Information On Demand
IMS Version 10

BPE External Trace

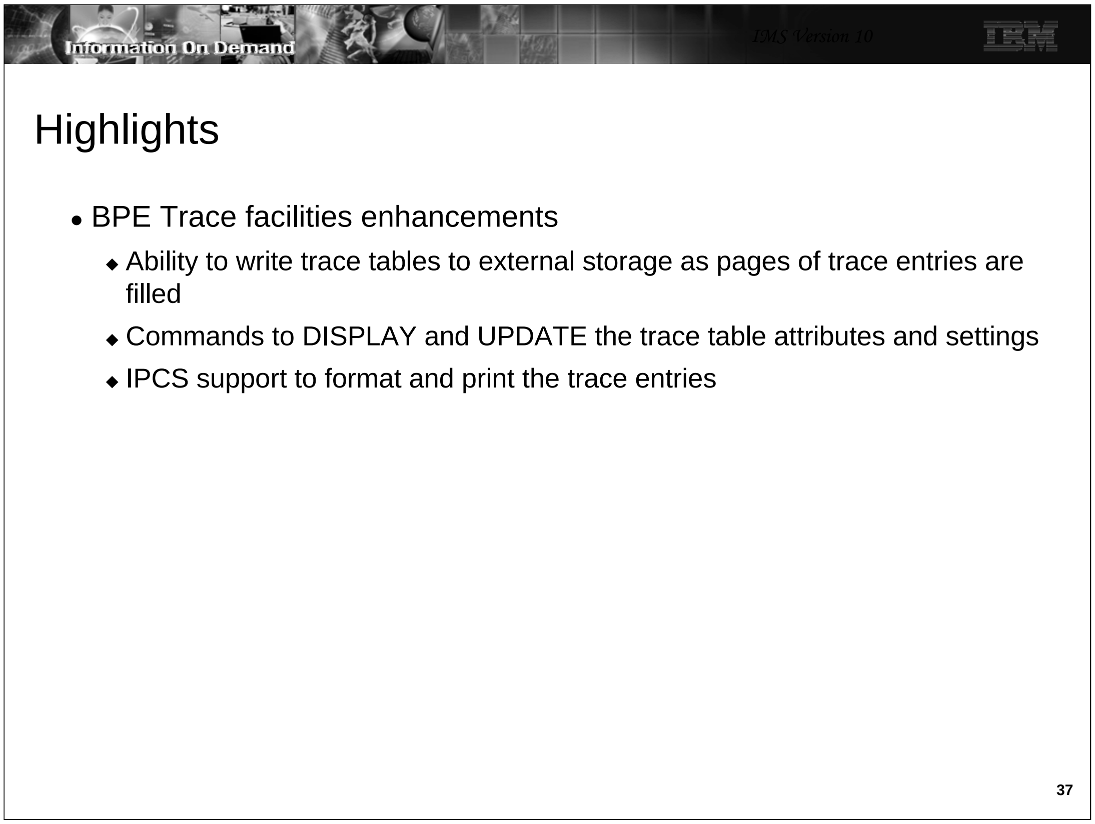

Information On Demand
IMS Version 10
IBM

Highlights

- BPE Trace facilities enhancements
  - Ability to write trace tables to external storage as pages of trace entries are filled
  - Commands to DISPLAY and UPDATE the trace table attributes and settings
  - IPCS support to format and print the trace entries

The Base Primitive Environment (BPE) is a critical component of several IMS-related components such as IMS Connect, the Common Queue Server, and several of the address spaces introduced with the Common Service Layer. In prior releases, problem determination using the BPE traces often required an increase in the number of trace pages and a corresponding requirement to recreate the problem and capture an MVS dump at the correct moment. This process increased the chance of multiple repeats, thereby decreasing productivity and increasing the time to find problem resolution.

IMS V10 enhances the BPE Trace facilities to provide a means to write the trace tables to external storage, Tape or DASD, as pages of trace entries are filled thereby increasing the capacity associated with tracing. Additionally, commands are provided to display the trace table attributes and update the trace level settings. To view data, trace formatting is provided through IPCS exits and formatting routines.

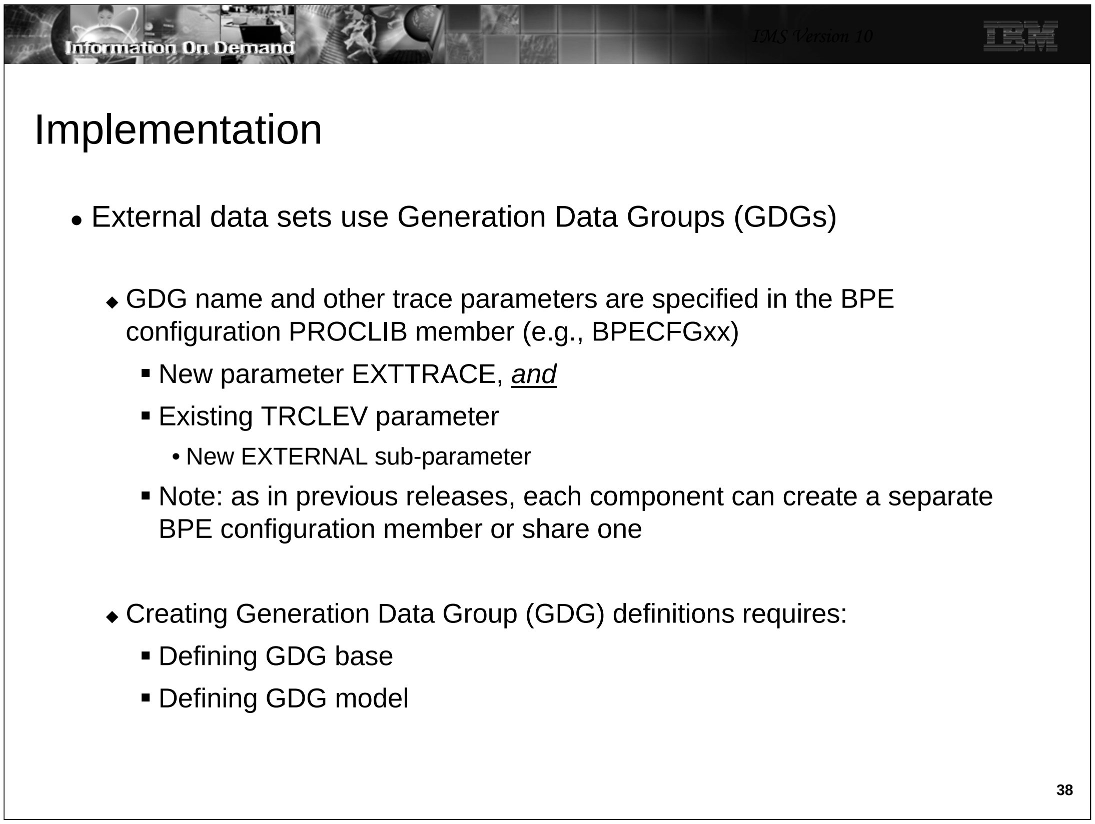
Information On Demand
IMS Version 10
IBM

Implementation

- External data sets use Generation Data Groups (GDGs)
  - GDG name and other trace parameters are specified in the BPE configuration PROCLIB member (e.g., BPECFGxx)
    - New parameter EXTTRACE, *and*
    - Existing TRCLEV parameter
      - New EXTERNAL sub-parameter
    - Note: as in previous releases, each component can create a separate BPE configuration member or share one
  - Creating Generation Data Group (GDG) definitions requires:
    - Defining GDG base
    - Defining GDG model

38

The BPE external trace data sets are defined as GDGs and must be specified in the BPE configuration member in PROCLIB. The primary tasks include the specification of new values in BPECFGxx as well as creating the GDG definitions.

Both TRCLEV with EXTERNAL=YES and the EXTTRACE statement have to be defined for external trace to start recording data. The existing TRCLEV statement is associated with a particular trace table and has been enhanced with the EXTERNAL=YES/NO parameter which specifies whether that trace table is to be externalized or not. The EXTTRACE statement is for the data set definition options and is needed for the function to be activated. If a user specifies, EXTERNAL=YES on a table and does not specify an EXTTRACE statement, an error message is produced.

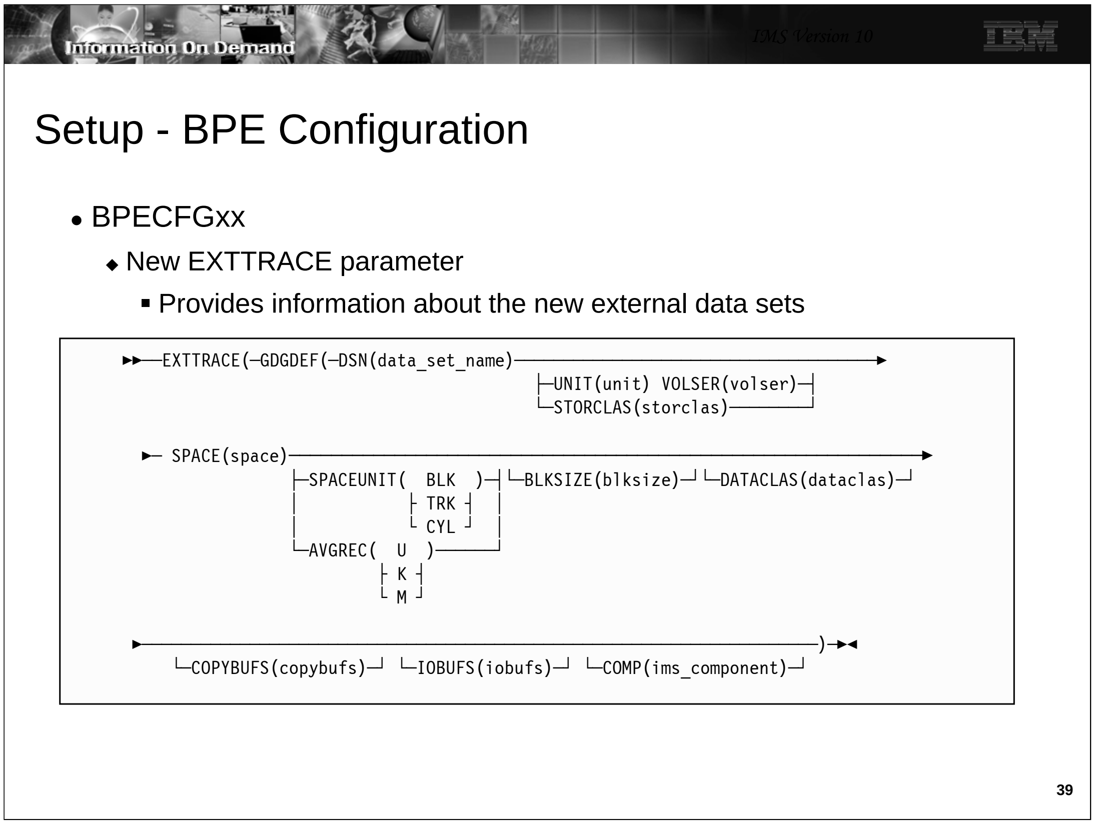

Information On Demand
IMS Version 10
IBM

Setup - BPE Configuration

- BPECFGxx
  - New EXTTRACE parameter
    - Provides information about the new external data sets

```
 ►►──EXTTRACE(─GDGDEF(─DSN(data_set_name)──────────────────────────────►
                                        ├─UNIT(unit) VOLSER(volser)─┤
                                        └─STORCLAS(storclas)────────┘

 ►─ SPACE(space)──────────────────────────────────────────────────────────►
               ├─SPACEUNIT(  BLK  )─┤ └─BLKSIZE(blksize)─┘ └─DATACLAS(dataclas)─┘
               │           ├ TRK ┤  │
               │           └ CYL ┘  │
               └─AVGREC(  U  )──────┘
                       ├ K ┤
                       └ M ┘

 ►────────────────────────────────────────────────────────────────)─►◄
    └─COPYBUFS(copybufs)─┘  └─IOBUFS(iobufs)─┘  └─COMP(ims_component)─┘
```

39

The EXTTRACE keyword definition is an enhancement to the BPECFGxx configuration definition set. One EXTTRACE keyword can be specified for each IMS component running with BPE. Alternatively, one EXTTRACE definition can be shared among multiple IMS components. If the COMP specification is omitted, the associated keyword applies to all components.

•GDGDEF(DSN(data_set_name),....) This is the name of the GDG data set. It does not include the relative generation number in parentheses. It does not include the absolute generation number. (For definition of absolute and relative generation number, refer z/OS JCL Reference).

•Location: Choices include STORCLAS(class_name), or, (UNIT(unit_name) and VOLSER(serial_number)

SPACE(space): specify number of units of space to be allocated for a data set

SPACEUNIT(BLK,TRK,CYL): quantity value of the SPACE operand as the number of blocks, tracks or cylinders.

AVGREC(U | K | M), in SMS, determines the size of the dataset allocation. The following are the valid values for AVGREC: (U - Use the primary space quantity specified on the SPACE operand; K - Multiply the primary space quantity specified on the SPACE operand by 1024 (1 K); M - Multiply the primary space quantity specified on the SPACE operand by 1,048,576 (1 M)).

•BLKSIZE(blksize): specifies the block size for the data set. The maximum allowable decimal value for block size is 32760.

•DATACLAS(dataclas): specifies the data class for the data set .

•COPYBUFS - are the buffers that will be used to copy data (trace table entries) from BPE tables. Data from these buffers are then copied into IOBUFS. The default value for number of buffers is 15 with max value of 64,000.

•IOBUFS - are the buffers that will copy data into external trace data sets. Number of IOBUFS defined by default is 2 and max value is 99.

•COMP(ims_component) - Specifies the IMS component to which the EXTTRACE statement applies. Possible values are: (CQS, HWS, OM, RM, SCI).

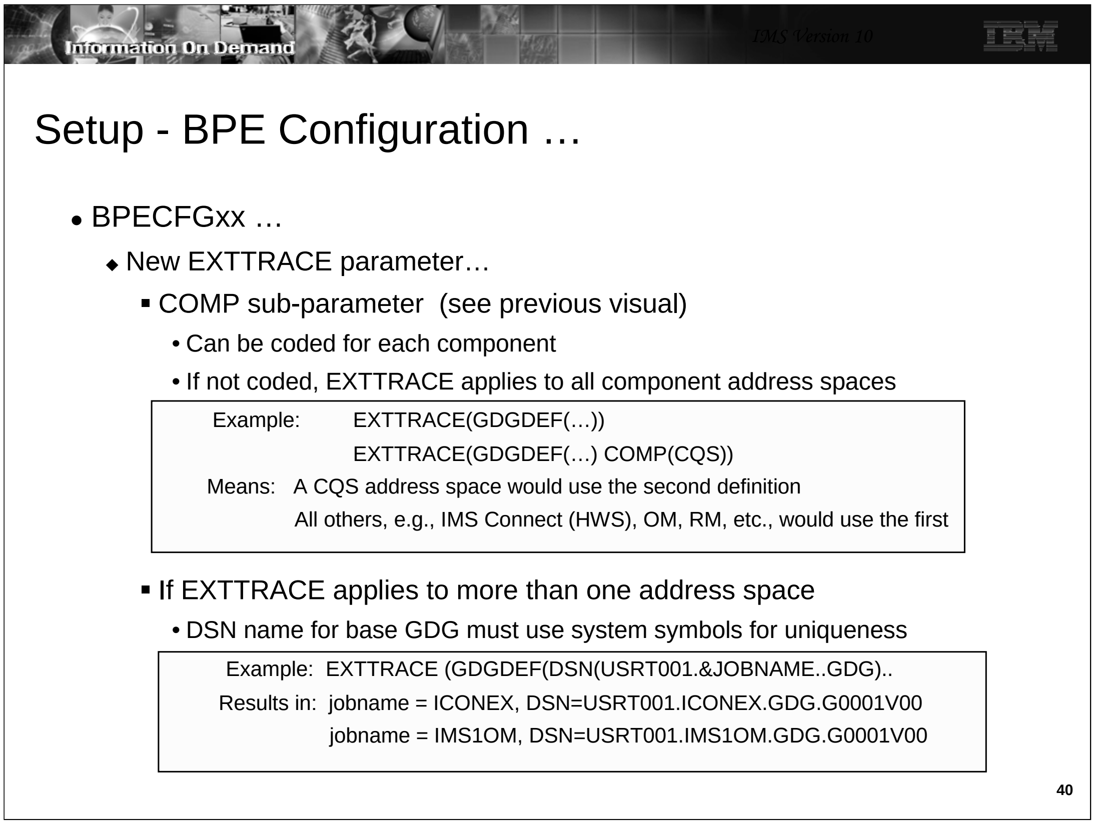
Information On Demand
IMS Version 10
IBM

Setup - BPE Configuration …

- BPECFGxx …
  - New EXTTRACE parameter…
    - COMP sub-parameter (see previous visual)
      - Can be coded for each component
      - If not coded, EXTTRACE applies to all component address spaces

```
 Example:       EXTTRACE(GDGDEF(…))
                EXTTRACE(GDGDEF(…) COMP(CQS))
Means:   A CQS address space would use the second definition
         All others, e.g., IMS Connect (HWS), OM, RM, etc., would use the first
```

- If EXTTRACE applies to more than one address space
  - DSN name for base GDG must use system symbols for uniqueness

```
  Example:  EXTTRACE (GDGDEF(DSN(USRT001.&JOBNAME..GDG)..
Results in:  jobname = ICONEX, DSN=USRT001.ICONEX.GDG.G0001V00
             jobname = IMS1OM, DSN=USRT001.IMS1OM.GDG.G0001V00
```

40

If COMP is not coded, then the EXTTRACE statement applies to all component address space types that are using the BPE proclib member, and that do not have an EXTTRACE statement with a specific COMP specified for their component type. For example, for a BPE configuration PROCLIB member with the following statements:

```
EXTTRACE( GDGDEF( ... ) )
EXTTRACE( GDGDEF( ... )  COMP(CQS))
```

If a CQS address space were started, it would use the external trace definitions from the second EXTTRACE statement. If any other type of address space were started, it would use the external trace definition from the first EXTTRACE statement. Note that if you have the EXTTRACE apply to more than one address space, you must use z/OS system symbols in the data set name (DSN) for the GDG base name, to ensure that each data set has a unique name. For example:

DSN(USRT001.&JOBNAME..GDG) and jobname=ICONEX then DSN=USRT001.ICONEX.GDG.G0001V00
DSN(USRT001.&JOBNAME..GDG) and jobname=IMS1OM then DSN=USRT001.IMS1OM.GDG.G0001V00
…

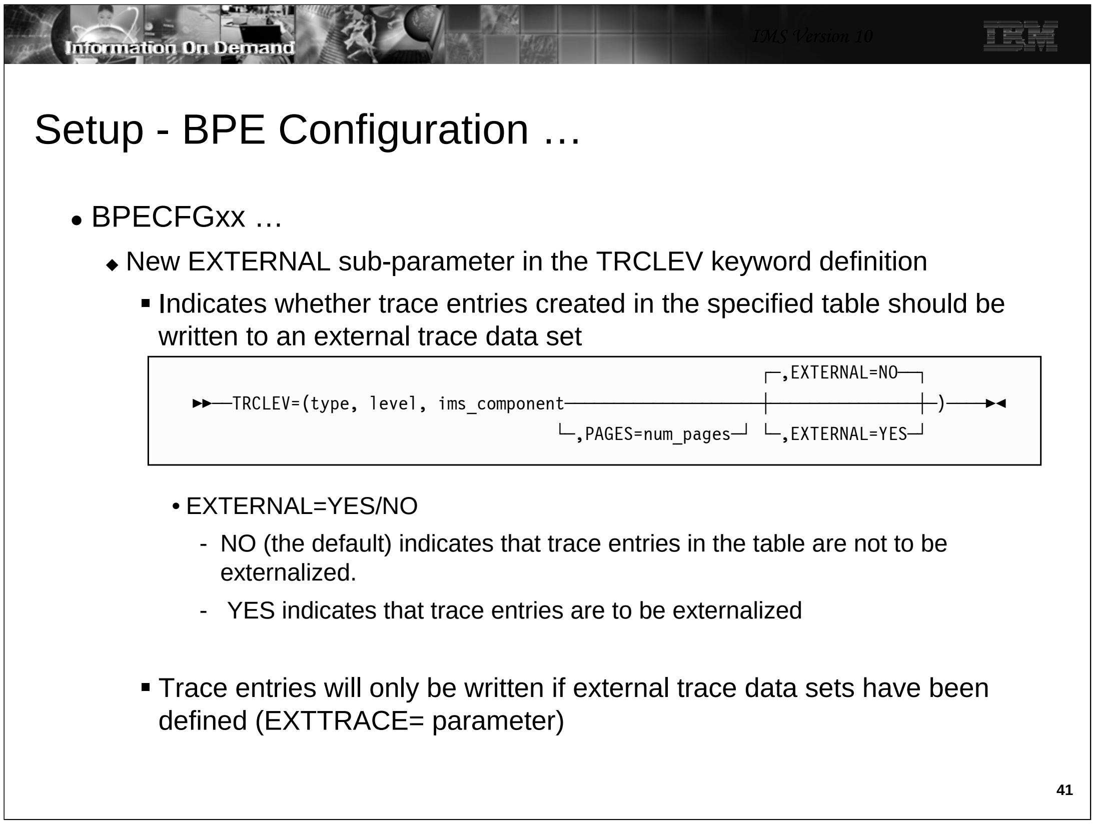
Information On Demand
IMS Version 10

Setup - BPE Configuration …

- BPECFGxx …
  - New EXTERNAL sub-parameter in the TRCLEV keyword definition
    - Indicates whether trace entries created in the specified table should be written to an external trace data set

```
                                                           ┌─,EXTERNAL=NO──┐
►►──TRCLEV=(type, level, ims_component─────────────────────┼───────────────┼─)───►◄
                                      └─,PAGES=num_pages─┘ └─,EXTERNAL=YES─┘
```

      - EXTERNAL=YES/NO
        - NO (the default) indicates that trace entries in the table are not to be externalized.
        - YES indicates that trace entries are to be externalized
    - Trace entries will only be written if external trace data sets have been defined (EXTTRACE= parameter)

The new EXTERNAL sub-parameter in the TRCLEV keyword definition provides the specification of whether or not the trace entries are to be externalized.

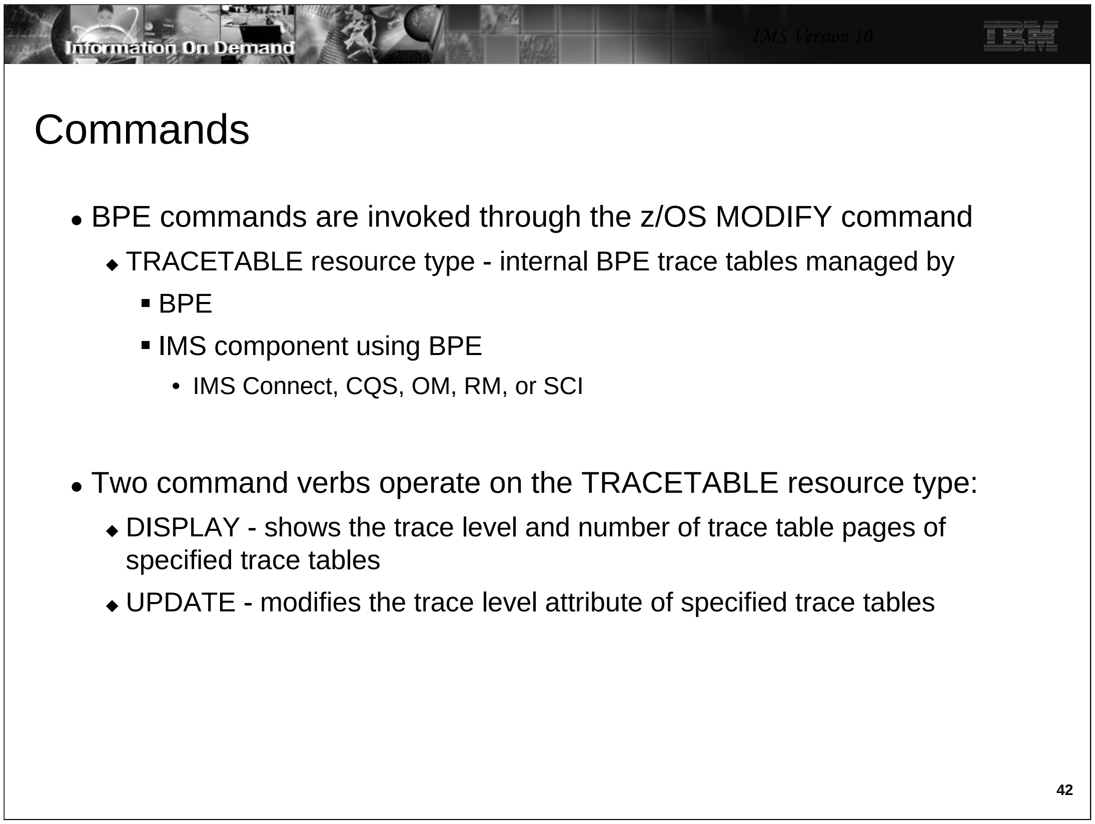

Commands

- BPE commands are invoked through the z/OS MODIFY command
  - TRACETABLE resource type - internal BPE trace tables managed by
    - BPE
    - IMS component using BPE
      - IMS Connect, CQS, OM, RM, or SCI

- Two command verbs operate on the TRACETABLE resource type:
  - DISPLAY - shows the trace level and number of trace table pages of specified trace tables
  - UPDATE - modifies the trace level attribute of specified trace tables

BPE commands are invoked through the z/OS MODIFY command.

The TRACETABLE resource type refers to the internal BPE-managed trace tables defined either by BPE (for example: DISP, CBS, STG, LATC), or by the IMS component using BPE (for example: IMS Connect, CQS, OM, RM, SCI).

Two command verbs operate on the TRACETABLE resource type: DISPLAY and UPDATE.

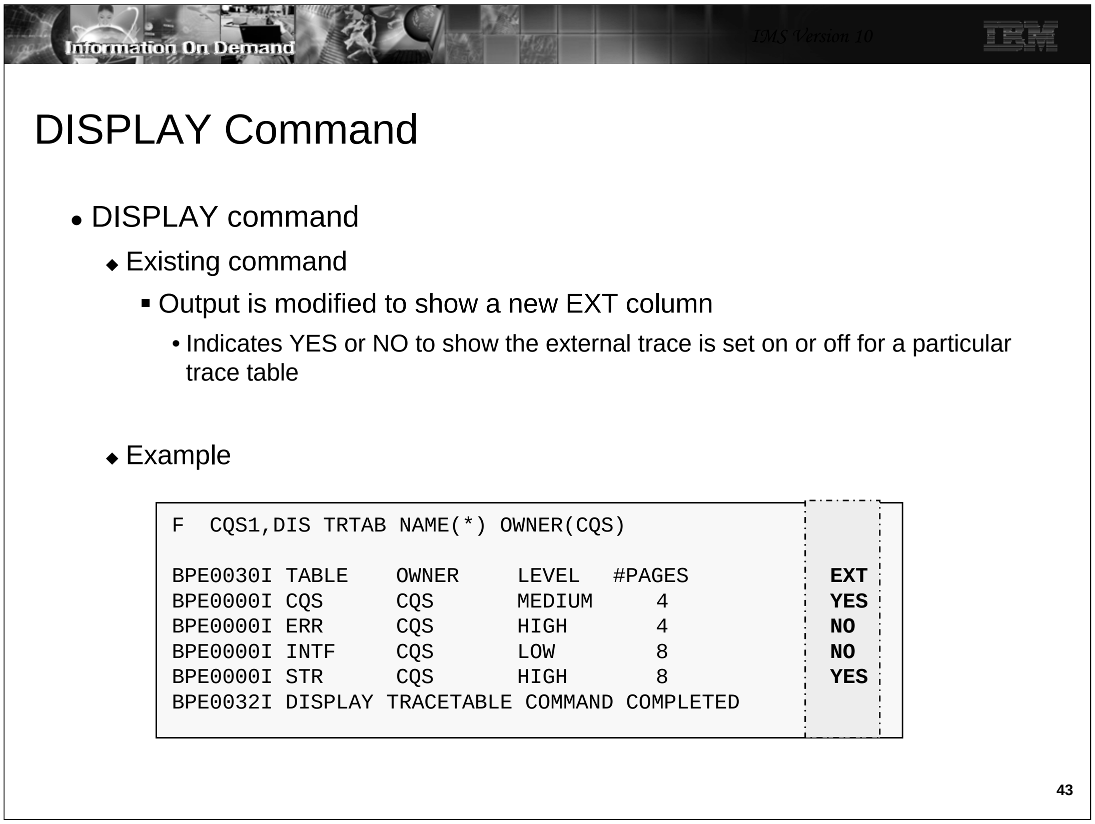
Information On Demand
IMS Version 10

DISPLAY Command

- DISPLAY command
  - Existing command
    - Output is modified to show a new EXT column
      - Indicates YES or NO to show the external trace is set on or off for a particular trace table
  - Example

```
F  CQS1,DIS TRTAB NAME(*) OWNER(CQS)

BPE0030I TABLE    OWNER     LEVEL   #PAGES          EXT
BPE0000I CQS      CQS       MEDIUM     4            YES
BPE0000I ERR      CQS       HIGH       4            NO
BPE0000I INTF     CQS       LOW        8            NO
BPE0000I STR      CQS       HIGH       8            YES
BPE0032I DISPLAY TRACETABLE COMMAND COMPLETED
```

The BPE DISPLAY TRACETABLE command is an existing BPE command that is used to display the current attribute settings for the requested trace tables. IMS V10 modifies the command output by adding a new column to indicate, YES or NO, whether the external trace is set on or off for a particular trace table.

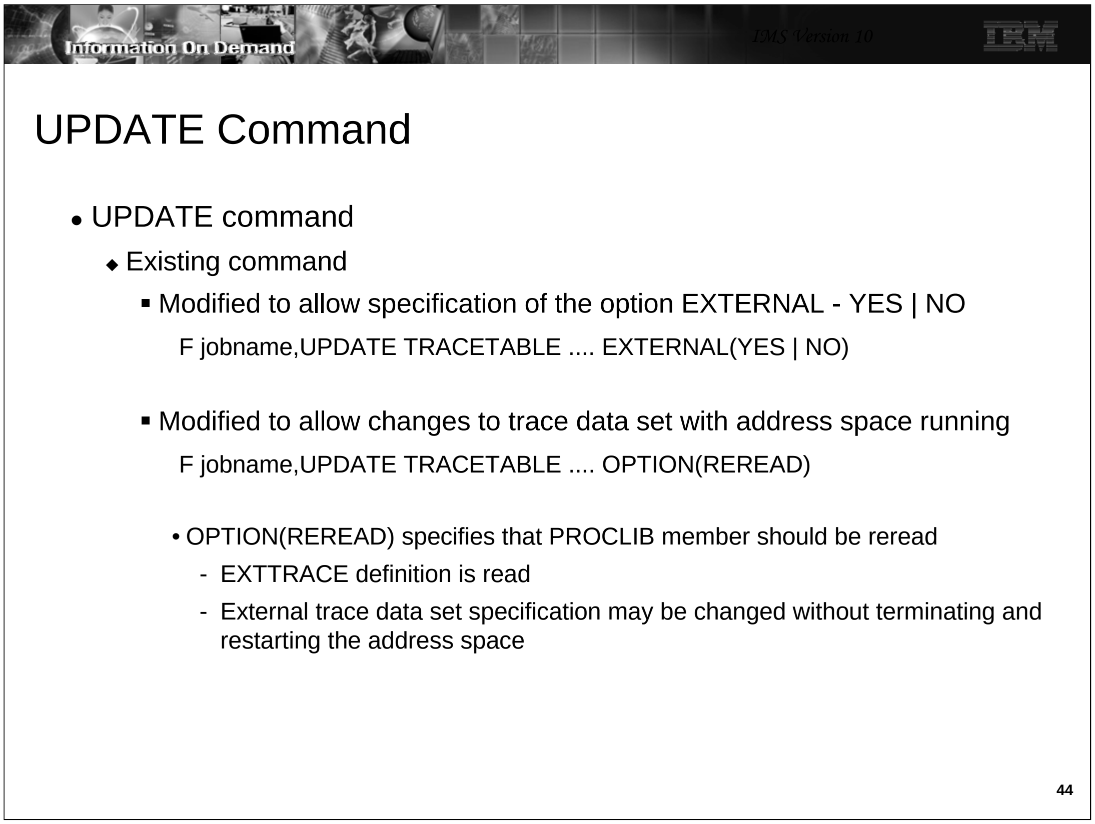

Information On Demand
IMS Version 10
IBM

UPDATE Command

- UPDATE command
  - Existing command
    - Modified to allow specification of the option EXTERNAL - YES | NO

      F jobname,UPDATE TRACETABLE .... EXTERNAL(YES | NO)

    - Modified to allow changes to trace data set with address space running

      F jobname,UPDATE TRACETABLE .... OPTION(REREAD)

      - OPTION(REREAD) specifies that PROCLIB member should be reread
        - EXTTRACE definition is read
        - External trace data set specification may be changed without terminating and restarting the address space

The BPE UPDATE TRACETABLE command is used to change the trace level setting for the requested trace tables. The command syntax for UPDATE TRACETABLE has been extended to allow specification of the option EXTERNAL from YES to NO or vice versa.

EXTERNAL is an optional parameter that specifies whether or not the trace entries for the trace tables affected by the command should be written to the external trace data set. If the keyword is not specified on an UPD TRTAB command, then there is no change made to the external trace setting for the affected tables. Note that specification of EXTERNAL(YES) requires that the BPE configuration PROCLIB member contain the parameters to define a trace data set.

The command syntax for UPDATE TRACETABLE also has been extended to allow specification of the option OPTION(REREAD). REREAD is the only valid value for OPTION. REREAD causes the PROCLIB member to be reread for the information in the specified EXTTRACE statement. This allows the external trace data set specification to be changed or added without terminating the address space and restarting it.

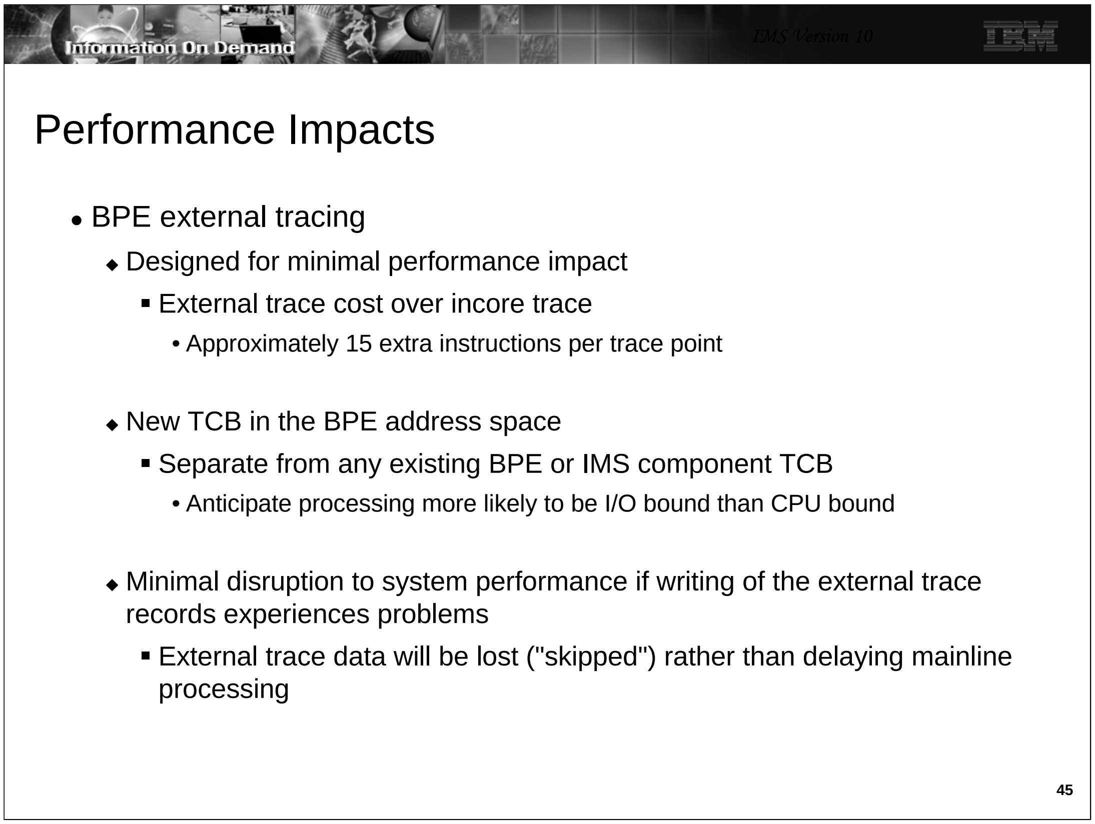

Performance Impacts

- BPE external tracing
  - Designed for minimal performance impact
    - External trace cost over incore trace
      - Approximately 15 extra instructions per trace point
  - New TCB in the BPE address space
    - Separate from any existing BPE or IMS component TCB
      - Anticipate processing more likely to be I/O bound than CPU bound
  - Minimal disruption to system performance if writing of the external trace records experiences problems
    - External trace data will be lost ("skipped") rather than delaying mainline processing

Performance is always a consideration especially when tracing is activated. The BPE external trace function is designed to minimize the performance overhead incurred at trace points in the mainline code. Running with external tracing enabled results in an approximate cost of 15 extra instructions per trace point compared to running with incore tracing.

The majority of the processing required to externalize BPE trace data occurs under a new TCB in the BPE address space. This TCB is separate from any existing BPE or IMS component TCB. Work under existing TCBs, therefore, is not directly impacted by work performed by the new external trace TCB. The processing under the external trace TCB is more likely to be I/O bound than CPU bound.

If the writing of the external trace records cannot keep up with the rate of the generation of the records, external trace data will be lost ("skipped"), rather than delaying mainline processing. The goal is to make running with BPE external trace enabled minimally disruptive to system performance.

Benefits

- BPE external trace enhancements that
  - Provide the ability to trace activities over longer time periods
  - Provide an IPCS interface to the trace data rather than having to take a console dump in order to look at the BPE traces for problem diagnosis

- Result in
  - More timely access to diagnostic data
  - Minimizes the need to recreate problem scenarios

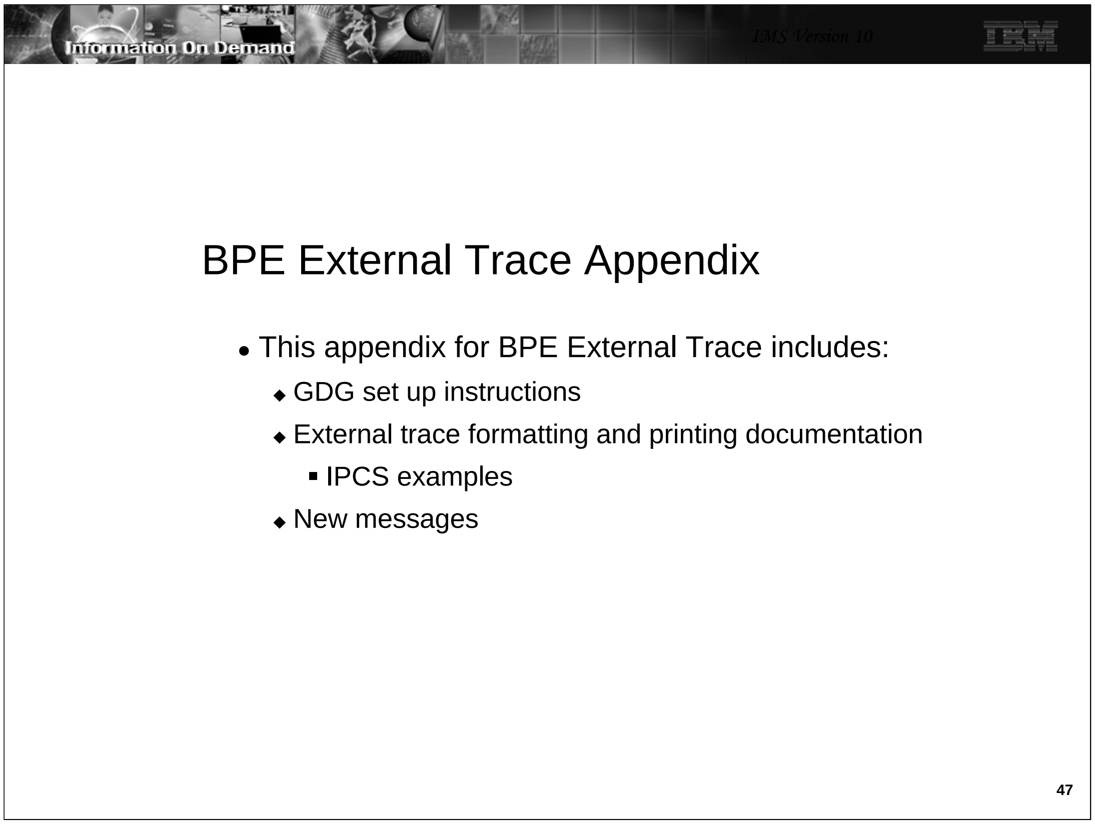
Information On Demand
IMS Version 10

BPE External Trace Appendix

- This appendix for BPE External Trace includes:
  - GDG set up instructions
  - External trace formatting and printing documentation
    - IPCS examples
  - New messages

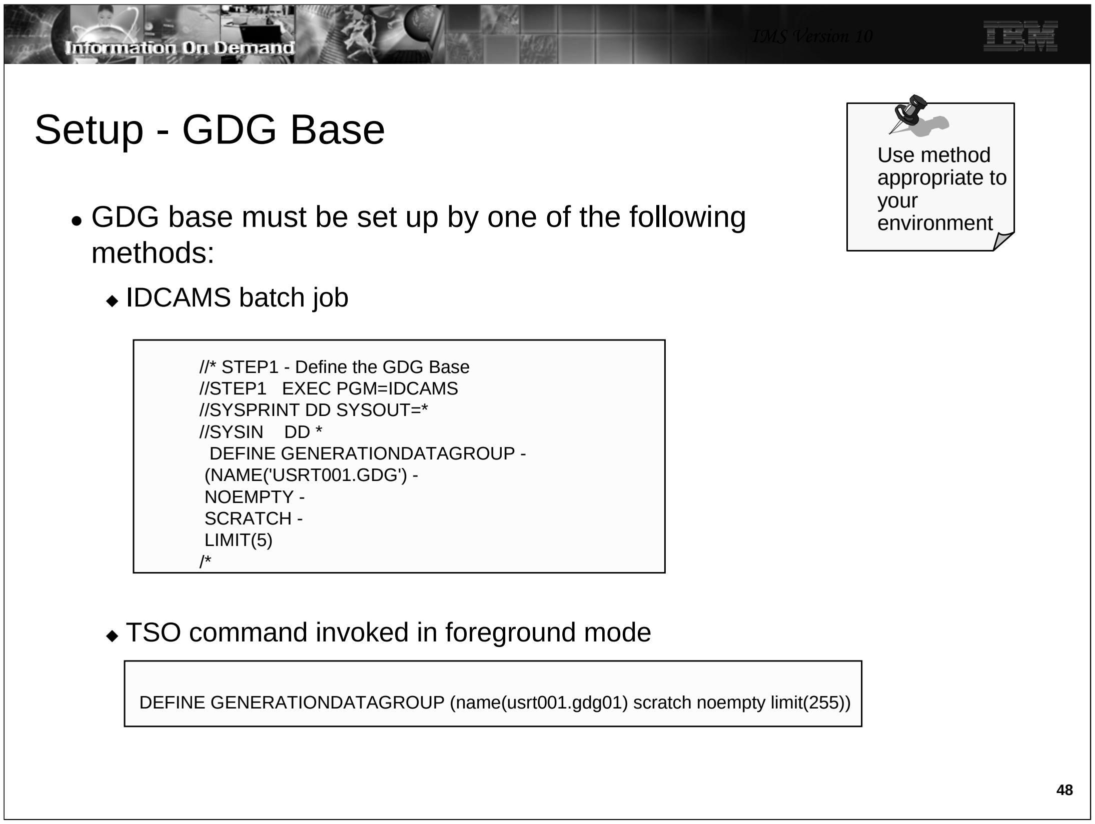

Setup - GDG Base

Use method appropriate to your environment

- GDG base must be set up by one of the following methods:
  - IDCAMS batch job

```
//* STEP1 - Define the GDG Base
//STEP1   EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
//SYSIN    DD *
  DEFINE GENERATIONDATAGROUP -
 (NAME('USRT001.GDG') -
 NOEMPTY -
 SCRATCH -
 LIMIT(5)
/*
```

  - TSO command invoked in foreground mode

```
DEFINE GENERATIONDATAGROUP (name(usrt001.gdg01) scratch noempty limit(255))
```

Users must define the GDG base prior to using BPE external trace functionality. This definition can be performed in any of the following ways:

1) IDCAMS batch job

2) TSO command invoked in the foreground

3) TSO command invoked in a batch job

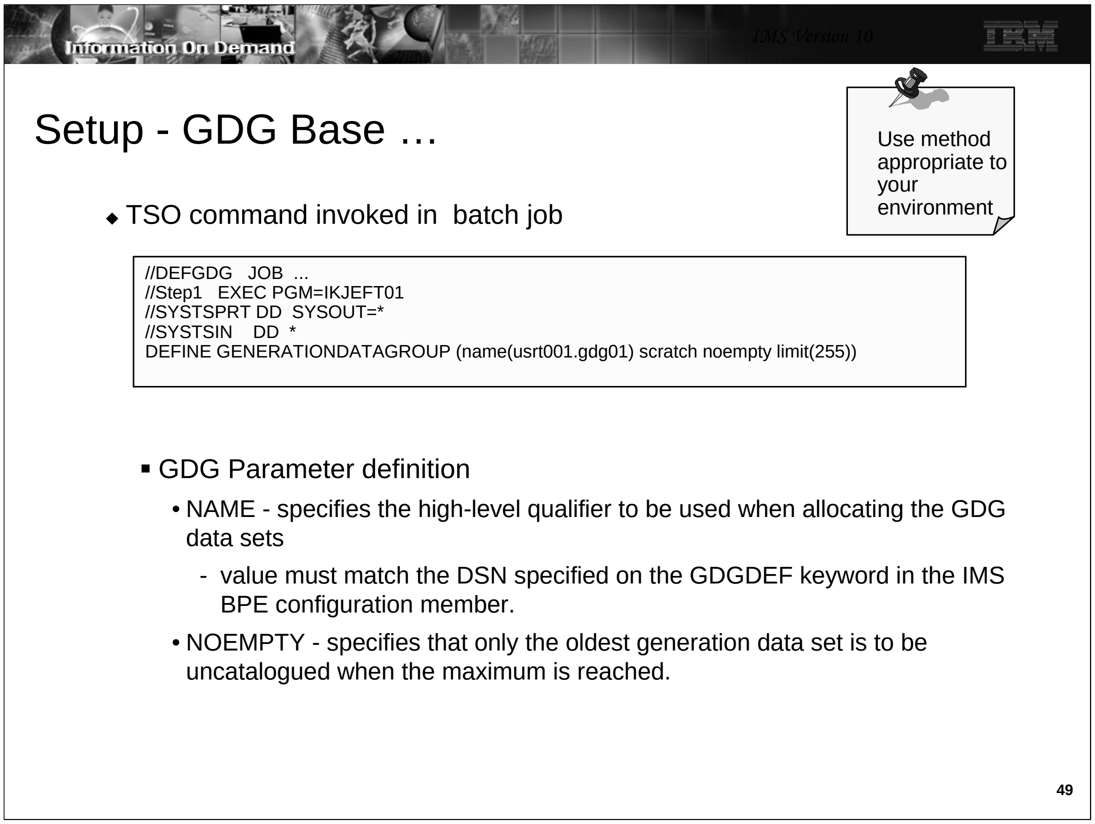

Setup - GDG Base …

Use method appropriate to your environment

- TSO command invoked in batch job

```
//DEFGDG   JOB  ...
//Step1   EXEC PGM=IKJEFT01
//SYSTSPRT DD  SYSOUT=*
//SYSTSIN    DD  *
DEFINE GENERATIONDATAGROUP (name(usrt001.gdg01) scratch noempty limit(255))
```

- GDG Parameter definition
  - NAME - specifies the high-level qualifier to be used when allocating the GDG data sets
    - value must match the DSN specified on the GDGDEF keyword in the IMS BPE configuration member.
  - NOEMPTY - specifies that only the oldest generation data set is to be uncatalogued when the maximum is reached.

Regardless of which method is chosen, the GDG parameter definition must specify a value for NAME which is the high level qualifier that will be used when allocating the DGD data sets. This value must match the value in the DSN parameter of the EXTTRACE keyword definition in the BPE configuration member.

The specification of NOEMPTY provides for the uncataloguing of the oldest GDG data set when the maximum number is reached.

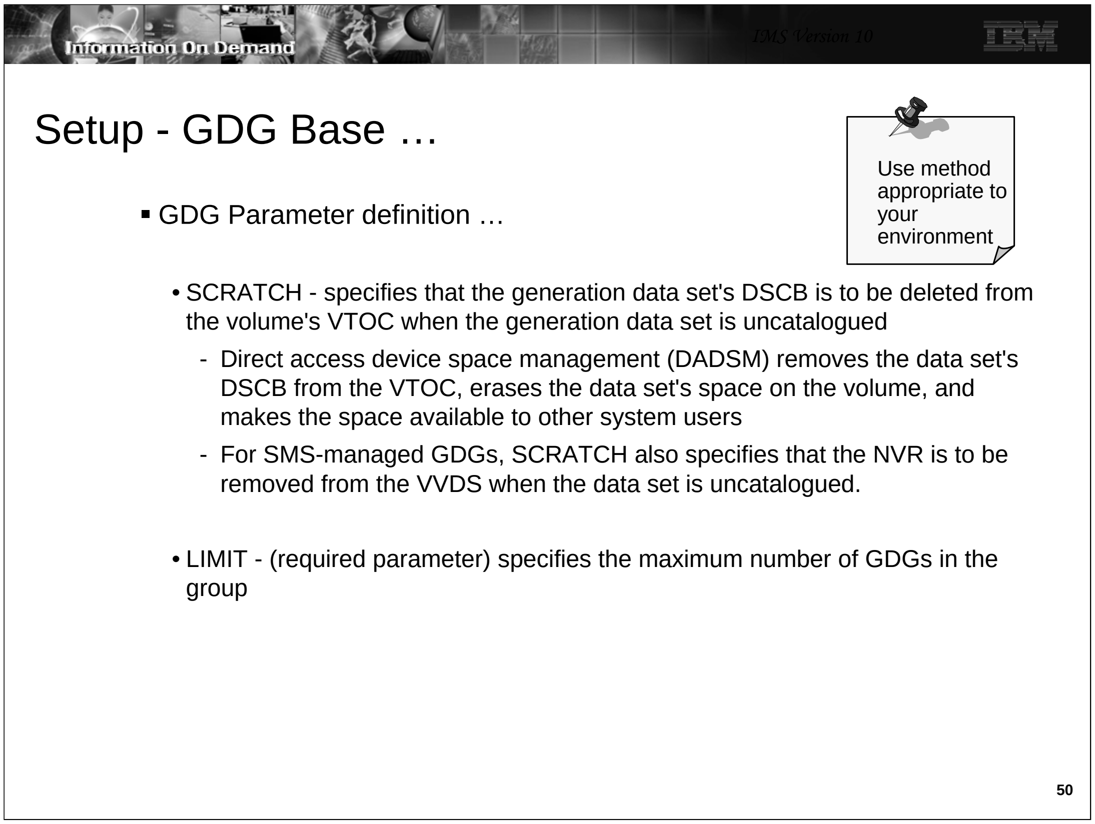

Information On Demand

IMS Version 10

Setup - GDG Base …

Use method appropriate to your environment

- GDG Parameter definition …
  - SCRATCH - specifies that the generation data set's DSCB is to be deleted from the volume's VTOC when the generation data set is uncatalogued
    - Direct access device space management (DADSM) removes the data set's DSCB from the VTOC, erases the data set's space on the volume, and makes the space available to other system users
    - For SMS-managed GDGs, SCRATCH also specifies that the NVR is to be removed from the VVDS when the data set is uncatalogued.
  - LIMIT - (required parameter) specifies the maximum number of GDGs in the group

50

SCRATCH specifies that the generation data set's DSCB is to be deleted from the volume's VTOC when the generation data set is uncatalogued. Direct access device space management (DADSM) removes the data set's DSCB from the VTOC, erases the data set's space on the volume, and makes the space available to other system users. The generation data set ceases to exist. For SMS-managed GDSs, SCRATCH also specifies that the NVR is to be removed from the VVDS when the data set is uncatalogued.

LIMIT is a required parameter and specifies the maximum number of GDGs in the group. In the example on the previous visuals, this value is 255.

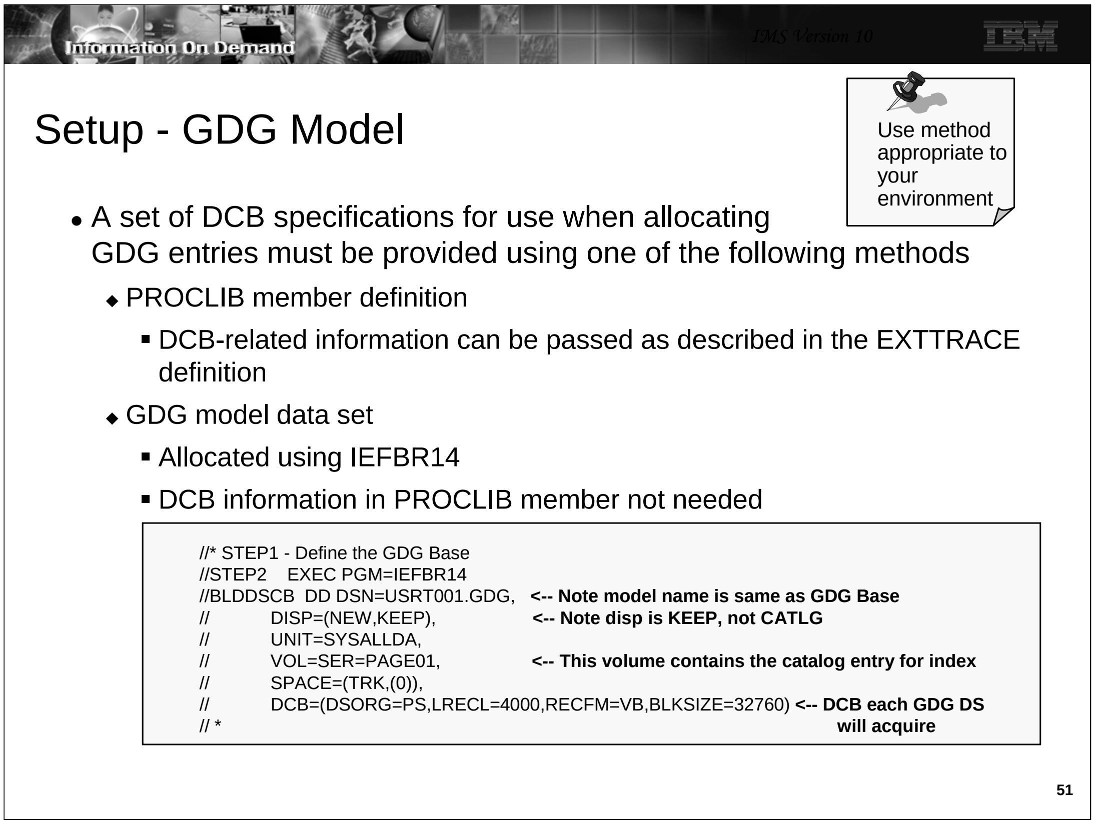

Information On Demand
IMS Version 10
IBM

Setup - GDG Model

Use method appropriate to your environment

- A set of DCB specifications for use when allocating GDG entries must be provided using one of the following methods
  - PROCLIB member definition
    - DCB-related information can be passed as described in the EXTTRACE definition
  - GDG model data set
    - Allocated using IEFBR14
    - DCB information in PROCLIB member not needed

```
//* STEP1 - Define the GDG Base
//STEP2    EXEC PGM=IEFBR14
//BLDDSCB  DD DSN=USRT001.GDG,   <-- Note model name is same as GDG Base
//           DISP=(NEW,KEEP),    <-- Note disp is KEEP, not CATLG
//           UNIT=SYSALLDA,
//           VOL=SER=PAGE01,     <-- This volume contains the catalog entry for index
//           SPACE=(TRK,(0)),
//           DCB=(DSORG=PS,LRECL=4000,RECFM=VB,BLKSIZE=32760) <-- DCB each GDG DS
// *                                                              will acquire
```

51

As with all GDGs, DCB information must be provided for use when the GDG entries are created. The DCB can be provided via one of the following methods:

- PROCLIB member definition - the DCB-related information can be passed in the EXTTRACE definition
- GDG model data set - a GDG model data set can be allocated using IEFBR14 and put on the same volume as the user catalog and GDG data sets. With a GDG model, providing the DCB information in EXTTRACE is not required. The visual gives an example of the IEFBR14 JCL.

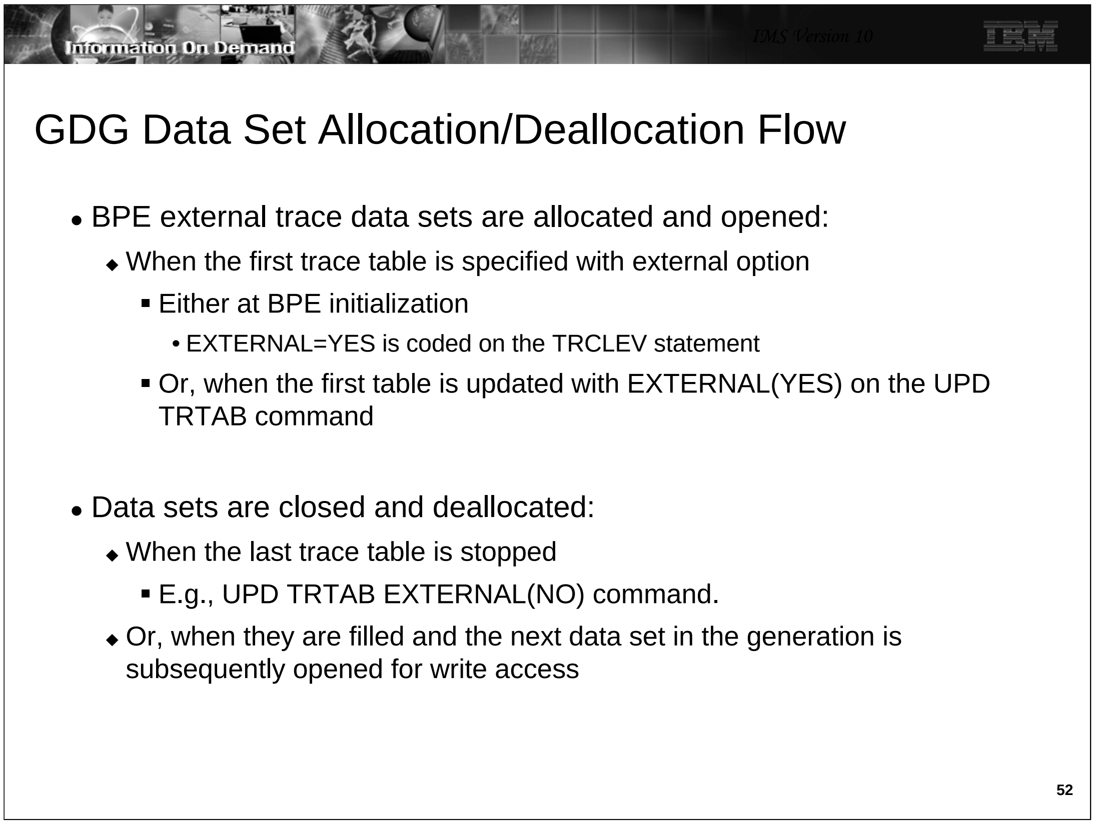

Information On Demand
IMS Version 10
IBM

GDG Data Set Allocation/Deallocation Flow

- BPE external trace data sets are allocated and opened:
  - When the first trace table is specified with external option
    - Either at BPE initialization
      - EXTERNAL=YES is coded on the TRCLEV statement
    - Or, when the first table is updated with EXTERNAL(YES) on the UPD TRTAB command
- Data sets are closed and deallocated:
  - When the last trace table is stopped
    - E.g., UPD TRTAB EXTERNAL(NO) command.
  - Or, when they are filled and the next data set in the generation is subsequently opened for write access

52

BPE external trace data sets are allocated and opened in one of two situations: during BPE initialization if the TRCLEV statement specifies EXTERNAL=YES, or as a result of the UPD TRTAB command and a specification of EXTERNAL(YES).

The trace data sets are closed and deallocated when the last trace table is stopped by a specification of EXTERNAL(NO) in the UPD TRTAB EXTERNAL command, or when the data sets are filled and the next data set in the generation is subsequently opened for write access.

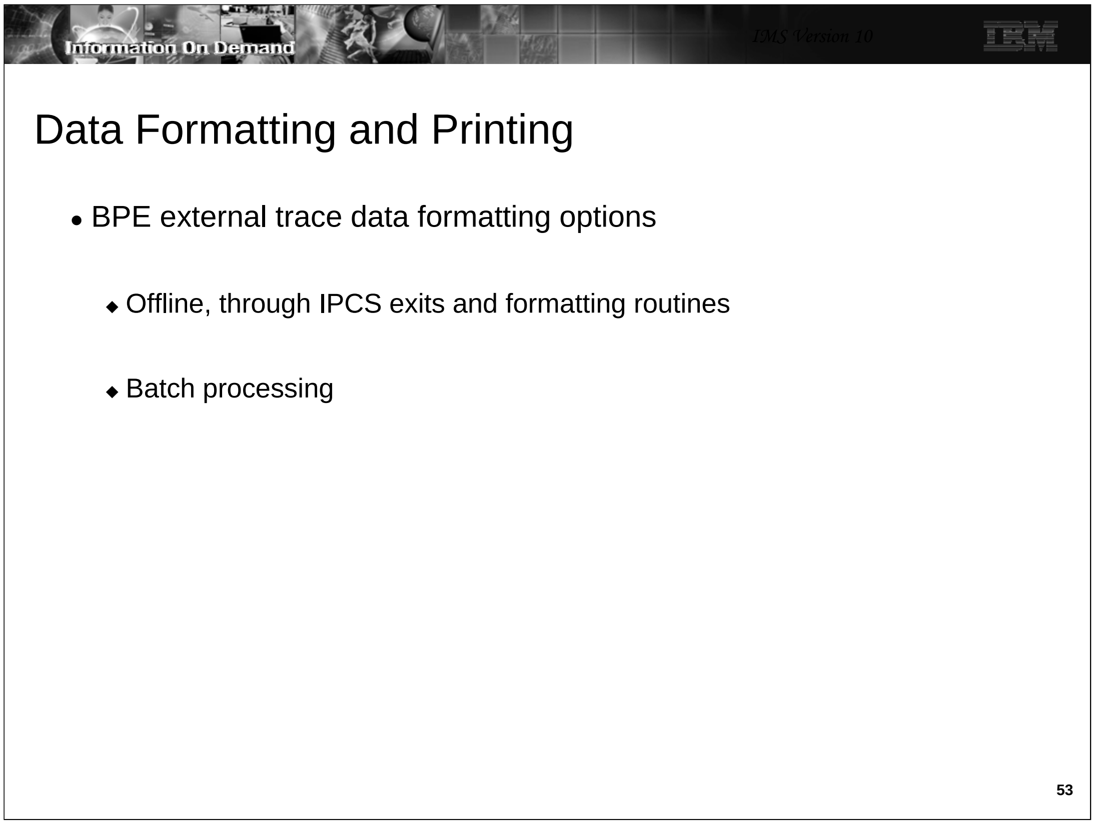

Information On Demand
IMS Version 10
IBM

Data Formatting and Printing

- BPE external trace data formatting options
  - Offline, through IPCS exits and formatting routines
  - Batch processing

53

There are two modes available to format and print the external trace data. The first is through IPCS routines. The second is through batch processing.

Data Formatting and Printing - IPCS

- IPCS
  - From the IPCS main menu, select option 0 to specify the GDG used
  - The selection brings up the IPCS Default Values panel, hit enter

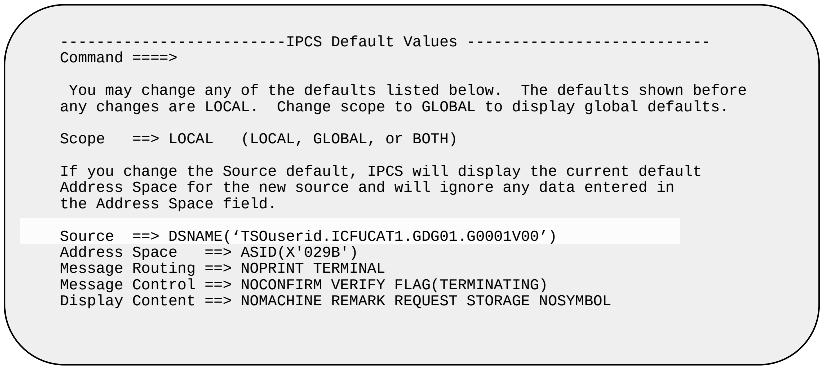

```
-------------------------IPCS Default Values ---------------------------
Command ====>

 You may change any of the defaults listed below.  The defaults shown before
any changes are LOCAL.  Change scope to GLOBAL to display global defaults.

Scope   ==> LOCAL   (LOCAL, GLOBAL, or BOTH)

If you change the Source default, IPCS will display the current default
Address Space for the new source and will ignore any data entered in
the Address Space field.

Source  ==> DSNAME(‘TSOuserid.ICFUCAT1.GDG01.G0001V00’)
Address Space   ==> ASID(X'029B')
Message Routing ==> NOPRINT TERMINAL
Message Control ==> NOCONFIRM VERIFY FLAG(TERMINATING)
Display Content ==> NOMACHINE REMARK REQUEST STORAGE NOSYMBOL
```

This and the next several visuals give an example of using IPCS to format and print the external data set trace records.

Data Formatting and Printing - IPCS …

- The IPCS MVS Dump Component Data Analysis panel shows the available options for analysis
- Select DFSAAMPR - IMS Interactive Dump Formatter

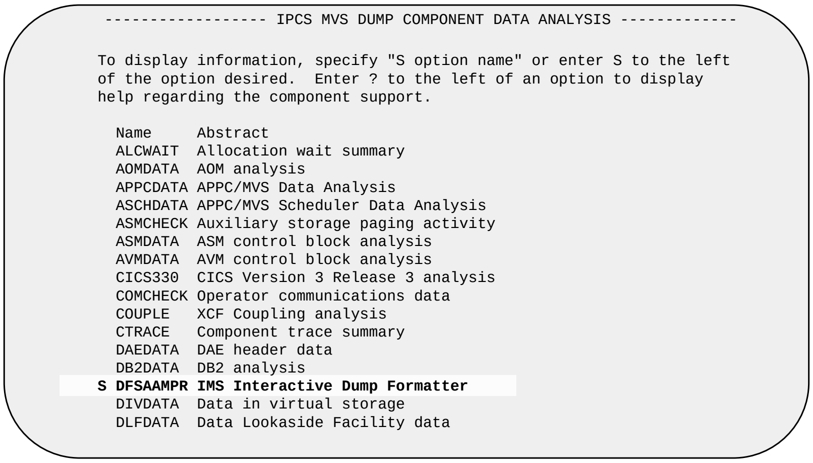

```
 ----------------- IPCS MVS DUMP COMPONENT DATA ANALYSIS -------------

To display information, specify "S option name" or enter S to the left
of the option desired.  Enter ? to the left of an option to display
help regarding the component support.

  Name     Abstract
  ALCWAIT  Allocation wait summary
  AOMDATA  AOM analysis
  APPCDATA APPC/MVS Data Analysis
  ASCHDATA APPC/MVS Scheduler Data Analysis
  ASMCHECK Auxiliary storage paging activity
  ASMDATA  ASM control block analysis
  AVMDATA  AVM control block analysis
  CICS330  CICS Version 3 Release 3 analysis
  COMCHECK Operator communications data
  COUPLE   XCF Coupling analysis
  CTRACE   Component trace summary
  DAEDATA  DAE header data
  DB2DATA  DB2 analysis
S DFSAAMPR IMS Interactive Dump Formatter
  DIVDATA  Data in virtual storage
  DLFDATA  Data Lookaside Facility data
```

The IMS Interactive Dump Formatter is named DFSAAMPR.

Data Formatting and Printing - IPCS …

- Select Option 6

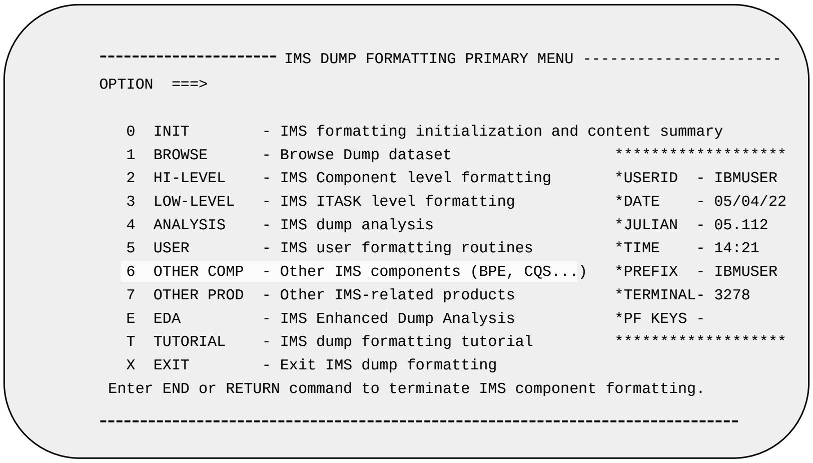

```
---------------------- IMS DUMP FORMATTING PRIMARY MENU ----------------------
OPTION  ===>

   0  INIT        - IMS formatting initialization and content summary
   1  BROWSE      - Browse Dump dataset               *******************
   2  HI-LEVEL    - IMS Component level formatting    *USERID  - IBMUSER
   3  LOW-LEVEL   - IMS ITASK level formatting        *DATE    - 05/04/22
   4  ANALYSIS    - IMS dump analysis                 *JULIAN  - 05.112
   5  USER        - IMS user formatting routines      *TIME    - 14:21
   6  OTHER COMP  - Other IMS components (BPE, CQS...)   *PREFIX  - IBMUSER
   7  OTHER PROD  - Other IMS-related products        *TERMINAL- 3278
   E  EDA         - IMS Enhanced Dump Analysis        *PF KEYS -
   T  TUTORIAL    - IMS dump formatting tutorial      *******************
   X  EXIT        - Exit IMS dump formatting
 Enter END or RETURN command to terminate IMS component formatting.
-----------------------------------------------------------------------
```

The IMS Dump Formatting Primary Menu shows the available IMS dump formatting options. Select Option 6 - ‘OTHER COMP’.

Data Formatting and Printing - IPCS …

- Select Option B (BPE)

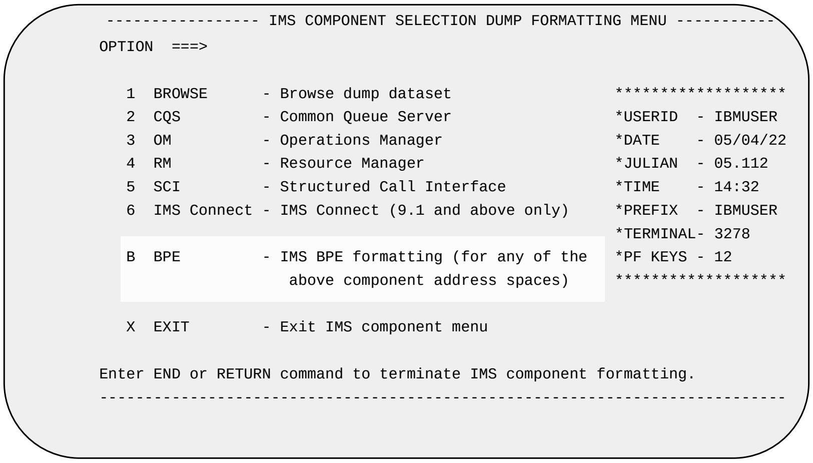

```
 ---------------- IMS COMPONENT SELECTION DUMP FORMATTING MENU -----------
OPTION  ===>

   1  BROWSE      - Browse dump dataset                  *******************
   2  CQS         - Common Queue Server                  *USERID  - IBMUSER
   3  OM          - Operations Manager                   *DATE    - 05/04/22
   4  RM          - Resource Manager                     *JULIAN  - 05.112
   5  SCI         - Structured Call Interface            *TIME    - 14:32
   6  IMS Connect - IMS Connect (9.1 and above only)     *PREFIX  - IBMUSER
                                                         *TERMINAL- 3278
   B  BPE         - IMS BPE formatting (for any of the   *PF KEYS - 12
                     above component address spaces)     *******************

   X  EXIT        - Exit IMS component menu

Enter END or RETURN command to terminate IMS component formatting.
---------------------------------------------------------------------------
```

The next panel identifies IMS BPE Formatting. Select B to access this formatter. Alternatively, each of the components can be chosen separately.

Data Formatting and Printing - IPCS …

- Select Option 4 for the BPE external trace formatting

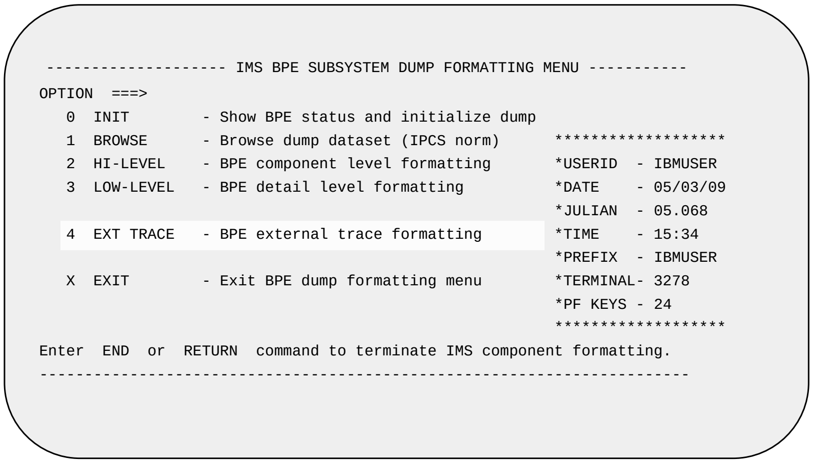

```
 -------------------- IMS BPE SUBSYSTEM DUMP FORMATTING MENU -----------
OPTION  ===>
   0  INIT        - Show BPE status and initialize dump
   1  BROWSE      - Browse dump dataset (IPCS norm)      *******************
   2  HI-LEVEL    - BPE component level formatting       *USERID  - IBMUSER
   3  LOW-LEVEL   - BPE detail level formatting          *DATE    - 05/03/09
                                                         *JULIAN  - 05.068
   4  EXT TRACE   - BPE external trace formatting        *TIME    - 15:34
                                                         *PREFIX  - IBMUSER
   X  EXIT        - Exit BPE dump formatting menu        *TERMINAL- 3278
                                                         *PF KEYS - 24
                                                         *******************
Enter  END  or  RETURN  command to terminate IMS component formatting.
------------------------------------------------------------------------
```

Option 4 provides the access to the new BPE External Trace formatting capability.

Data Formatting and Printing - IPCS …

- Select the traces to view

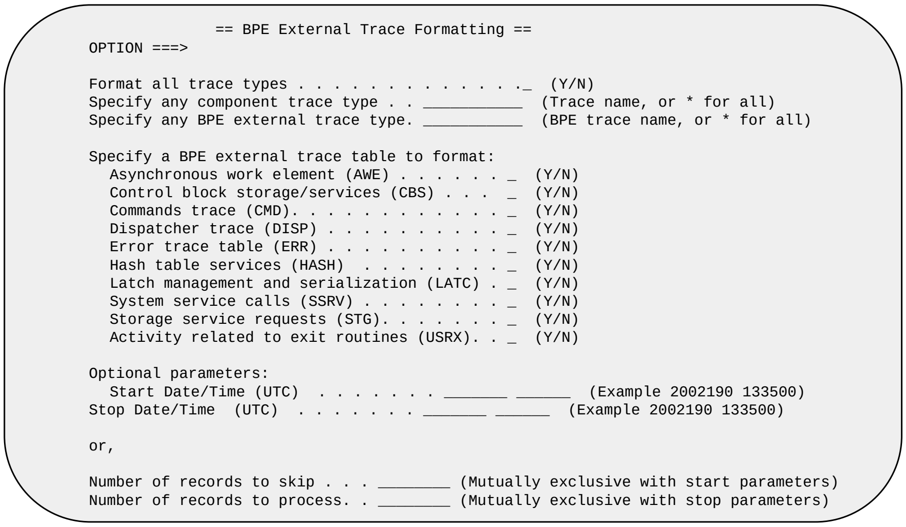

```
              == BPE External Trace Formatting ==
OPTION ===>

Format all trace types . . . . . . . . . . . . . ._  (Y/N)
Specify any component trace type . . ___________  (Trace name, or * for all)
Specify any BPE external trace type. ___________  (BPE trace name, or * for all)

Specify a BPE external trace table to format:
  Asynchronous work element (AWE) . . . . . . . _  (Y/N)
  Control block storage/services (CBS) . . .  _  (Y/N)
  Commands trace (CMD). . . . . . . . . . . . . _  (Y/N)
  Dispatcher trace (DISP) . . . . . . . . . . . _  (Y/N)
  Error trace table (ERR) . . . . . . . . . . . _  (Y/N)
  Hash table services (HASH)  . . . . . . . . . _  (Y/N)
  Latch management and serialization (LATC) . _  (Y/N)
  System service calls (SSRV) . . . . . . . . . _  (Y/N)
  Storage service requests (STG). . . . . . . . _  (Y/N)
  Activity related to exit routines (USRX). . _  (Y/N)

Optional parameters:
  Start Date/Time (UTC)  . . . . . . . _______ ______  (Example 2002190 133500)
Stop Date/Time  (UTC)  . . . . . . . _______ ______  (Example 2002190 133500)

or,

Number of records to skip . . . ________ (Mutually exclusive with start parameters)
Number of records to process. . ________ (Mutually exclusive with stop parameters)
```

The BPE External Trace Formatting panel provides the options to select how the trace is to be produced.

Data Formatting and Printing - IPCS …

- Each of the components can also be chosen separately
  - For example, Option 6 will request IMS Connect formatting

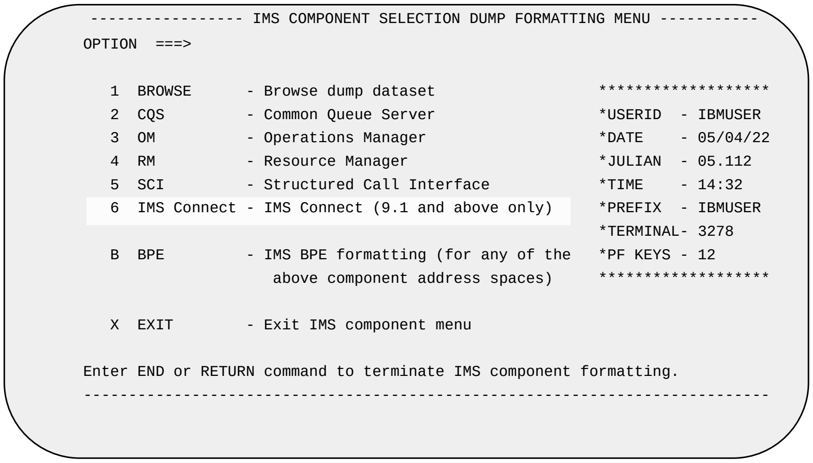

```
 ----------------- IMS COMPONENT SELECTION DUMP FORMATTING MENU -----------
OPTION  ===>

   1  BROWSE      - Browse dump dataset                  *******************
   2  CQS         - Common Queue Server                  *USERID  - IBMUSER
   3  OM          - Operations Manager                   *DATE    - 05/04/22
   4  RM          - Resource Manager                     *JULIAN  - 05.112
   5  SCI         - Structured Call Interface            *TIME    - 14:32
   6  IMS Connect - IMS Connect (9.1 and above only)     *PREFIX  - IBMUSER
                                                         *TERMINAL- 3278
   B  BPE         - IMS BPE formatting (for any of the   *PF KEYS - 12
                     above component address spaces)     *******************

   X  EXIT        - Exit IMS component menu

Enter END or RETURN command to terminate IMS component formatting.
-----------------------------------------------------------------------------
```

IPCS also supports accessing each component separately. This next example will show the panels associated with IMS Connect.

Data Formatting and Printing - IPCS …

- Select Option 4 for External Trace formatting

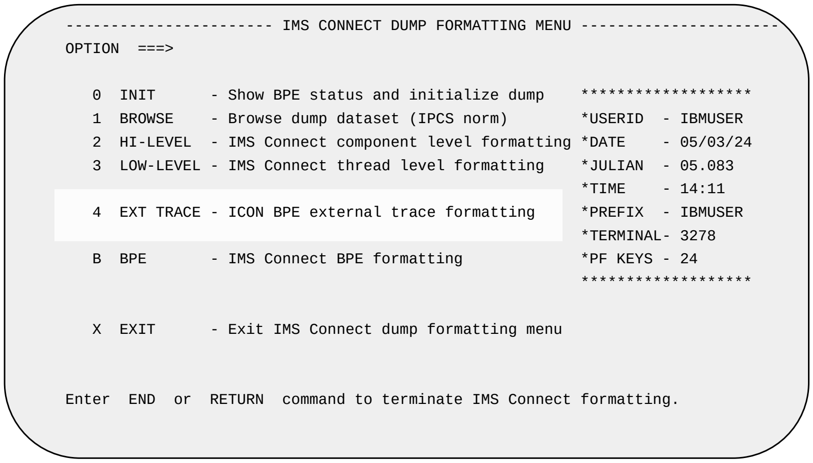

```
----------------------- IMS CONNECT DUMP FORMATTING MENU ----------------------
OPTION  ===>

   0  INIT      - Show BPE status and initialize dump    *******************
   1  BROWSE    - Browse dump dataset (IPCS norm)        *USERID  - IBMUSER
   2  HI-LEVEL  - IMS Connect component level formatting *DATE    - 05/03/24
   3  LOW-LEVEL - IMS Connect thread level formatting    *JULIAN  - 05.083
                                                         *TIME    - 14:11
   4  EXT TRACE - ICON BPE external trace formatting     *PREFIX  - IBMUSER
                                                         *TERMINAL- 3278
   B  BPE       - IMS Connect BPE formatting             *PF KEYS - 24
                                                         *******************


   X  EXIT      - Exit IMS Connect dump formatting menu



Enter  END  or  RETURN  command to terminate IMS Connect formatting.
```

The IMS Connect Dump Formatting panel provides several options. Option 4 is new and shows how to access the BPE external trace formatting for IMS Connect.

Data Formatting and Printing - IPCS …

- New panel

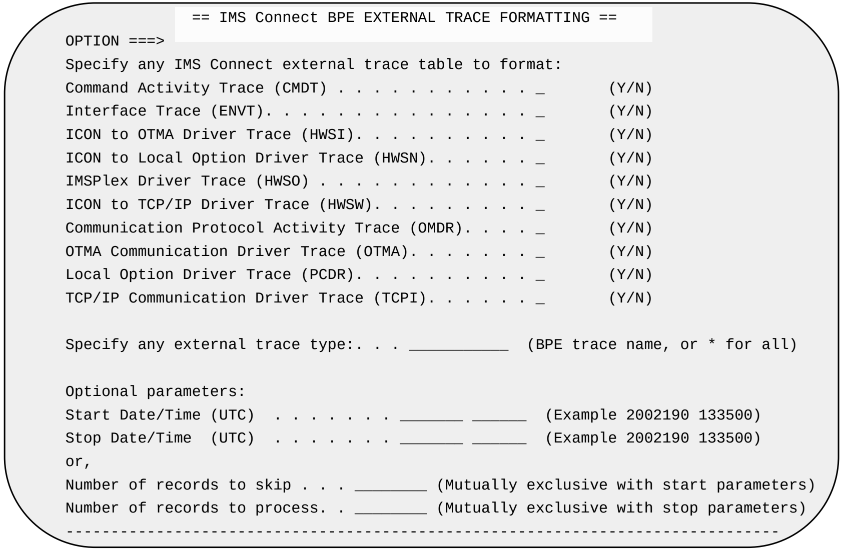

```
                == IMS Connect BPE EXTERNAL TRACE FORMATTING ==
OPTION ===>
Specify any IMS Connect external trace table to format:
Command Activity Trace (CMDT) . . . . . . . . . . . _       (Y/N)
Interface Trace (ENVT). . . . . . . . . . . . . . . _       (Y/N)
ICON to OTMA Driver Trace (HWSI). . . . . . . . . . _       (Y/N)
ICON to Local Option Driver Trace (HWSN). . . . . . _       (Y/N)
IMSPlex Driver Trace (HWSO) . . . . . . . . . . . . _       (Y/N)
ICON to TCP/IP Driver Trace (HWSW). . . . . . . . . _       (Y/N)
Communication Protocol Activity Trace (OMDR). . . . _       (Y/N)
OTMA Communication Driver Trace (OTMA). . . . . . . _       (Y/N)
Local Option Driver Trace (PCDR). . . . . . . . . . _       (Y/N)
TCP/IP Communication Driver Trace (TCPI). . . . . . _       (Y/N)

Specify any external trace type:. . . ___________  (BPE trace name, or * for all)

Optional parameters:
Start Date/Time (UTC)  . . . . . . . _______ ______  (Example 2002190 133500)
Stop Date/Time  (UTC)  . . . . . . . _______ ______  (Example 2002190 133500)
or,
Number of records to skip . . . ________ (Mutually exclusive with start parameters)
Number of records to process. . ________ (Mutually exclusive with stop parameters)
-------------------------------------------------------------------------------
```

The visual shows the new panel for IMS Connect. Note that the other components, e.g., CQS, RM, OM and SCI also have panels that relate specifically to each of those functions.

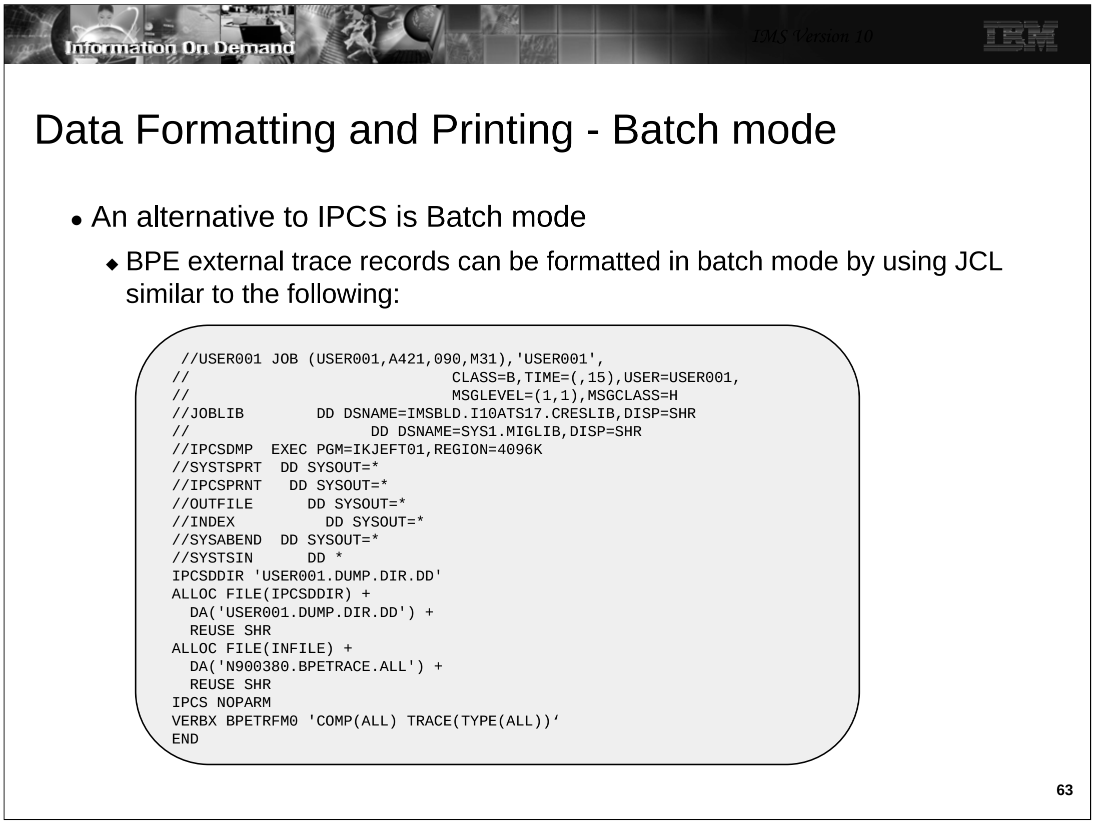

Data Formatting and Printing - Batch mode

- An alternative to IPCS is Batch mode
  - BPE external trace records can be formatted in batch mode by using JCL similar to the following:

```
 //USER001 JOB (USER001,A421,090,M31),'USER001',
//                             CLASS=B,TIME=(,15),USER=USER001,
//                             MSGLEVEL=(1,1),MSGCLASS=H
//JOBLIB        DD DSNAME=IMSBLD.I10ATS17.CRESLIB,DISP=SHR
//                    DD DSNAME=SYS1.MIGLIB,DISP=SHR
//IPCSDMP  EXEC PGM=IKJEFT01,REGION=4096K
//SYSTSPRT  DD SYSOUT=*
//IPCSPRNT   DD SYSOUT=*
//OUTFILE      DD SYSOUT=*
//INDEX          DD SYSOUT=*
//SYSABEND  DD SYSOUT=*
//SYSTSIN      DD *
IPCSDDIR 'USER001.DUMP.DIR.DD'
ALLOC FILE(IPCSDDIR) +
  DA('USER001.DUMP.DIR.DD') +
  REUSE SHR
ALLOC FILE(INFILE) +
  DA('N900380.BPETRACE.ALL') +
  REUSE SHR
IPCS NOPARM
VERBX BPETRFM0 'COMP(ALL) TRACE(TYPE(ALL))‘
END
```

Batch invocation is an alternative to IPCS. This visual gives an example of how to run a batch job for data formatting and printing.

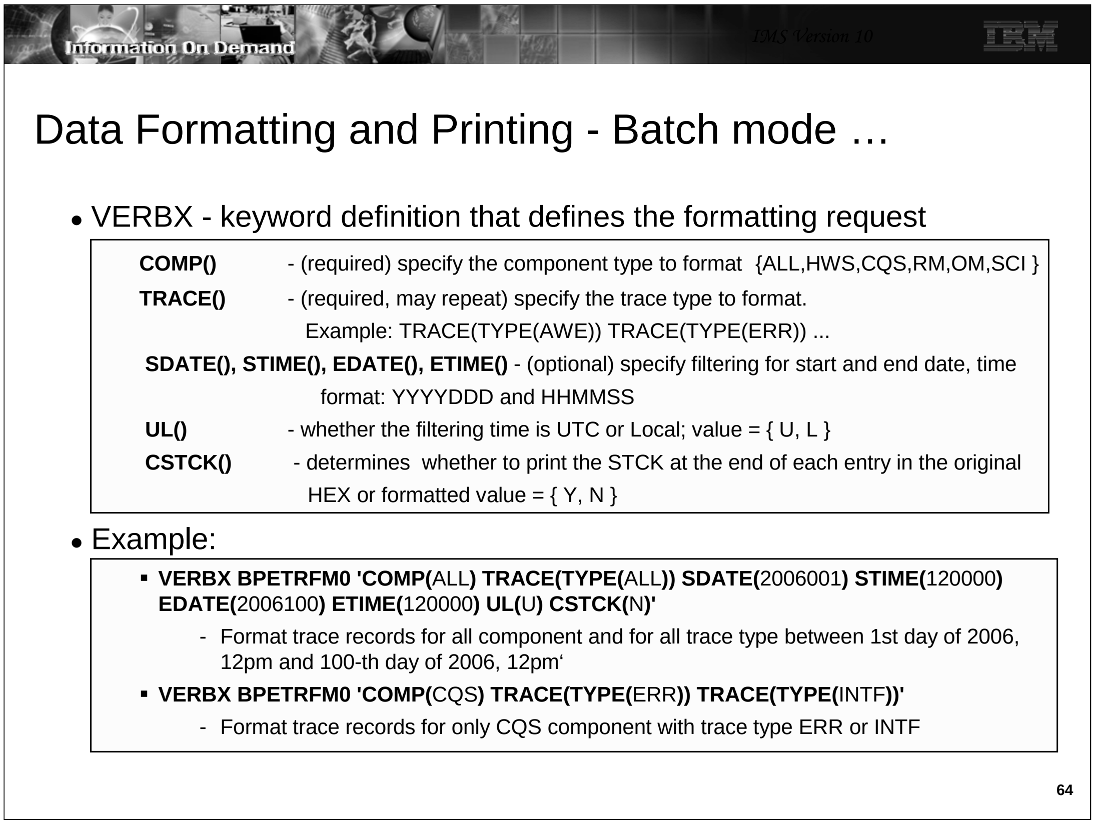

Information On Demand
IMS Version 10

Data Formatting and Printing - Batch mode …

- VERBX - keyword definition that defines the formatting request

COMP() - (required) specify the component type to format {ALL,HWS,CQS,RM,OM,SCI }

TRACE() - (required, may repeat) specify the trace type to format.
Example: TRACE(TYPE(AWE)) TRACE(TYPE(ERR)) ...

SDATE(), STIME(), EDATE(), ETIME() - (optional) specify filtering for start and end date, time format: YYYYDDD and HHMMSS

UL() - whether the filtering time is UTC or Local; value = { U, L }

CSTCK() - determines whether to print the STCK at the end of each entry in the original HEX or formatted value = { Y, N }

- Example:
  - **VERBX BPETRFM0 'COMP(**ALL**) TRACE(TYPE(**ALL**)) SDATE(**2006001**) STIME(**120000**) EDATE(**2006100**) ETIME(**120000**) UL(**U**) CSTCK(**N**)'**
    - Format trace records for all component and for all trace type between 1st day of 2006, 12pm and 100-th day of 2006, 12pm‘
  - **VERBX BPETRFM0 'COMP(**CQS**) TRACE(TYPE(**ERR**)) TRACE(TYPE(**INTF**))'**
    - Format trace records for only CQS component with trace type ERR or INTF

The VERBX keyword provides the definitions that request the type of formatting to be done along with filtering requirements.

Formatted Data - Example

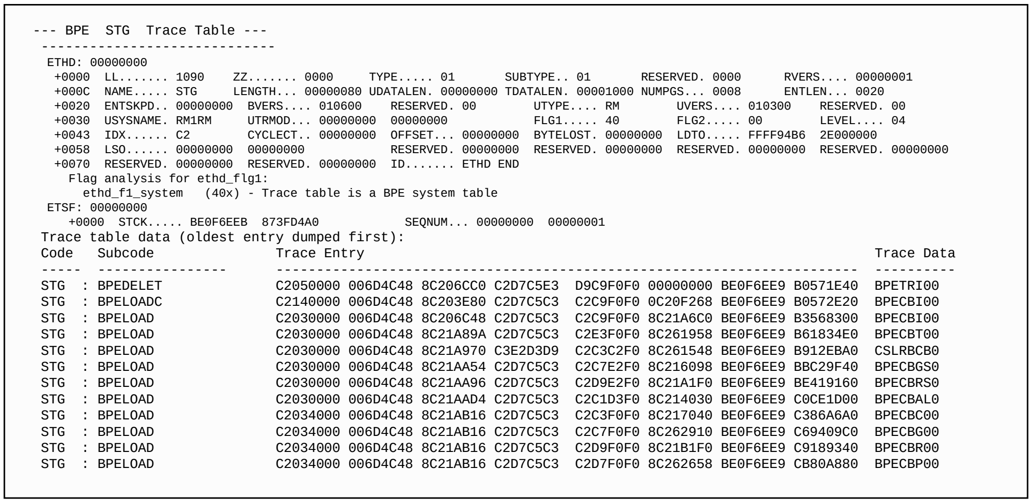

```
 --- BPE  STG  Trace Table ---
  -----------------------------
   ETHD: 00000000
    +0000  LL....... 1090    ZZ....... 0000     TYPE..... 01       SUBTYPE.. 01       RESERVED. 0000      RVERS.... 00000001
    +000C  NAME..... STG     LENGTH... 00000080 UDATALEN. 00000000 TDATALEN. 00001000 NUMPGS... 0008      ENTLEN... 0020
    +0020  ENTSKPD.. 00000000  BVERS.... 010600    RESERVED. 00        UTYPE.... RM        UVERS.... 010300    RESERVED. 00
    +0030  USYSNAME. RM1RM     UTRMOD... 00000000  00000000            FLG1..... 40        FLG2..... 00        LEVEL.... 04
    +0043  IDX...... C2        CYCLECT.. 00000000  OFFSET... 00000000  BYTELOST. 00000000  LDTO..... FFFF94B6  2E000000
    +0058  LSO...... 00000000  00000000            RESERVED. 00000000  RESERVED. 00000000  RESERVED. 00000000  RESERVED. 00000000
    +0070  RESERVED. 00000000  RESERVED. 00000000  ID....... ETHD END
      Flag analysis for ethd_flg1:
        ethd_f1_system   (40x) - Trace table is a BPE system table
   ETSF: 00000000
      +0000  STCK..... BE0F6EEB  873FD4A0            SEQNUM... 00000000  00000001
  Trace table data (oldest entry dumped first):
  Code   Subcode               Trace Entry                                                               Trace Data
  -----  ----------------      ----------------------------------------------------------------------    ----------
  STG  : BPEDELET              C2050000 006D4C48 8C206CC0 C2D7C5E3  D9C9F0F0 00000000 BE0F6EE9 B0571E40  BPETRI00
  STG  : BPELOADC              C2140000 006D4C48 8C203E80 C2D7C5C3  C2C9F0F0 0C20F268 BE0F6EE9 B0572E20  BPECBI00
  STG  : BPELOAD               C2030000 006D4C48 8C206C48 C2D7C5C3  C2C9F0F0 8C21A6C0 BE0F6EE9 B3568300  BPECBI00
  STG  : BPELOAD               C2030000 006D4C48 8C21A89A C2D7C5C3  C2E3F0F0 8C261958 BE0F6EE9 B61834E0  BPECBT00
  STG  : BPELOAD               C2030000 006D4C48 8C21A970 C3E2D3D9  C2C3C2F0 8C261548 BE0F6EE9 B912EBA0  CSLRBCB0
  STG  : BPELOAD               C2030000 006D4C48 8C21AA54 C2D7C5C3  C2C7E2F0 8C216098 BE0F6EE9 BBC29F40  BPECBGS0
  STG  : BPELOAD               C2030000 006D4C48 8C21AA96 C2D7C5C3  C2D9E2F0 8C21A1F0 BE0F6EE9 BE419160  BPECBRS0
  STG  : BPELOAD               C2030000 006D4C48 8C21AAD4 C2D7C5C3  C2C1D3F0 8C214030 BE0F6EE9 C0CE1D00  BPECBAL0
  STG  : BPELOAD               C2034000 006D4C48 8C21AB16 C2D7C5C3  C2C3F0F0 8C217040 BE0F6EE9 C386A6A0  BPECBC00
  STG  : BPELOAD               C2034000 006D4C48 8C21AB16 C2D7C5C3  C2C7F0F0 8C262910 BE0F6EE9 C69409C0  BPECBG00
  STG  : BPELOAD               C2034000 006D4C48 8C21AB16 C2D7C5C3  C2D9F0F0 8C21B1F0 BE0F6EE9 C9189340  BPECBR00
  STG  : BPELOAD               C2034000 006D4C48 8C21AB16 C2D7C5C3  C2D7F0F0 8C262658 BE0F6EE9 CB80A880  BPECBP00
```

The formatted data is printed in a format similar to what is shown on this visual.

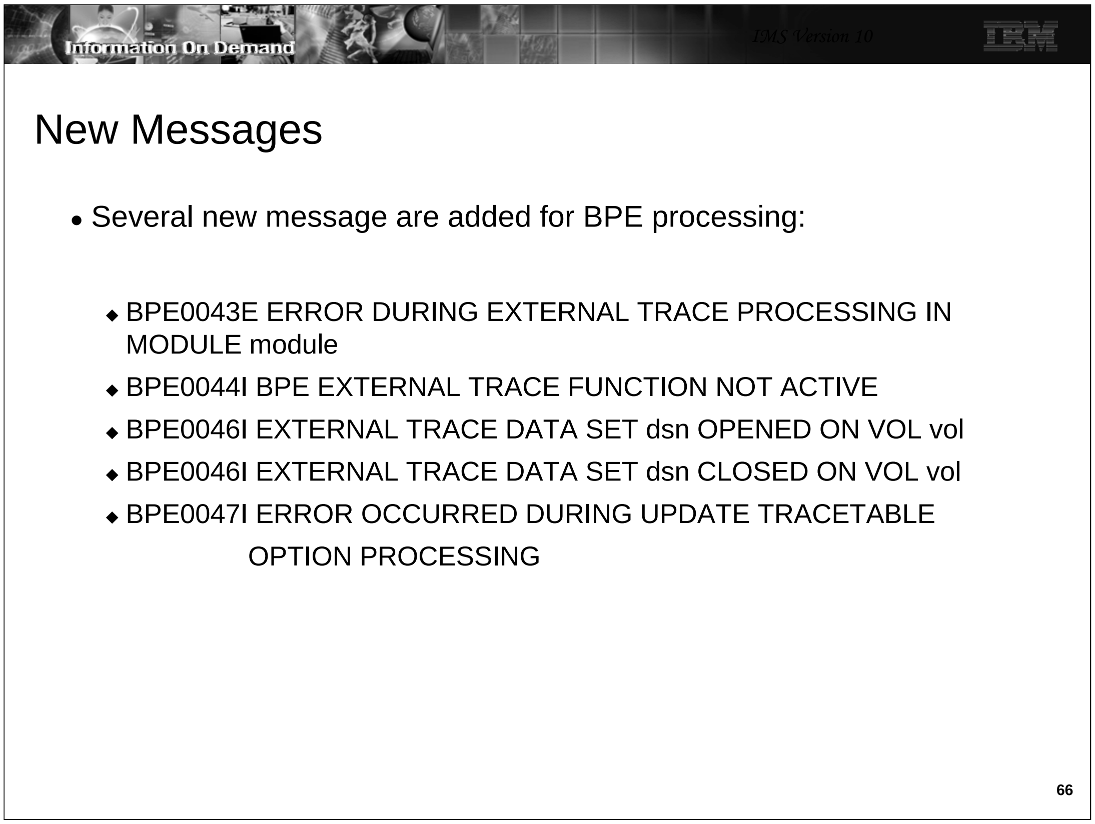

New Messages

- Several new message are added for BPE processing:
  - BPE0043E ERROR DURING EXTERNAL TRACE PROCESSING IN MODULE module
  - BPE0044I BPE EXTERNAL TRACE FUNCTION NOT ACTIVE
  - BPE0046I EXTERNAL TRACE DATA SET dsn OPENED ON VOL vol
  - BPE0046I EXTERNAL TRACE DATA SET dsn CLOSED ON VOL vol
  - BPE0047I ERROR OCCURRED DURING UPDATE TRACETABLE OPTION PROCESSING

If the BPE external trace function is enabled, an installation's Operations Procedures documentation should be updated to reflect the possibility of several new error messages associated with external tracing. This visual documents the new messages.

If an error occurs during UPDATE TRACETABLE processing, the text of the message produces one of the following:

•UNABLE TO GET NECESSARY STORAGE

•ERROR READING BPE CONFIGURATION PROCLIB MEMBER member address space

•ERROR PARSING BPE CONFIGURATION PROCLIB MEMBER member processing

•UNKNOWN ERROR rc FROM PARSING MODULE BPEPCF10

•LOAD FAILED FOR BPEPCF10, BPELOADC RC=rc

•BPE CONFIGURATION PROCLIB MEMBER NAME WAS NOT CODED ON STARTUP PARMS

•PREVIOUS UPDATE COMMAND REREAD PROCESSING IN PROGRESS

•NO MATCHING EXTTRACE STATEMENT FOUND FOR ADDRESS SPACE

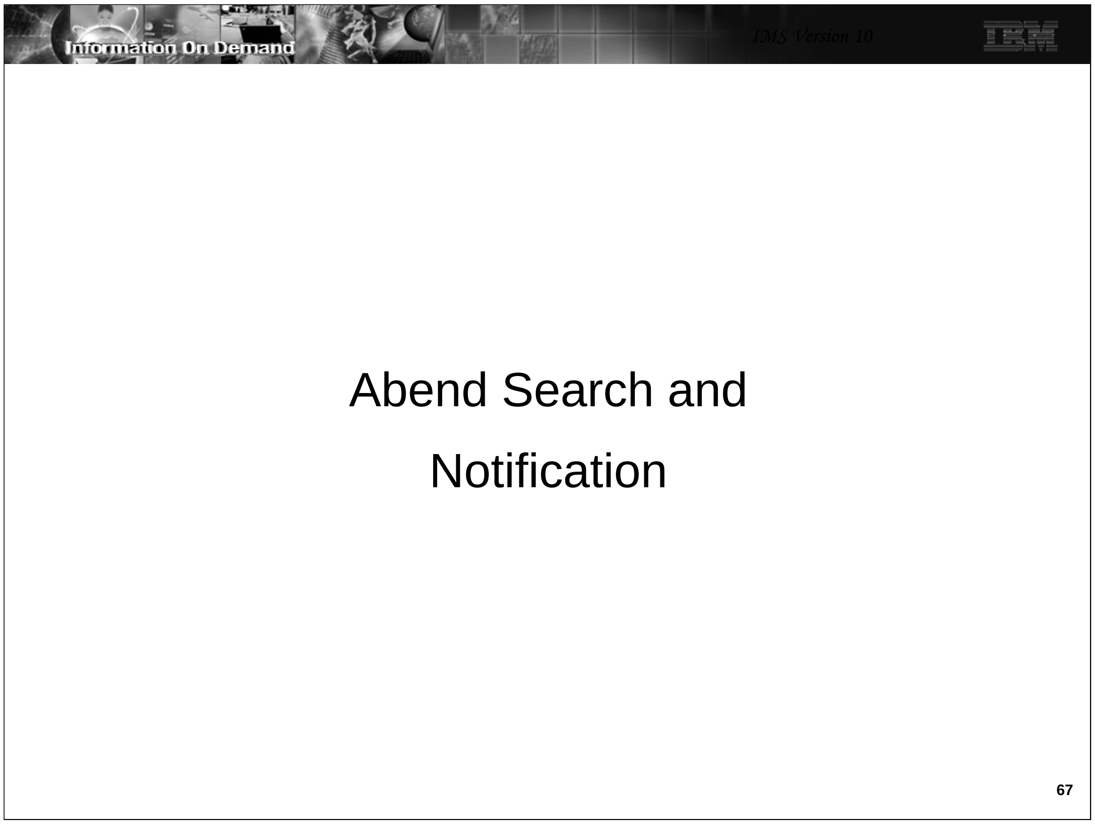

Abend Search and Notification

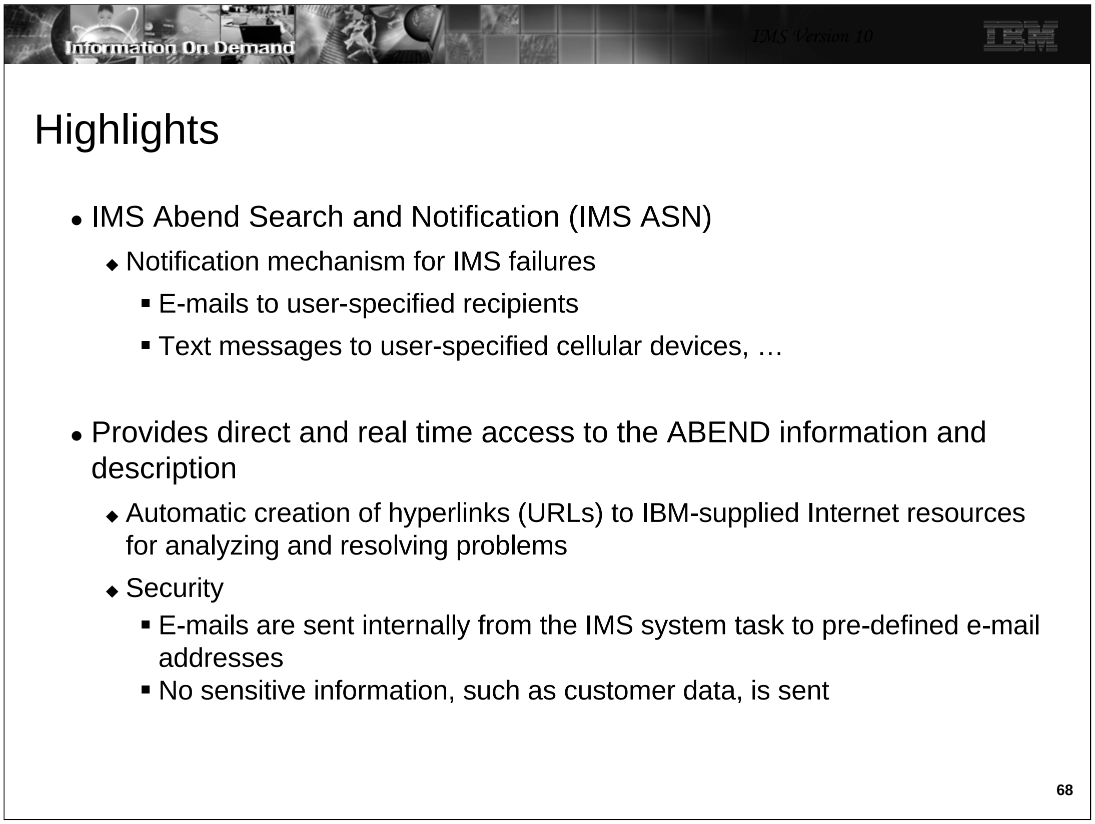

Highlights

- IMS Abend Search and Notification (IMS ASN)
  - ◆ Notification mechanism for IMS failures
    - ▪ E-mails to user-specified recipients
    - ▪ Text messages to user-specified cellular devices, …
- Provides direct and real time access to the ABEND information and description
  - ◆ Automatic creation of hyperlinks (URLs) to IBM-supplied Internet resources for analyzing and resolving problems
  - ◆ Security
    - ▪ E-mails are sent internally from the IMS system task to pre-defined e-mail addresses
    - ▪ No sensitive information, such as customer data, is sent

The IMS Abend Search and Notification (IMS ASN) function in IMS V10 provides a mechanism for 'real-time' automatic notification of IMS problems, such as abnormal terminations, to designated personnel who may not have immediate or direct access to the operator consoles.

When the abend occurs (or when the system fails), IMS ASN gathers information about the abend, the return code, the module affected, and, when possible, the APAR level.

Note that security is not an issue because all associated e-mails are sent internally from the IMS system task to the list of pre-defined e-mail addresses. No sensitive information, such as customer data, is ever sent.

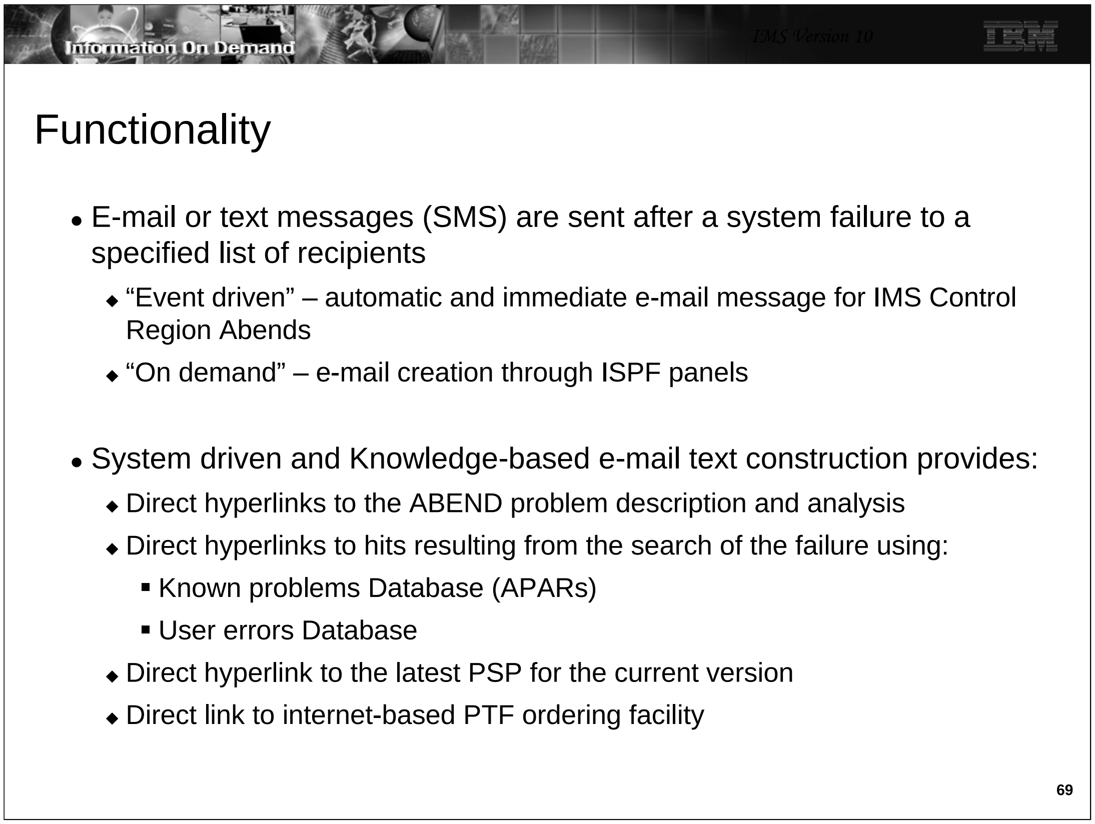
Information On Demand
IMS Version 10

Functionality

- E-mail or text messages (SMS) are sent after a system failure to a specified list of recipients
  - “Event driven” – automatic and immediate e-mail message for IMS Control Region Abends
  - “On demand” – e-mail creation through ISPF panels
- System driven and Knowledge-based e-mail text construction provides:
  - Direct hyperlinks to the ABEND problem description and analysis
  - Direct hyperlinks to hits resulting from the search of the failure using:
    - Known problems Database (APARs)
    - User errors Database
  - Direct hyperlink to the latest PSP for the current version
  - Direct link to internet-based PTF ordering facility

69

The goal of IMS ASN is to notify one or more pre-defined recipients, via mechanisms such as e-mail or a text message to a cellular device, that an abend just occurred in a given subsystem. There are two ways to initiate an IMS ASN e-mail. The first is “event driven” which occurs based upon an event such as an IMS sytem abend. The second is “on demand” by initiating a request that is sent from the ISPF panels that are a part of the IMS ASN support.

IMS ASN also attempts to construct a message that provides more information on the specific nature of the problem along with potential solutions. The function automatically attempts to create hyperlinks to:

•An abend explanation in the IMS Messages and Codes manual

•Pertinent INTERNET database search engines (in practical terms, the hyperlink will be a URL with a pre-built search argument). The following are examples of links but note that, since they are links, then they are subject to change:

Technotes: http://www-1.ibm.com/support/search.wss?rs=81&tc=SSEPH2&dc=DB520+D800+D900+DA900+

DA800&dtm

PSP bucket: https://techsupport.services.ibm.com/server/390.psp390
Messages: http://publibfp.boulder.ibm.com/cgi-bin/bookmgr/BOOKS/dfsmc1g1/CCONTENTS
PMR: https://www-925.ibm.com/software/support/esr/esr_home.do

List of possible related APARs (fix), or APARs in error (PE APAR).
Note: this hyperlink will be present only if the APAR information is available to IMS ASN components at time of error.
The APAR hyperlink is http://www-1.ibm.com/support/search.wss?rs=81&tc=SSEPH2&dc=DB550+D100&dtm

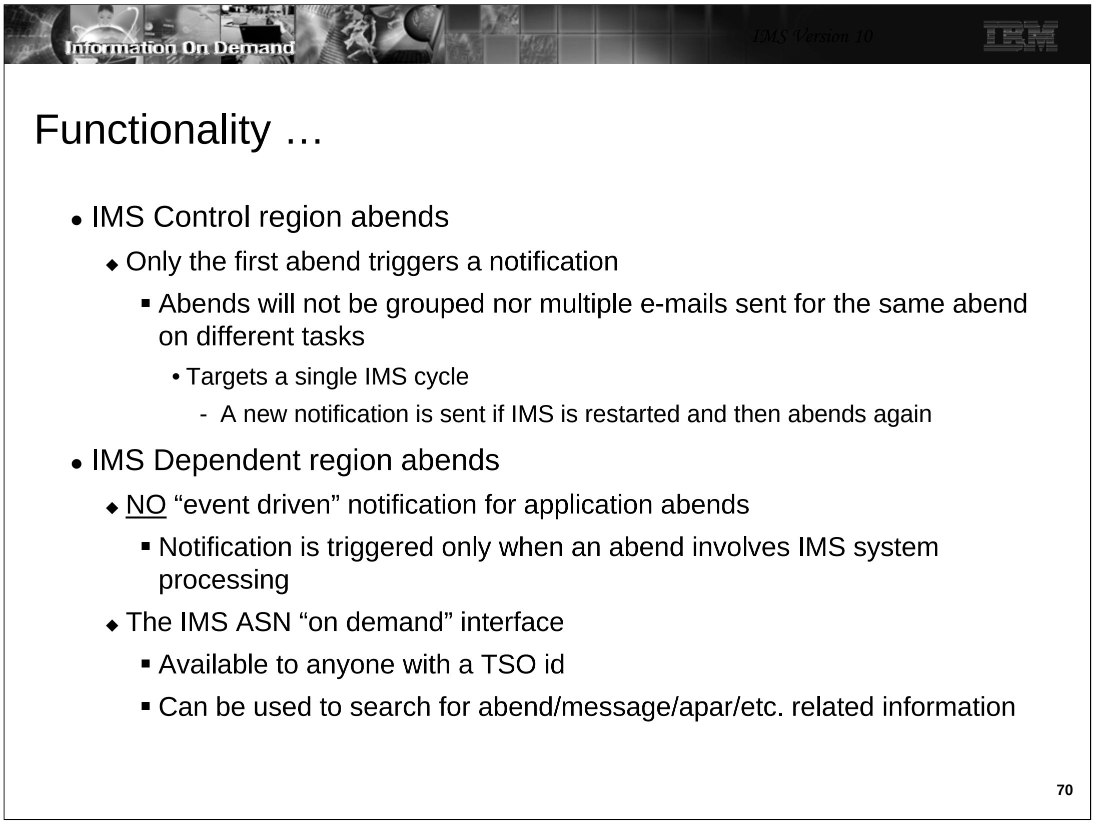
Information On Demand

IMS Version 10

Functionality …

- IMS Control region abends
  - Only the first abend triggers a notification
    - Abends will not be grouped nor multiple e-mails sent for the same abend on different tasks
      - Targets a single IMS cycle
        - A new notification is sent if IMS is restarted and then abends again
- IMS Dependent region abends
  - NO "event driven" notification for application abends
    - Notification is triggered only when an abend involves IMS system processing
  - The IMS ASN "on demand" interface
    - Available to anyone with a TSO id
    - Can be used to search for abend/message/apar/etc. related information

70

Abend notification is triggered when an abend occurs that causes the IMS Control Region to abnormally terminate. Only the first abend triggers a notification. Additionally, IMS ASN does not group abends, nor does it send multiple e-mails for the same abend that might occur in different TCBs. The notification is intended for a single IMS cycle so if IMS is restarted and then abends again, IMS ASN sends a new notification.

Before creating the notification, IMS ASN logic checks a new dump suppression table in the dump module DFSFDMP0 to determine if an abend requires notification. If it does, then IMS ASN gathers information about the abend, the return code, the module affected, and, when possible, the APAR level. On the other hand, there are abends that may have notification interest, but for which no search is meaningful or necessary. The following list shows abends for which either no notification, or notification without hyperlinks is generated:

- U0020 - notification only, no hyperlinks with search arguments provided
- U0150 - no notification necessary. CTL region abend already generated a notification
- U0604 - no notification necessary. /SWITCH SYSTEM FORCE/ACTIVE entered (XRF)
- U4095 - no notification necessary. CTL region abend already generated a notification
- SFFF - non-IBM program termination. No notification necessary
- SFFE - non-IBM program termination. No notification necessary

Note that dependent regions abends (MPP or BMP) trigger the notification and search "event driven" function only if the abend involves IMS system processing. Application abends are not recognized by the "event driven" capability.

Dependent Region abend searches, however, can be accomplished using the IMS ASN functionality. The IMS ASN "on demand interface", available to anyone with a TSO id, can be used to search for abend/message/apar/etc. related information.

Example of E-mail Notification - Graphics

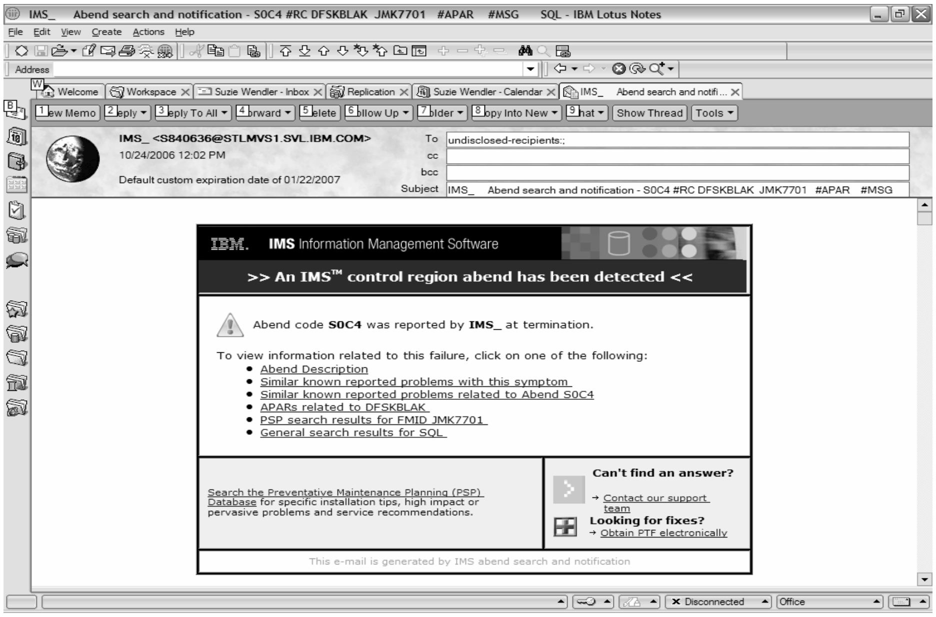

IMS_ <S840636@STLMVS1.SVL.IBM.COM>
10/24/2006 12:02 PM

Default custom expiration date of 01/22/2007

To: undisclosed-recipients:;
cc:
bcc:
Subject: IMS_ Abend search and notification - S0C4 #RC DFSKBLAK JMK7701 #APAR #MSG

IBM. **IMS** Information Management Software

>> An IMS™ control region abend has been detected <<

Abend code **S0C4** was reported by **IMS_** at termination.

To view information related to this failure, click on one of the following:

- Abend Description
- Similar known reported problems with this symptom
- Similar known reported problems related to Abend S0C4
- APARs related to DFSKBLAK
- PSP search results for FMID JMK7701
- General search results for SQL

Search the Preventative Maintenance Planning (PSP) Database for specific installation tips, high impact or pervasive problems and service recommendations.

Can't find an answer?
→ Contact our support team

Looking for fixes?
→ Obtain PTF electronically

This e-mail is generated by IMS abend search and notification

This visual gives an actual example of an e-mail notification in an environment that supports full graphics.

Example of E-mail Notification - Text Only

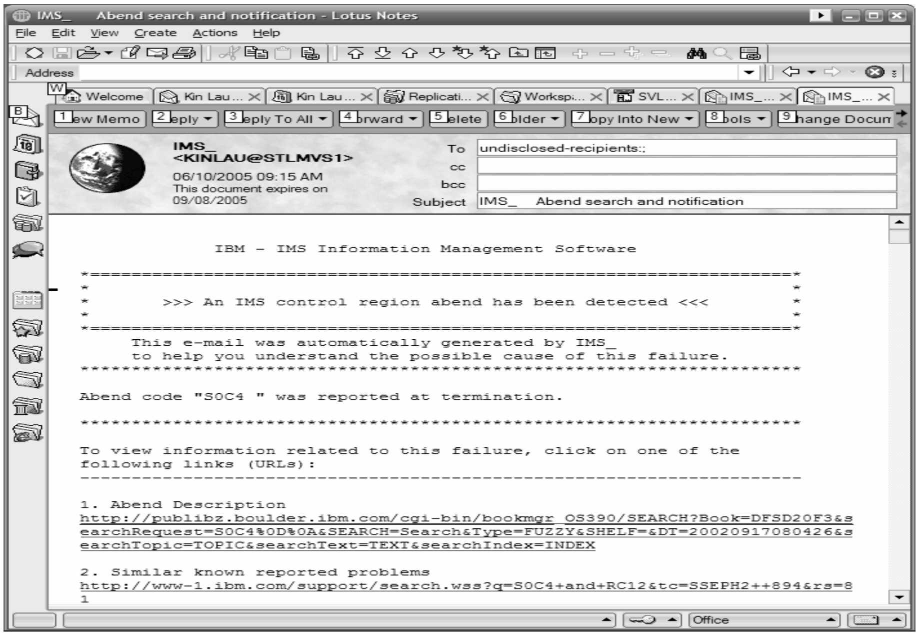

IMS_
<KINLAU@STLMVS1>
06/10/2005 09:15 AM
This document expires on
09/08/2005

To undisclosed-recipients:;
cc
bcc
Subject IMS_ Abend search and notification

```
               IBM - IMS Information Management Software

*===================================================================*
*                                                                   *
*        >>> An IMS control region abend has been detected <<<      *
*                                                                   *
*===================================================================*
     This e-mail was automatically generated by IMS_
     to help you understand the possible cause of this failure.
*********************************************************************

Abend code "S0C4 " was reported at termination.

*********************************************************************

To view information related to this failure, click on one of the
following links (URLs):
---------------------------------------------------------------------

1. Abend Description
http://publibz.boulder.ibm.com/cgi-bin/bookmgr_OS390/SEARCH?Book=DFSD20F3&s
earchRequest=S0C4%0D%0A&SEARCH=Search&Type=FUZZY&SHELF=&DT=20020917080426&s
earchTopic=TOPIC&searchText=TEXT&searchIndex=INDEX

2. Similar known reported problems
http://www-1.ibm.com/support/search.wss?q=S0C4+and+RC12&tc=SSEPH2++894&rs=8
1
```

For environments with limited graphics capabilities, text-only notification is available. This visual provides an example of the text-only e-mail format.

Setup

- Setup overview
  - Allocate two data sets
    - Runtime data set and skeletal data set
  - Invoke "IMS ASN System Setup"
    - Specifies data sets
    - Specifies message recipients
    - Specifies SMTP server
  - Copy procedure produced by IMS ASN System setup to z/OS Proclib
  - Specify procedure name in DFSDFxxx member

```
/* Diagnostic_STATISTICS Section                                      */
/*********************************************************************/
<SECTION=DIAGNOSTICS_STATISTICS>
IASNPROC=procname   /* IMS abend search and notification PROCLIB member */
```

- Detailed information is in the *System Definition Guide*

IMS ASN has a set of ISPF panels that are used to set up the environment. The application can be invoked either by keying in an exec command in the TSO ISPF Option 6 panel or selecting option 9 in the IMS Application Menu program.

On Demand Interface

1

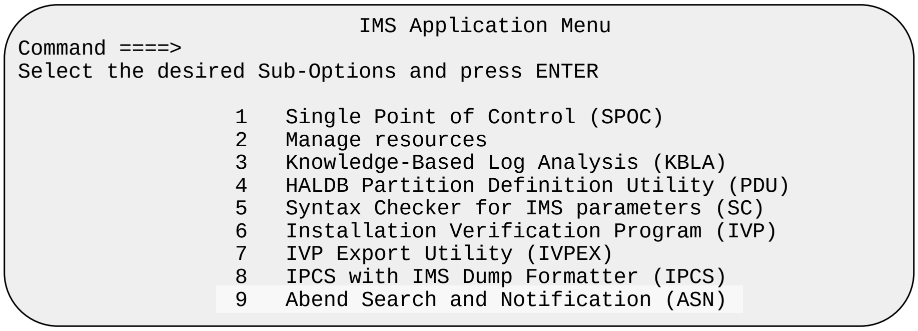

```
                           IMS Application Menu
Command ====>
Select the desired Sub-Options and press ENTER

                 1   Single Point of Control (SPOC)
                 2   Manage resources
                 3   Knowledge-Based Log Analysis (KBLA)
                 4   HALDB Partition Definition Utility (PDU)
                 5   Syntax Checker for IMS parameters (SC)
                 6   Installation Verification Program (IVP)
                 7   IVP Export Utility (IVPEX)
                 8   IPCS with IMS Dump Formatter (IPCS)
                 9   Abend Search and Notification (ASN)
```

2

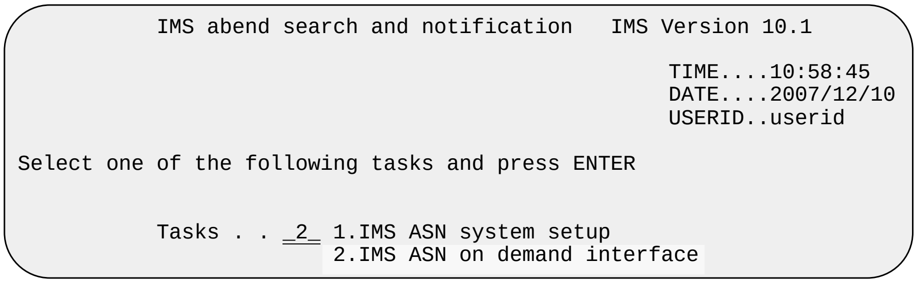

```
           IMS abend search and notification   IMS Version 10.1

                                                    TIME....10:58:45
                                                    DATE....2007/12/10
                                                    USERID..userid

Select one of the following tasks and press ENTER


           Tasks . . _2_ 1.IMS ASN system setup
                         2.IMS ASN on demand interface
```

To invoke the "On Demand" interface, you first go to the IMS Application Menu. ASN is option 9 on this screen. This invokes the IMS abend search and notification screen. The "on demand" interface is option 2 on this screen.

On Demand Interface …

- Specifications for “on demand” processing

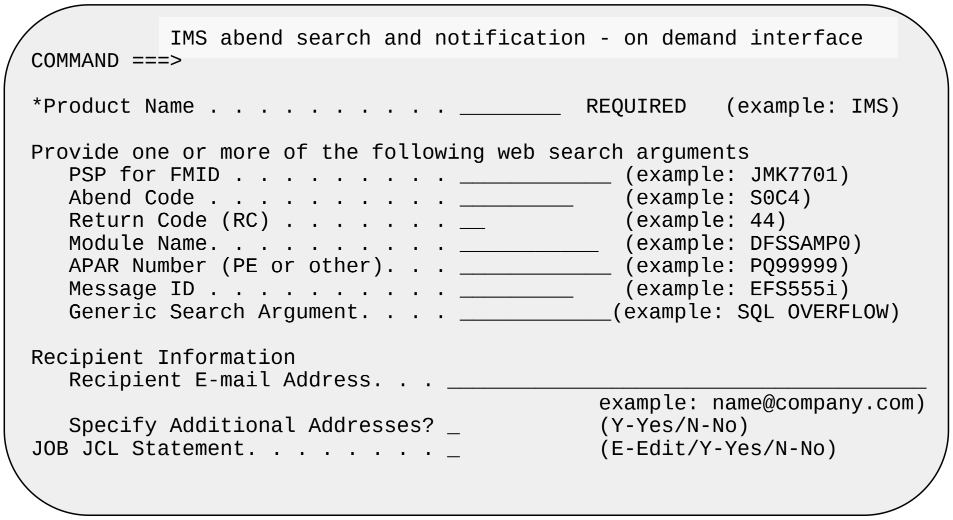

```
                 IMS abend search and notification - on demand interface
COMMAND ===>

*Product Name . . . . . . . . . . ________  REQUIRED   (example: IMS)

Provide one or more of the following web search arguments
   PSP for FMID . . . . . . . . . ____________ (example: JMK7701)
   Abend Code . . . . . . . . . . _________    (example: S0C4)
   Return Code (RC) . . . . . . . __           (example: 44)
   Module Name. . . . . . . . . . ___________  (example: DFSSAMP0)
   APAR Number (PE or other). . . ____________ (example: PQ99999)
   Message ID . . . . . . . . . . _________    (example: EFS555i)
   Generic Search Argument. . . . ____________(example: SQL OVERFLOW)

Recipient Information
   Recipient E-mail Address. . . ______________________________________
                                            example: name@company.com)
   Specify Additional Addresses? _          (Y-Yes/N-No)
JOB JCL Statement. . . . . . . . _          (E-Edit/Y-Yes/N-No)
```

The panel for the “on demand” interface is used to activate IMS ASN dynamically from a TSO session. It also provides the ability to specify search criteria, or to replicate the search criteria which was initiated by an abend-driven invocation of IMS ASN.

The program driving this panel parses and assembles the search criteria into a JCL stream. The user can then submit the job generated by IMS ASN.

If ‘Y’ is specified in the ‘Specify Additional Addresses?’ field above, the user will be prompted in edit mode on a member of a PDS where an unlimited number of e-mail recipients can be entered.

If ‘E’ is specified in the “JOB JCL Statement” field, a JOB statement is presented for the user to customize. It is retained for future use. If ‘Y’ is specified, the previously customized job statement is used.

‘Product Name’ is a required field and at least one search argument must be provided on the panel.

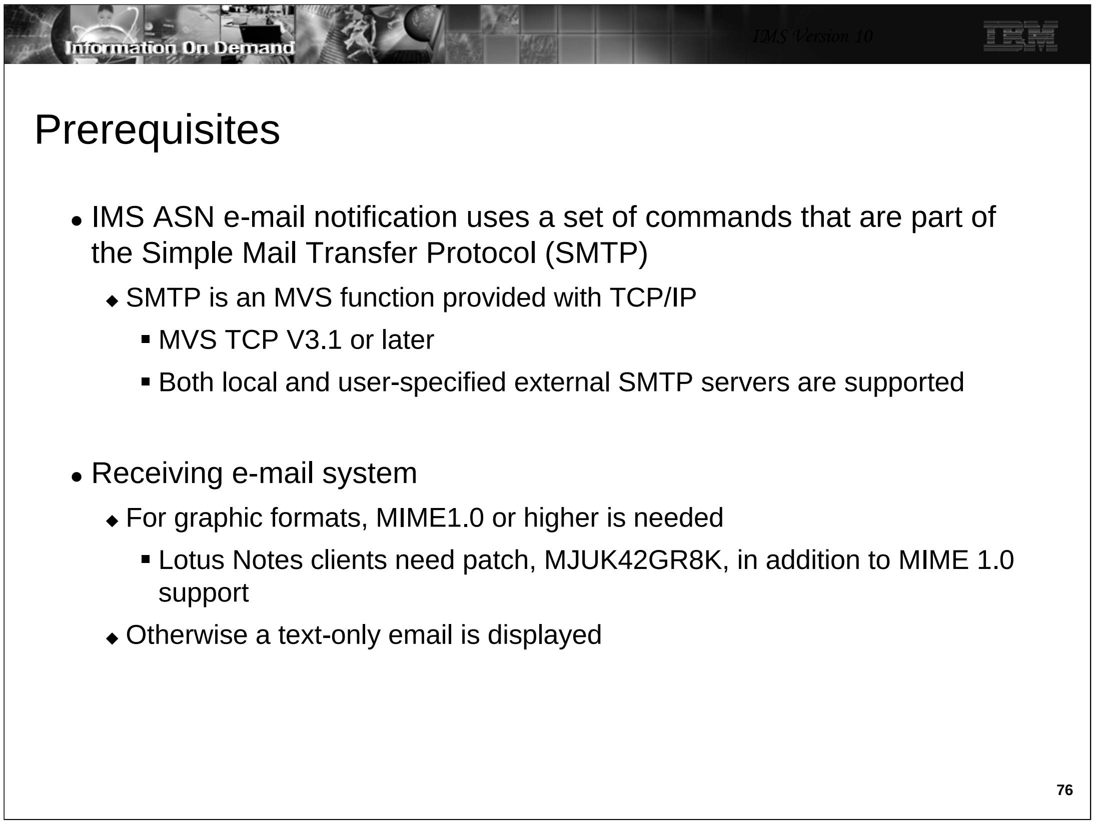
Information On Demand
IMS Version 10

Prerequisites

- IMS ASN e-mail notification uses a set of commands that are part of the Simple Mail Transfer Protocol (SMTP)
  - SMTP is an MVS function provided with TCP/IP
    - MVS TCP V3.1 or later
    - Both local and user-specified external SMTP servers are supported
- Receiving e-mail system
  - For graphic formats, MIME1.0 or higher is needed
    - Lotus Notes clients need patch, MJUK42GR8K, in addition to MIME 1.0 support
  - Otherwise a text-only email is displayed

Since the e-mail is sent using SMTP, the IMS system must reside on a z/OS system that has enabled MVS TCP/IP. SMTP is a standard application that is delivered with the TCP/IP stack. For optimal results, the e-mail manager should support graphics and images. With consideration for the small screen size of cellular devices, the sample text message does not include the creation of URLs. The purpose of the text message is to notify users of the occurrence of the abend.

Benefits

- The IMS Abend Search and Notification function
  - Provides
    - Timely and automatic notification of IMS system problems
    - Quick access to explanation of the error and possible causes
    - Quick and direct access to IBM-supplied diagnostics
  - Facilitates problem determination
    - Inclusion of links that are of interest to the problem
  - Reduces problem resolution time

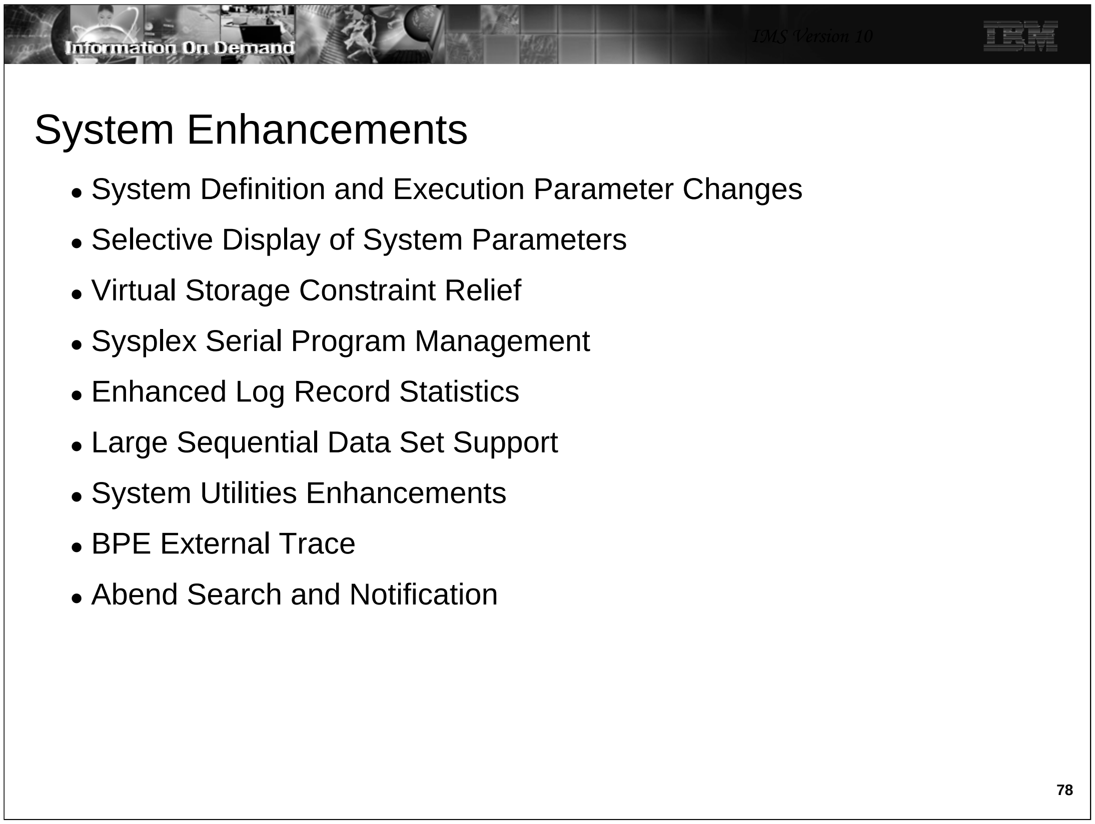

System Enhancements

- System Definition and Execution Parameter Changes
- Selective Display of System Parameters
- Virtual Storage Constraint Relief
- Sysplex Serial Program Management
- Enhanced Log Record Statistics
- Large Sequential Data Set Support
- System Utilities Enhancements
- BPE External Trace
- Abend Search and Notification