IBM IMS Buffer Pool Analyzer for z/OS
1.4

User's Guide

IBM

Note:

Before using this information and the product it supports, read the information in "Notices" on page 121.

Contents

About this information

IBM® IMS Buffer Pool Analyzer for z/OS® (also referred to as IMS Buffer Pool Analyzer, or simply, Buffer Pool Analyzer) is an IMS Tools product that provides a way to analyze database buffer pool performance for both IMS batch jobs (DLI and DBB) and IMS subsystems (DBCTL and DB/DC).

These topics provide instructions for installing, configuring, and using IMS Buffer Pool Analyzer.

To use these instructions, you must have already installed IMS Buffer Pool Analyzer by completing the instructions in the *Program Directory for IBM IMS Buffer Pool Analyzer for z/OS* (GI10-8486-03), which is included with the product media and is also available on the IMS Tools Product Documentation page.

These topics are designed to help database administrators, system programmers, application programmers, and system operators perform the following tasks:

- Understand the capabilities of the functions that are associated with IMS Buffer Pool Analyzer
- Install and operate IMS Buffer Pool Analyzer
- Customize your IMS Buffer Pool Analyzer environment
- Diagnose and recover from IMS Buffer Pool Analyzer problems
- Use IMS Buffer Pool Analyzer with other IMS products

To use these topics, you should have a working knowledge of:

- The z/OS operating system
- ISPF
- SMP/E
- IMS

Always refer to the IMS Tools Product Documentation web page for complete product documentation resources:

https://www.ibm.com/support/pages/node/712955

The IMS Tools Product Documentation web page includes:

- Links to IBM Documentation for the user guides ("HTML")
- PDF versions of the user guides ("PDF")
- Program Directories for IMS Tools products
- Recent updates to the user guides, referred to as "Tech docs" ("See updates to this information!")
- Technical notes from IBM Software Support, referred to as "Tech notes"
- White papers that describe product business scenarios and solutions

Chapter 1. IMS Buffer Pool Analyzer overview

IBM IMS Buffer Pool Analyzer for z/OS provides a way to analyze database buffer pool performance for both IMS batch jobs (DLI and DBB) and IMS subsystems (DBCTL and DB/DC).

Topics:

What's new in IMS Buffer Pool Analyzer

This topic summarizes the technical changes for this edition.

New and changed information is indicated by a vertical bar (|) to the left of a change. Editorial changes that have no technical significance are not noted.

Revision markers follow these general conventions:

- Only technical changes are marked; style and grammatical changes are not marked.
- If part of an element, such as a paragraph, syntax diagram, list item, task step, or figure is changed, the entire element is marked with revision markers, even though only part of the element might have changed.
- If a topic is changed by more than 50%, the entire topic is marked with revision markers (so it might seem to be a new topic, even though it is not).

Revision markers do not necessarily indicate all the changes made to the information because deleted text and graphics cannot be marked with revision markers.

SC19-3980-02 - August 2022

Description	Related APARs
Documentation changes related to the description of the GTF trace report. The following topic has been updated: • "GTF trace report" on page 34	N/A

SC19-3980-01 - March 2021

Description	Related APARs
The following message has been updated: BPL687E	PH33675
This APAR fixes several problems. The following topic and some of its subtopics were updated: • Chapter 5, "IMS buffer pool reports," on page 27	PH30457

Description	Related APARs
Support for IMS-managed ACBs. If IMS Buffer Pool Analyzer is used in an IMS-managed ACBs environment, it can obtain ACB information from the IMS directory. No changes are required to existing IMS Buffer Pool Analyzer JCL. The following topics were updated or created: • "Security requirements for access to data sets" on page 12 • "Customizing buffer pool reports using SYSIN statements" on page 29 • "Varied hit ratio report" on page 43 • Updated messages: BPL109E, BPL189I • New messages: BPL724I to BPL730E	PH29309
The fix for this APAR provides a new function to support enhanced IMS Tools Knowledge Base (ITKB) product registration capability. The following topic was updated: • "Support for IMS Tools Knowledge Base" on page 12	PH18012
This APAR supports RESLIB concatenation to LINKLIST. The following topics were updated or created: • "BPLTRACE started task parameters" on page 15 • Updated messages: BPL159E, BPL702E	PI95633
Documentation changes regarding working under the IMS Dynamic database buffer pools environment. The following topics were updated or created: • "What does IMS Buffer Pool Analyzer do?" on page 3 • "How to gather data" on page 15	N/A
Documentation changes for APAR PI87559 and other updates related to multiple traces. APAR PI87559 improves configuration trace record handling. The following topics were updated or created: • "Combining multiple traces into a single report" on page 19 • "Generating buffer pool reports" on page 28 • "BPLCF JCL" on page 60 • New message: BPL723W	PI87559

SC19-3980-00 - August 2013

Description	Related APARs
• Library names, member names, and JCL samples are updated following changes made for IMS Buffer Pool Analyzer 1.4. • "Allocating data sets" on page 9 is updated. • "Support for IMS Tools Knowledge Base" on page 12 is updated. • Some messages are updated and new messages are added in "Runtime messages (BPL)" on page 77.	N/A

What does IMS Buffer Pool Analyzer do?

IMS Buffer Pool Analyzer provides more information than just IMS database buffer pool hit ratios and I/O rates. It provides a way to determine the impact of buffer pool changes before they are made.

Product features

IMS Buffer Pool Analyzer provides the following features:

- Reviews your buffer pool environment and recommends changes to the number of buffers in each subpool to improve performance. Specific recommendations are made for the number of buffers for each subpool, and the resulting change in storage usage and reduction in number of database reads.
- Provides information about the efficiency of OSAM cached buffers in a coupling facility (CF), and shows the effect of changing the amount of storage allocated to OSAM buffers in the CF cache structure.
- Reviews the databases that are allocated to each subpool and document subpools that include databases with I/O access ratios that are inconsistent, providing the information that you can use to change your subpool configuration.
- Identifies storage that is wasted because the most efficient buffer sizes are not available for database data sets.
- Determines if adding or subtracting buffers will improve the performance of a selected buffer pool.
- Models buffer pool usage to determine I/O rates for various numbers of buffers in each pool.
- Identifies the databases that most heavily use each of the database subpools.
- Performs scenario analysis, including analyzing the impact of creating new buffer pools or changing the block size of a database.
- Reports on IMS OSAM, VSAM, and DEDB lookaside buffer pools, and IMS OSAM buffers that are cached in a coupling facility structure. The reports provide information that is required to make informed decisions about implementing changes to the structure and sizes of your database buffer pools and the information that should be used with the tuning methods that are also described in this document.
- Requires no changes to any IMS control region or batch job during IMS Buffer Pool Analyzer installation. IMS can continue running while you are implementing the product. IMS Buffer Pool Analyzer dynamically accesses IMS data sets and control blocks from the trace address space to prepare to gather the information that is required to produce buffer pool reports.
- IMS Buffer Pool Analyzer works under an IMS dynamic database buffer pools environment with several restrictions. For details, see "How to gather data" on page 15.

Product benefits

IMS Buffer Pool Analyzer provides the following benefits:

- Eliminates guesswork when making changes to your buffer pool configuration so that you can understand the impact of changes to your system before you implement them.
- Reduces IMS outages that are needed to make changes to your buffer pool configuration and achieve your buffer pool tuning objectives without having to try one change, and then another, and so on.
- Reduces IMS real storage and virtual storage usage without the fear of severe impact on IMS performance and user response time. You can use IMS Buffer Pool Analyzer to model changes and see what the impact of subpool storage reductions will be before you implement changes.
- Determines which buffer subpools have the most space for additional storage and hence, will result in the highest reduction in I/O rate when the subpool is increased in size.
- Makes changes in your buffer pool configuration by adding and removing subpools or changing databases that are using subpools by understanding what the impact of the change will be on your system. IMS Buffer Pool Analyzer allows you to model different buffer pool configurations and see what the effect is on overall system performance.
- Balances your IMS subpools to improve individual database data set hit ratios. IMS Buffer Pool Analyzer identifies database data sets with much higher or lower hit ratios than other database data sets in each

subpool. You can use this information to create new subpools that include database data sets that are more similar in access patterns and improve overall buffer pool performance.

- When making changes to your database access method or block size or CI size, you can model the changes and understand the impact of database changes on your buffer pool configuration. You can have IMS Buffer Pool Analyzer make recommendations on changes to your buffer pool configurations to accommodate the changes that are being made to your database data sets.
- Includes OSAM coupling facility caching information to allow you analysis of the benefit of caching OSAM buffers in a coupling facility.
- Provides support for IMS Tools Knowledge Base to store and retrieve IMS Buffer Pool Analyzer reports.
- Supports IMS-managed ACBs. In an IMS-managed ACBs environment, IMS can manage the runtime application control blocks (ACBs) for databases and program views for you. When IMS manages ACBs, IMS no longer requires DBD, PSB, and ACB libraries. In an IMS-managed ACBs environment, IMS Buffer Pool Analyzer can obtain ACB information from the IMS directory instead of from ACB libraries.

Hardware and software prerequisites

IMS Buffer Pool Analyzer operates on any hardware supported by IMS. IMS Buffer Pool Analyzer supports several IMS software releases that run on OS/390 or z/OS.

Complete information about installation requirements, prerequisites, and procedures for IMS Buffer Pool Analyzer is located in the Program Directory for IMS Buffer Pool Analyzer for z/OS.

Hardware prerequisites

IMS Buffer Pool Analyzer operates on any hardware that is supported by IMS.

Software prerequisites

IMS Buffer Pool Analyzer is designed to operate with any version of z/OS that supports the version of IMS that you are running. All supported releases of IMS are supported by IMS Buffer Pool Analyzer.

Product installation

The installation process allocates and populates two libraries required for product use:

- SBPLSAMP library, which contains sample JCL
- SBPLLINK library, which contains the executable programs

Components

IMS Buffer Pool Analyzer has two parts: a data gathering component, which uses the Generalized Trace Facility (GTF), and batch reporting components, which read the output from one or more data gathering traces.

The data gathering component uses the GTF to record database data set I/O activity. The GTF records that are recorded in IMS Buffer Pool Analyzer traces include:

- VSAM buffer trace records (GTF F61 records)
- IMS OSAM trace records (GTF F4F records)
- Buffer Pool Analyzer trace records (GTF 440 records)

Buffer Pool Analyzer trace records record information about DEDB lookaside buffer activity along with IMS buffer pool configuration information, including subpools, database data sets, coupling facility structures, and IMS options.

The batch reporting components provide batch reports that can be used to assist you in making changes to your IMS database buffer pools. The reports provide recommendations for buffer pool changes, and provides the details that were used to determine recommended changes.

In addition to making recommendations for changes in the number of buffers, the Buffer Pool report allows you to model changes to your buffer pool configuration. You can add new subpools and assign databases that will use the new subpools. These modeling abilities allow you make changes to your DL/I batch or IMS control region database buffer pools with confidence in the resulting performance.

IMS Buffer Pool Analyzer process flow

Collectively, the report types for IMS Buffer Pool Analyzer are OSAM, VSAM and DEDB lookaside buffer pools. Summary reports (included in the Buffer Pool report) provide recommendations for changes to your database buffer pool environment.

IMS Buffer Pool Analyzer has two parts:

- A data gathering component that uses the Generalized Trace Facility (GTF)
- Batch reporting components that read the output from one or more of the traces

The data gathering component uses GTF to record IMS subsystem configuration information and database data set buffer activity. Buffer Pool Analyzer uses GTF to record the following trace record types.

- VSAM buffer activity (GTF F61 records)
- IMS OSAM GTF trace records (GTF F4F records)
- Buffer Pool Analyzer trace records (GTF 440 records) that include IMS configuration information and DEDB lookaside buffer pool activity

GTF records this information in a trace data set, which is used by the modeling component to produce the database buffer pool modeling reports.

The batch reporting components provide batch reports that you can use to assist you in making changes to your IMS database buffer pools. The reports provide recommendations for buffer pool changes, and the details that were used to determine the recommended changes.

The following figure illustrates the flow of data gathering in an IMS Buffer Pool Analyzer environment:
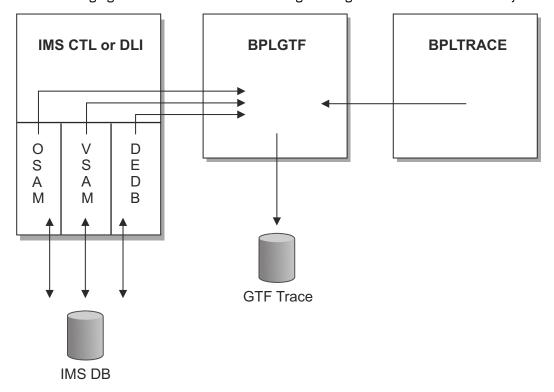

Figure 1. IMS Buffer Pool Analyzer process flow

Buffer pool report definitions

IMS Buffer Pool Analyzer uses several measures to analyze the performance of database buffer subpools. These measures are used in almost all reports and, because they are critical to understanding the data in the reports, their definitions are presented here.

Requests per second

The number of requests per second is an average for the number of database buffer requests generated by IMS. It includes only buffer requests that represent read activity; buffer write activity is not counted. The number of requests per second provides a way to compare subpool and database data set activity.

Reads per second

The number of reads per second is an average for the number of DASD I/Os required to fulfill the database buffer requests generated by IMS. The number of reads per second provides a way to compare the I/O rates for different subpools and database data sets.

Hits per second

The number of hits per second is an average for the number of database buffer requests generated by IMS that did not require any DASD I/O. The number of hits per second is calculated as the difference between the number of requests per second and the number of reads per second.

Hit ratio

The hit ratio for a database data set or subpool is the percent of database buffer requests that were satisfied from the buffer pool. That is, the percent of database buffer requests that did not require any I/O to fulfill. The following figure shows how the hit ratio is calculated:

```
Number of requests – Number of reads
------------------------------------   x   100%
Number of requests
```

Figure 2. Calculating the hit ratio

Buffer life

Buffer life is the average length of time that a buffer remains in the buffer pool between requests for that specific block or CI. Buffer life reflects how quickly the buffers in a pool are discarded to make space in the pool for a new block or CI request. The following figure shows how buffer life is calculated:

```
Elapsed time x Number of buffers
--------------------------------
Number of reads
```

Figure 3. Calculating buffer life

Buffer life is not calculated for specific database data sets. It is an indication of how quickly buffers in a specific subpool are discarded, and is therefore only reported at the subpool level.

Marginal reduction

Marginal reduction statistics attempt to indicate the improvement in performance of a buffer pool. Marginal reduction is calculated only in the subpool models. Marginal reduction applies when changing the number of buffers from the prior row in the model to the current row; marginal reduction shows the number of DASD I/Os per hour that will be reduced per kilobyte (1024 bytes) of storage added to the buffer pool.

The following figure shows how marginal reduction is calculated:

```
Reduction in I/Os per hour (from the prior row)
-----------------------------------------------
    Increase in storage in KB (from prior row)
```

Figure 4. Calculating marginal reduction

Marginal reduction is calculated so that you can compare the relative improvement in performance when adding buffers to a buffer pool. For example, assume that one subpool has a marginal reduction of two and another has a marginal reduction of 200. If the same amount of storage was added to each pool, the reduction in DASD reads for the second pool would be 100 times as great as the reduction in reads for the first pool.

Service updates and support information

Service updates and support information for this product, including software fix packs, PTFs, frequently asked questions (FAQs), technical notes, troubleshooting information, and downloads, are available from the web.

To find service updates and support information, see the following website:

IBM Support: IMS Buffer Pool Analyzer for z/OS

Product documentation and updates

IMS Tools information is available at multiple places on the web. You can receive updates to IMS Tools information automatically by registering with the IBM My Notifications service.

Information on the web

Always refer to the IMS Tools Product Documentation web page for complete product documentation resources:

https://www.ibm.com/support/pages/node/712955

The IMS Tools Product Documentation web page includes:

- Links to IBM Documentation for the user guides ("HTML")
- PDF versions of the user guides ("PDF")
- Program Directories for IMS Tools products
- Recent updates to the user guides, referred to as "Tech docs" ("See updates to this information!")
- Technical notes from IBM Software Support, referred to as "Tech notes"
- White papers that describe product business scenarios and solutions

IBM Redbooks® publications that cover IMS Tools are available from the following web page:

http://www.redbooks.ibm.com

The IBM Information Management System website shows how IT organizations can maximize their investment in IMS databases while staying ahead of today's top data management challenges:

https://www.ibm.com/software/data/ims/

Receiving documentation updates automatically

To automatically receive emails that notify you when new technote documents are released, when existing product documentation is updated, and when new product documentation is available, you can register with the IBM My Notifications service. You can customize the service so that you receive information about only those IBM products that you specify.

To register with the My Notifications service:

1. Go to http://www.ibm.com/support/mynotifications

2. Enter your IBM ID and password, or create one by clicking **register now**.

3. When the My Notifications page is displayed, click **Subscribe** to select those products that you want to receive information updates about. The IMS Tools option is located under **Software** > **Information Management**.

4. Click **Continue** to specify the types of updates that you want to receive.

5. Click **Submit** to save your profile.

How to send your comments

Your feedback is important in helping us provide the most accurate and highest quality information. If you have any comments about this or any other IMS Tools information, you can take one of the following actions:

• Click the Feedback button at the top of the IBM Documentation topic that you are commenting on.

• Click the Contact Us tab at the bottom of any IBM Documentation topic.

• Send an email to ibmdocs@us.ibm.com. Be sure to include the book title, topic or section title, specific text, and your comment.

To help us respond quickly and accurately, please include as much information as you can about the content you are commenting on, where we can find it, and what your suggestions for improvement might be.

Accessibility features

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use a software product successfully.

The major accessibility features in this product enable users to perform the following activities:

• Use assistive technologies such as screen readers and screen magnifier software. Consult the assistive technology documentation for specific information when using it to access z/OS interfaces.

• Customize display attributes such as color, contrast, and font size.

• Operate specific or equivalent features by using only the keyboard. Refer to the following publications for information about accessing ISPF interfaces:

– *z/OS ISPF User's Guide, Volume 1*

– *z/OS TSO/E Primer*

– *z/OS TSO/E User's Guide*

These guides describe how to use the ISPF interface, including the use of keyboard shortcuts or function keys (PF keys), include the default settings for the PF keys, and explain how to modify their functions.

Chapter 2. Customizing IMS Buffer Pool Analyzer

In order to run IMS Buffer Pool Analyzer effectively, you should customize it for your environment.

Topics:

APF authorization

The IMS Buffer Pool Analyzer library SBPLLINK must be APF-authorized on each MVS system on which the data gathering component runs.

Allocating data sets

In addition to the data sets in the SMP/E maintained libraries, IMS Buffer Pool Analyzer requires two additional data sets: *hlq*.PARMLIB and *hlq*.TRACE. SBPLSAMP library member BPLALCDS contains a sample job to allocate these two data sets.

hlq.**PARMLIB**
The BPLTRACE started task populates this PDS with parameters used by the Generalized Trace Facility (GTF) during initialization. The data set needs only to be allocated properly before execution. See sample job BPLALCDS for the allocation parameters.

hlq.**TRACE**
The TRACE data set is used to collect GTF data. The trace data set name is specified in the BPLGTF PROC as well as in the report jobs. It is recommended that the data set be allocated without LRECL or BLKSIZE parameters, allowing GTF to choose optimal parameters based on the hardware on which the data set resides. See sample job BPLALCDS for the allocation parameters.

Space allocation is variable and depends on the:

- Length of time that GTF trace data is gathered
- Amount of database I/O performed by the job or IMS subsystem being monitored
- Amount of VSAM I/O activity on the MVS system

A busy IMS subsystem can easily fill a 1000-cylinder 3390 trace data set in minutes. Also, note that GTF uses no secondary space extents. When allocating the *hlq*.TRACE data set, ensure that you specify the **CONTIG** parameter in the SPACE= keyword of the DD statement.

If GTF fills a DASD trace data set, it wraps; that is, it writes over the beginning of the data set. It is important for the trace data set to be large enough to hold the entire trace.

You can avoid space problems, such as wrapping the trace data set and any impact on your DASD subsystem, by using tape for the GTF trace data. On a very busy system, tape output can be more efficient and reduce the chance that GTF might have to discard trace records due to a lack of GTF trace buffer space.

Customizing the sample started tasks

Two started tasks (tasks started by an operator command) are required for the data gathering component. Sample PROCs for the started tasks are distributed in the SBPLSAMP library.

About this task

The names of the sample PROCs are BPLTRACE and BPLGTF.

Procedure

To customize the started tasks, follow these steps:

1. Copy BPLTRACE and BPLGTF to a system PROCLIB data set known to JES for started tasks PROCs.

   You can change the names of the started task PROCs to conform to your installation's standards. The name of the BPLTRACE started task is used only in the operator command used to start the data collection process. If you change the name of the BPLTRACE started task, simply use the new name in place of BPLTRACE in the **START** command.

   The name of the BPLGTF PROC is specified in the BPLTRACE started task PROC. To change the name of the BPLGTF PROC, simply change the GTFPROC= keyword value from BPLGTF to the new name of the started task in the BPLTRACE PROC.

   The BPLGTF started task PROC must be reviewed and updated as described below. The following figure shows the distributed sample BPLGTF PROC.

```
//BPLGTF PROC RGN=4096K,
//         PARMLIB=BPL.PARMLIB,        DSN OF PARMLIB
//         TRACE=BPL.TRACE,            DSN OF TRACE
//           BLOK=256K,         AMOUNT OF ECSA TRACE RECORD STORAGE
//           NCP=64,            NUMBER OF CONCURRENT TRACE I/O
//           SIZE='SIZE=4M'   ONLY VALID FOR z/OS 1.6 AND ABOVE
//*
//* This PROC is used by IMS Buffer Pool Analyzer.
//*
//BPLGTF      EXEC PGM=AHLGTF,REGION=&RGN,
//            PARM='NOPROMPT,BLOK=&BLOK,&SIZE'
//SYSLIB      DD   DSN=&PARMLIB(BPLTRACE),DISP=SHR
//IEFRDER     DD   DSN=&TRACE,DISP=SHR,DCB=NCP=&NCP
//*
```

Figure 5. Sample BPLGTF PROC

where:

RGN *symbolic parameter*
The RGN symbolic parameter defines the region size used by the BPLGTF started task. The default value of 4096K should be sufficient for most installations.

PARMLIB *symbolic parameter*
This parameter defines the data set name of the PARMLIB data set whose allocation is described in "Allocating data sets" on page 9.

TRACE *symbolic parameter*
This parameter defines the data set name of the TRACE data set whose allocation is described in "Allocating data sets" on page 9.

BLOK *symbolic parameter*
This parameter identifies the amount of ECSA that GTF will use to store trace records before writing them to the trace data set. ECSA can range from 40K to 99999K.

Note: Ensure that the value specified will not cause a shortage of ECSA; otherwise, severe z/OS errors can result.

The sample PROC specifies BLOK=256K, which should be adequate for most installations. This value can be increased if GTF records are lost during the data gathering process, as reported in the GTF trace report included in the report jobs.

NCP *symbolic parameter*
This symbolic parameter specifies the number of channel programs that the output data set uses for the GTF trace data set. The value specified in the sample PROC of NCP=64 should be adequate for most installations. You can increase this value if GTF records are lost during the data gathering process, as reported in the GTF trace report included in the report jobs.

SIZE= PROC *symbolic parameter*
This symbolic parameter specifies the amount of GTF buffer space to be used and allocated in the private area of the BPLGTF started task. This symbolic parameter can be used only on an MVS system of a supported release of z/OS. If you are running an earlier version of z/OS, specify **SIZE**= with no value in your BPLGTF PROC. For higher levels of z/OS, the value in the sample PROC of SIZE='SIZE-2M' should be adequate for most installations. You can increase this value if GTF records are lost during the data gathering process, as reported in the GTF trace report included in the report jobs.

IEFRDER DD
This DD describes the data set used by GTF when writing trace records. If you choose to use a DASD data set for your trace data, the sample PROC can be used as distributed. If you choose to use a TAPE data set for your trace data, change the DD statement to allocate a new tape output data set. For example, you could change the IEFRDER DD to the following to use a UNIT name of TAPE as shown in the following figure:

```
//IEFRDER     DD   DSN=&TRACE,DISP=(NEW,CATLG),
//                 UNIT=TAPE,LABEL=RETPD=30,
//                 DCB=NCP=&NCP
```

Figure 6. Replacement IEFRDER DD when writing GTF trace to tape

2. Review and update the BPLTRACE started task PROC as described below. The following figure shows the distributed sample BPLTRACE PROC.

```
//BPLTRACE PROC RGN=4M,          REGION SIZE

//   SBPLLINK=BPL140.SBPLLINK,  DSN OF SBPLLINK LIBRARY
//   PARMLIB=BPL.PARMLIB,         DSN OF BPL PARMLIB
//           GTFPROC=BPLGTF, GTF PROC NAME TO BE STARTED
//           OPT=,            OPTIONS
//           IMSID=,          NAME OF IMS CTL RGN IMSID
//           JOB=,            NAME OF DLI/DBB BATCH JOB
//           DT=,             MINUTES TO WAIT BEFORE STARTING
//           WT=,             MAX MINUTES TO WAIT FOR JOB/IMS TO START
//           GT=15,            MINUTES TO RUN GTF TRACE
//*
//* This PROC is used by the IMS Buffer Pool Analyzer
//*
//BPLTRACE EXEC PGM=BPLTRACE,REGION=&RGN,
//       PARM=('IMSID=&IMSID,JOB=&JOB,GTFPROC=&GTFPROC',
//       'GTFTIME=&GT,WAITTIME=&WT,DELAY=&DT,&OPT')
//STEPLIB  DD DSN=&SBPLLINK,DISP=SHR
//SYSLIB   DD DSN=&PARMLIB(BPLTRACE),DISP=SHR
//SYSPRINT DD SYSOUT=X
//SYSUDUMP DD SYSOUT=X
```

Figure 7. Sample BPLTRACE PROC

where:

RGN *symbolic parameter*
The RGN symbolic parameter defines the region size used by the BPLTRACE started task. The default value of 4M should be sufficient for most installations.

SBPLLINK *symbolic parameter*
The SBPLLINK= value must specify an APF authorized library which contains the IMS Buffer Pool Analyzer load modules.

PARMLIB *symbolic parameter*
The data set name must be changed to reflect the data set name specified in the BPLALCDS sample job (the *hlq*.PARMLIB data set)

GTFPROC

The name of the BPLGTF started task. The distributed name of the started task is BPLGTF. If you chose to change the name, as described in the preceding section, then update the BPLTRACE PROC to specify the PROC name you created.

The remaining symbolic parameters in the BPLTRACE PROC are typically supplied on the MVS START command. These parameters are documented in .

Security requirements for access to data sets

There are no security requirements for IMS Buffer Pool Analyzer other than typical data set access security requirements.

The **BPLTRACE** PROC requires:

- Update access to the PARMLIB data set.
- Control access to the RECON data sets used by the batch job or IMS subsystem being monitored.
- In an ACBLIB-managed ACBs environment: Read access to the active ACBLIB data sets used by any IMS subsystems to be monitored. For batch (DLI/DBB) jobs, read access is not required.
- In an IMS-managed ACBs environment: Read access to the IMS catalog and IMS directory used by any IMS subsystems to be monitored. For batch (DLI/DBB) jobs, read access is not required.

The BPLGTF PROC requires:

- Update access to the TRACE data set specified in the IEFRDER DD statement in the PROC.
- Read access to the PARMLIB data set specified in the SYSLIB DD statement.

Support for IMS Tools Knowledge Base

If your installation currently uses IMS Tools Knowledge Base, which is a component of IBM IMS Tools Base for z/OS (5655-V93), IMS Buffer Pool Analyzer can take advantage of the product's ability to archive and retrieve reports. The use of IMS Tools Knowledge Base is optional.

About this task

If you choose to use IMS Tools Knowledge Base, you must prepare the IMS Tools Knowledge Base repository for IMS Buffer Pool Analyzer reports.

To retain a history of IMS Buffer Pool Analyzer reports in IMS Tools Knowledge Base, use the product administration utility to define the report retention period. After IMS Buffer Pool Analyzer is registered, you can edit the report retention period. By default, the reports are retained for 30 days. A sample job is provided in the SBPLSAMP library. The following figure shows the sample job that is provided in the BPLITKB member of the SBPLSAMP library.

Important: Do not use the PRODUCTID=IB option when you register IMS Buffer Pool Analyzer with IMS Tools Knowledge Base. IMS Buffer Pool Analyzer is not defined in the default definition table, so only the most recent reports will be stored in the repository if you specify PRODUCTID=IB.

```
//BPLITKB  JOB (ACCT#),BPL140,NOTIFY=&SYSUID,
// class="K",MSGCLASS=A,TIME=10
//*----------------------------------------------------------------
//*
//*  LICENSED MATERIALS - PROPERTY OF IBM
//*
//*  5697-H77
//*
//*  Copyright IBM Corp. 2002, 2013
//*  ALL RIGHTS RESERVED.
//*
//*  Copyright Rocket Software, Inc. 2014
//*  ALL RIGHTS RESERVED.
//*
//*------------------------------------------------------------------
//*
//* THIS OPTIONAL JOB DEFINES IMS BUFFER POOL ANALYZER TO THE
//* IMS TOOLS KNOWLEDGE BASE.  IF YOU DO NOT HAVE THE IMS TOOLS
//* KNOWLEDGE BASE YOU SHOULD SKIP THIS PROCESS.
//*
//*********************************************************************
//* TOOL:     IMS BUFFER POOL ANALYZER                             *
//* DESC:     SAMPLE JCL TO DEFINE PRODUCTS TO IMS TOOLS KB        *
//* SOURCE:   SBPLSAMP(BPLITKB)                                    *
//*********************************************************************
//*
//* THIS JOB WILL DEFINE THE ITKB SERVER TO THE ABOVE TOOL.
//*
//* INSTRUCTIONS:
//*  1) MODIFY THE JOBCARD TO MEET YOUR STANDARDS.
//*  2) CHANGE 'ITKBSRVR=itkbsrvr'
//*     TO YOUR ITKB SERVER NAME.
//*  3) CHANGE 'your.SHKTLOAD' ALL
//*     TO YOUR DATA SET NAME FOR THE ITKB LOAD LIBRARY.
//*  4) CHANGE SBPLLINK=sbpllink
//*     TO YOUR DATA SET NAME FOR THE BPA LOAD LIBRARY.
//*  5) IF YOU USE IMS TOOLS KNOWLEDGE BASE V1.6 OR LATER, REMOVE THE
//*     ATTERISK AT THE BEGINNING OF THE LISTX COMMAND.
//*
//   SET  SBPLLINK=sbpllink
//   SET  SHKTLOAD=shktload
//   SET  ITKBSRVR=itkbsrvr
//*
//* THE DATA SET NAME SPECIFIED FOR THE SHKTLOAD KEYWORD IS THE DATA
//* SET THAT CONTAINS THE IMS TOOLS KNOWLEDGE BASE SOFTWARE.
//*
//* ----------------------------------------------------------------
//* DEFINE IMS BUFFER POOL ANALYZER TO ITKB
//* ----------------------------------------------------------------
//ADDPROD EXEC PGM=HKTAPRA0,PARM='ITKBSRVR=&ITKBSRVR'
//STEPLIB  DD DISP=SHR,DSN=&SBPLLINK
//         DD DISP=SHR,DSN=&SHKTLOAD
//SBPLLINK DD DISP=SHR,DSN=&SBPLLINK
//TEMPLATE DD DUMMY
//SYSPRINT DD SYSOUT=*
//OUTRPT   DD SYSOUT=*
//OUTRPTX  DD SYSOUT=*
//SYSIN    DD *
  ADDPROD TABLE=BPL$RIB0,REPLACE=YES
  LIST    PRODUCTID=IB,REPORTID=*
* LISTX   PRODUCTID=IB
/*
```

Figure 8. Sample BPLITKB JCL

Note that in addition to updating the JOB card to meet your installation's standards, you must supply appropriate values for the three symbolic variables.

SBPLLINK *symbolic parameter*
This symbolic parameter must be set to the data set name of the IMS Buffer Pool Analyzer SBPLLINK library.

SHKTLOAD *symbolic parameter*
This symbolic parameter must be set to the data set name of the IMS Tools Knowledge Base SHKTLOAD load library.

ITKBSRVR *symbolic parameter*

> This symbolic parameter must be set to the XCF group name for the IMS Tools Knowledge Base servers.

You should expect to receive the following messages from execution of the first step in the above job:

`HKT2178I Attempting to add an IBM product`

`HKT2158I Request completed successfully`

These messages are expected and normal. The second step in the sample job lists the IMS Buffer Pool Analyzer definitions in the IMS Tools Knowledge Base repository. The output from this step will show the definitions as follows.

```
PRODUCT  ID: IB RELEASE  010300 LONG NAME:  IMS BUFFER POOL ANALYZER
   REPORT ID    LONG TITLE
    DB           DATABASE INFORMATION
    DS           DATA SET INFORMATION
    GT           GTF TRACE STATISTICS
    MD           MODEL SUBPOOL REPORTS
    SP           SUBPOOL REPORTS
```

Any other messages that occur when defining IMS Buffer Pool Analyzer to the IMS Tools Knowledge Base should be reviewed using the *IMS Tools Base IMS Tools Knowledge Base User's Guide and Reference*.

In order to use IMS Tools Knowledge Base when creating reports, you must supply a PARM field in the report job step EXEC statement. The PARM field must specify the name of the IMS Tools Knowledge Base server. Further information can be found in "Generating buffer pool reports" on page 28. If the ITKBSRVR specification is not included in the PARM of the report JCL, IMS Buffer Pool Analyzer will not archive report output in the knowledge base repository.

Chapter 3. Gathering trace data

You gather trace data in order to discover the performance levels of your database buffer pools. This data is recorded in trace data sets and used to produce buffer pool modeling reports.

Topics:

How to gather data

The data gathering component of Buffer Pool Analyzer uses the GTF to gather data for modeling reports.

You can use multiple GTF trace data sets to generate a single report. When creating multiple trace data sets, the trace data must be consistent in order to generate a meaningful report.

When generating a report using multiple GTF trace data sets, the following information must be the same:

- Buffer pool definitions - This includes not only the same pool buffer sizes and subpool definitions, but the same number of buffers.
- Database definitions - This includes the number and names of the databases in the batch or online subsystem, the DD names associated with each DBD, block sizes or CI sizes, and the subpool assignments for each data set.

Use the BPLTRACE started task to initiate the data gathering component.

You can customize the BPLTRACE started task parameters in the started task PROC, or you can override them on the MVS START command.

Restrictions for gathering trace data: IMS Buffer Pool Analyzer works under an IMS dynamic database buffer pools environment with the following restrictions:

- You must stop BPLTRACE before entering the IMS **UPDATE POOL TYPE(DBAS) SECTION(xxxxxxxx)** command. Then, after that command is completed, restart BPLTRACE by using another trace data set.
- To generate reports, run a job by using the trace data set that was obtained before the IMS **UPDATE POOL** command was entered, and run another job by using the trace data set that was obtained after the IMS **UPDATE POOL** command was entered.

BPLTRACE started task parameters

To perform the task of gathering trace data, the BPLTRACE started task uses the following parameters:**IMSID**, **JOB**, **DT**, **WT**, and **GT**. There are four additional parameters in the BPLTRACE PROC that must also be specified, but these parameters (SBPLLINK, PARMLIB, GTFPROC, and RGN) should be set in the actual PROC JCL, and should not need to be specified as part of the **START** command.

The following parameters are available for the BPLTRACE started task.

IMSID

The IMSID of an IMS DB/DC or DBCTL subsystem to be monitored. Either the **IMSID** or **JOB** parameter must be specified, but not both.

When monitoring DEDB lookaside buffer pools, IMS RESLIB must be specified in STEPLIB/JOBLIB DD of the IMS subsystem or IMS RESLIB must be concatenated to LINKLIST.

JOB

> The name of an IMS batch (DLI or DBB region type) job to be monitored. Either the **IMSID** or **JOB** parameter must be specified, but not both.

DT

> Delay time. The number of minutes to wait before starting to look for the IMS subsystem or batch job name. The number of minutes can range from 0 to 9999. Zero (or no specification) results in no delay time.

> The **DT** parameter allows you to gather GTF data at a later time. For example, to gather data from an IMS control region during heavy nighttime BMP processing, start BPLTRACE at 4:00 PM with DT=600; this delays the start of GTF for 10 hours, until 2:00 AM the following morning.

WT

> Wait time. The number of minutes to wait for the named JOB to start (after the delay time). The number of minutes can range from 0 to 9999. Zero (or no specification) results in no wait time.

> When monitoring a DLI or DBB batch job, when BPLTRACE is started and the delay time has expired, it searches for an IMS batch subsystem with JOB=*jobname* specified. If the job name is not found, BPLTRACE searches the active jobs 10 times per minute until the wait time expires.

> For example, to gather data for a DLI batch job that is expected to start four hours from the current time, use a WT=210 and DT=60. BPLTRACE will wait 3.5 hours to begin looking for the job, and continues searching for the next hour. If the job is not found at the end of the WT period, BPLTRACE issues a warning message and ends.

GT

> GTF time. The number of minutes to gather GTF data. The number of minutes can range from 1 to 9999.

Tip: For batch DLI or DBB jobs, BPLTRACE stops the trace activity automatically when the job step ends.

Examples of starting a trace

An MVS **START** command is used to start a data collection trace. The examples in this section show the **START** commands required to collect data.

These examples assume that the **SBPLLINK**, **PARMLIB**, **GTFPROC**, and **RGN** parameters are set appropriately in the PROC.

Examples

Example 1

Start collecting data immediately for an IMS subsystem with IMSID=IMSA. Collect data for 30 minutes.

```
START BPLTRACE,IMSID=IMSA,GT=30
```

Example 2

Collect data for an IMS subsystem with IMSID=IMSA. Begin data collection in 6 hours, and collect data for 20 minutes.

```
START BPLTRACE,IMSID=IMSA,DT=360,GT=20
```

Example 3

Collect data for a batch job named DLIJOB. The job is currently running. Collect data for 60 minutes or until the job ends.

```
START BPLTRACE,JOB=DLIJOB,GT=60
```

Example 4

Collect data for a batch job named DBBJOB, which is expected to start in 60 - 90 minutes. Collect data for 120 minutes or until the first IMS step in the job ends.

```
START BPLTRACE,JOB=DBBJOB,DT=60,WT=30,GT=120
```

Functions of the BPLTRACE started task

The BPLTRACE task starts and stops the GTF trace; it also collects data from the target address space and passes the additional data to GTF.

Specifically, the BPLTRACE task performs these actions:

1. Waits for the delay time, if requested.
2. Searches for the target job or IMS subsystem.
3. If the target job is not found, retries the search every six seconds until the wait time expires.
4. Builds GTF control cards based on the IMS environment.
5. Obtains configuration information from the target job or IMS subsystem, including DBD names, DD names, data set names, buffer pool allocation, and subpool-to-DBDNAME/DDNAME relationships.
6. Dynamically plants an intercept in IMS DEDB lookaside pool processing that creates a GTF trace record describing each DEDB buffer request, and whether the request was satisfied from memory or the coupling facility.
7. Starts GTF and verifies that GTF initialized successfully.
8. Traces the target job or IMS subsystem's configuration information gathered in step "5" on page 17.
9. After the GTF time interval completes or the batch job ends, turns off the IMS OSAM trace and stops the GTF address space.

Controlling BPLTRACE

The BPLTRACE started task can remain active for long periods of time, depending on the delay time and wait time parameters you choose.

You can use the MVS MODIFY command to obtain the current status of the trace, stop the trace, or bypass wait times requested at the start of the trace.

Obtaining the status of a trace

The BPLTRACE **STATUS** command provides output to the system console that describes the status of the trace and how much time remains in the current state.

Procedure

To obtain the status of a trace, perform the following steps:

1. To obtain the status of a trace, issue the following MVS command:

```
MODIFY BPLTRACE,STATUS
```

2. Verify the output. The following example shows the output of a **STATUS** command:

```
F BPLTRACE,STATUS
BPL063I TRACE REQUESTED FOR IMSID IMS8
BPL066I IMS OSAMGTF TRACE STATUS: STARTED
BPL067I GTF TRACE IN PROGRESS - 0006:52 OF 10 MINUTES REMAINING
BPL068I IMS CONFIGURATION STATUS: COMPLETE
BPL069I STATUS DISPLAY COMPLETE
```

Figure 9. Output of the STATUS command

Results

This status report indicates that the trace is in progress for IMS subsystem IMS8.

The configuration information has already been collected and recorded. Other status information might also appear as a result of the STATUS command. All messages that can be displayed in the status report are described in "Runtime messages (BPL)" on page 77. For the information that is provided in the status report, see "Functions of the BPLTRACE started task " on page 17.

Stopping a trace

BPLTRACE can be stopped at any time, although if the trace is gathering configuration information, an abend might occur in the configuration subtask.

About this task

If the GTF trace has started and the configuration information is complete, stopping the trace provides a valid trace data set for use in reporting.

Procedure

To stop a trace, complete the following steps:

1. Issue the MVS **MODIFY** or MVS **STOP** command.

   The following list shows the MVS commands that you can use to stop the trace:

   - **STOP** BPLTRACE
   - **P** BPLTRACE
   - **MODIFY** BPLTRACE,**STOP**
   - **F** BPLTRACE,**STOP**

   Stopping the trace stops both the BPLTRACE address space and the BPLGTF address space (if it is active). If the GTF trace has started and the configuration information is complete, stopping the trace provides a valid trace data set for use in reporting.

2. Check the log to ensure that the trace was stopped with no errors.

Bypassing the delay time

If a trace was started with a delay time (**DT=**), the delay time can be bypassed.

About this task

For example, if a batch job starts before it was expected to start, and before the delay time expires, you can issue an MVS **MODIFY** command to bypass the delay and start the GTF trace.

To bypass a delay, issue the following MVS **MODIFY** command:

```
F BPLTRACE,START
```

BPLTRACE immediately searches for the job and, once found, starts the GTF trace.

Trace overhead

Overhead is associated with gathering trace data for VSAM, OSAM, and DEDB 1/O activity, as well as for gathering trace data for configuration information. The trace data set should be allocated on a volume that has no busy system data sets.

IMS Buffer Pool Analyzer uses established means to gather trace data.

- For VSAM I/O activity, VSAM GTF record collection for user F61 record types is used.
- For OSAM I/O activity, the IMS OSAM GTF trace is enabled, which generates GTF user F4F records.

- For configuration information and DEDB lookaside buffer information, the BPLTRACE task creates GTF user 440 type records.

The overhead of collecting OSAM/VSAM GTF data has been estimated at 5-10% of the CPU time for the tasks for which data is being collected. Additionally, there is I/O overhead to allow GTF to write trace records to the output data set.

Trace recommendations

The accuracy and usefulness of any report depends on the data gathered. Consider the points presented in this section when gathering trace data.

Sufficient trace data must be gathered. The time required to collect trace data varies based on I/O rates and on database I/O patterns.

It is important to not let GTF trace DASD data sets wrap because some IMS Buffer Pool Analyzer trace records could be lost.

To determine how quickly a DASD trace data set fills, you can use an MVS performance monitor to show how many I/Os occur to the IEFRDER data set of the BPLGTF started task. For a 1000-cylinder trace data set with half-track blocking, 30,000 I/Os would fill the trace data set.

For a batch job with consistent call patterns throughout the job, a 5-minute trace interval could provide more than enough trace data to produce valid reports and projections. On the other hand, an online IMS subsystem with hundreds of databases and I/O patterns that vary by time of day will require additional trace data.

Depending on the amount of DASD available for GTF trace data, an online IMS subsystem could be monitored for 10, 20, or 30 minute intervals. Because an online subsystem might have a batch workload occurring at night which has characteristics different from the daytime workload, several trace data sets should be created for an online system. All the trace data sets can be included in a single execution of the reporting utility to produce projections that include the differing workloads.

To see differences in the projections by workload type, run separate reports with only one workload type included in the reporting job.

Trace data sets which are full can be copied to another DASD device or to tape by using the IEBGENER utility. You can then reuse the trace data set and reduce the amount of DASD space used for trace data.

If you experience any problems with DASD trace data sets, you can simply update the BPLGTF PROC to log the trace data directly to tape.

Combining multiple traces into a single report

When creating multiple trace data sets that will be combined into a single report, the trace data must be consistent to generate meaningful reports. In addition, all trace data sets must include the trace data that was acquired from the same IMS or the same DL/I batch job.

The information that must be the same includes:

- Buffer pool definitions

  This includes not only the same pool buffer sizes and subpool definitions, but the same number of buffers

- Database definitions

  Definitions include:

  - number and names of the databases in the subsystem (batch or online)
  - DD names associated with each DBD
  - block sizes
  - CI size
  - subpool assignments for each database data set

Also, ensure that you do not concatenate trace data sets in report JCL. Each trace data set you want to use as input to one of the report jobs must have a different DD name, each name starting with the word TRACE.

Chapter 4. Block Analysis reports introduction

Block analysis reports provide details about individual database block access requests.

The Block analysis reports can be used to review which blocks are most frequently referenced for each DBDS or subpool. This report information might be helpful when you review database access patterns, database organization, or randomizer effectiveness.

These reports were not designed to assist in the buffer pool tuning process, but they might be useful when you perform other tasks.

Topics:

Block analysis JCL

The Block analysis report uses the same trace data as the buffer analysis report that is produced by running the block analysis JCL.

The JCL to run the Block analysis report is described in this topic.

```
//BPLBLAN JOB
//*
// SET BPL140.SBPLLINK=                    DSN of BPA LOAD LIBRARY
// SET TRACE=BPL.TRACE                TRACE DATA SET NAME
//*
//BPLBLAN EXEC  PGM=BPLBLAN,REGION=64M
//STEPLIB  DD    DSN=&SBPLLINK,DISP=SHR
//SYSPRINT DD    SYSOUT=*
//GTFRPT   DD    SYSOUT=*
//DBDRPT   DD    SYSOUT=*
//SPLRPT   DD    SYSOUT=*
//SYSOUT   DD    SYSOUT=*
//SYSUDUMB DD    SYSOUT=*
//BPLTEMP  DD    UNIT=SYSDA,SPACE=(CYL,(20,20))
//TRACE    DD    DSN=&TRACE,DISP=SHR,DCB=BUFNO=50
//SYSIN    DD    *
```

Figure 10. Block analysis JCL

This sample JCL is included in the SBPLSAMP library in the member BPLBLAN. The DD names used by the Block Analysis report utility are described in the following list.

STEPLIB
  Must reference a data set that contains the Buffer Pool Analyzer load modules.

SYSPRINT
  Report output file. This is a required DD statement. The SYSPRINT file contains error messages and any user specifications as supplied in the SYSIN DD statement. The DCB attributes of this output file need not be specified. They are set as LRECL=133, RECFM=FBA, and DSORG=PS.

GTFRPT
  Report output file. If this DD statement is omitted, the trace file reports will not be produced. The GTFRPT file contains reports that describe the trace input files for this execution of the Block Analysis report utility. The DCB attributes of this output file need not be specified. They are set as LRECL=133, RECFM=FBA, and DSORG=PS.

DBDRPT
  Report output file. This statement is required only if you specify a DBD= statement in the SYSIN DD. The DBDRPT file contains reports that document the most heavily used blocks for database data sets

that you request through statements that are supplied in the SYSIN DD. The DCB attributes of this output file need not be specified. They are set as LRECL=133, RECFM=FBA, and DSORG=PS.

SPLRPT

Report output file. If this DD statement is omitted, the subpool reports are not produced. The SPLRPT file contains reports that document the most heavily used blocks for each subpool used by the IMS subsystem or batch job. The DCB attributes of this output file need not be specified. They are set as LRECL=133, RECFM=FBA, and DSORG=PS.

SYSOUT

Sort report output file. This DD statement is required for Sort to run successfully. The SYSOUT file is used by DFSORT for report and error messages. Your installation might use a different name for its Sort utility report file.

SYSUDUMP

Dump diagnostic output file. This DD statement is not required, but its use is recommended. The SYSUDUMP file will contain a diagnostic dump if the block analysis utility ends abnormally.

BPLTEMP

Temporary disk file that is used between sort phases. This is a required DD statement. If you encounter an X'37' abend in the sort output file, increase the amount of space that is allocated to this temporary file.

TRACE*xxx*

Trace input files. At least one trace DD statement is required. Any DD statement that begins with TRACE is used by the block analysis report as an input file. The trace files must have been created by the IMS Buffer Pool Analyzer BPLTRACE task, and all trace files that are supplied as input to running the Block Analysis report must be from the same IMS subsystem or batch job. Each trace file must be allocated to a different DD name. And, you cannot concatenate trace files to one DD name.

For other considerations, see "Combining multiple traces into a single report" on page 19.

SYSIN

Control statement input file. If this DD statement is not included in your JCL, then no Block Analysis reports are generated, but the subpool is generated. The DCB information for this DD statement must be fixed with a record length of 80. The format of the control statement is documented in "Block analysis control statements" on page 22.

Block analysis control statements

The Block Analysis report provides several options for user customization. User control statements allow you to specify databases for which Block Analysis reports should be created. All control statements must be specified in the SYSIN DD.

Control statements are specified as KEYWORD=*VALUE*. Blanks can be included between words in the statement. Continued statements are not permitted, but multiple statements are allowed.

The following statement is supported:

DBD

This statement allows you to specify the names of databases for which Block Analysis reports will be created. You can specify a database name or ALL to request that all databases be reported.

The following examples show valid block analysis control statements.

```
DBD=DI21PART
  DBD = DBFSAMD1
      DBD =        DBFSAMD2
```

Figure 11. Block analysis control statements

Block analysis report formats

The block analysis utility reads the trace files that you supply and determines each database's block access requests as the access requests are recorded in the trace files. When the block analysis utility

runs, it creates the reports that document the input trace files and the reports that document block requests.

Samples of the Block analysis report formats and their descriptions are provided in the following sections.

SYSPRINT report

The SYSPRINT output contains a listing of any SYSIN control statements along with any error messages that might have resulted from the control statements or from processing the trace files.

An example of SYSPRINT output is shown as follows:

```
PAGE        1                         IMS BUFFER POOL ANALYZER 1.4.0
DATABASE BLOCK ANALYSIS REPORT                      BLOCK ANALYZER

USER SYSIN SPECIFICAITONS;
        DBD = ITEMMASG
        DBD = PRODCNTG
```

Figure 12. Sample report: SYSPRINT output

GTF trace report

The GTFRPT output file documents the IMS Buffer Pool Analyzer trace files that were supplied as input to the Block analysis report.

Note: Any DD name in the report JCL that begins with TRACE is assumed to be a trace file. Also, you cannot concatenate trace files in the JCL. You must use different DD names for each trace file.

An example of the GTF trace report is shown in the following figure. A separate GTF trace report is generated for each trace file that is included in the JCL for the Block Analysis report job.

```
PAGE        1                               IMS BUFFER POOL ANALYZER 1.4.0
GTF TRACE REPORTS                           IMS IMB1 BLOCK ANALYZER

GTF TRACE DATA FOR INPUT DATASET DDNAME=TRACES  DSN=P390M.BPL.TRACE.SVL5

                      NUMBER                                      ELAPSED
   GTF RECORD TYPE    RECORDS     START TIME         END TIME     SECONDS

   CONTROL              220   2021.074 14:26:12   2021.074  14.30.08    236.1
   BFR POOL ANAL      4,157   2021.074 14.26.14   2021.074  14.30.08    233.9
   OSAM                   0
   VSAM              52,891   2021.074 14.27.37   2021.074  14.29.12     95.8
     GET             50,269
     RPLREQ=0F        2,522
   UNKNOWN TYPE           0
   LOST RECORDS           0
   LOST BLOCKS            0
   TOTAL             57,268
```

Figure 13. Sample report: GTF trace

Database block analysis report

The Database block analysis reports are written to the DBDRPT output file. Database block analysis reports are generated only if they are requested using the DBD control statement in the SYSIN DD.

The database reports are produced for each database data set. For each data set, IMS Buffer Pool Analyzer documents data set information. In addition, information about the database data set is documented based on trace information, such as the subpool that was used for the database data set, the hit ratio, the number of buffer requests in the trace, and the number of unique blocks accessed during the trace interval. The report then shows the 200 most frequently accessed blocks in the database data set, including the block number, the number of times that block number was requested, and the number of times the block was found in the buffer pool (the number of hits). Note that the block number, the number of requests, and the number of hits are all shown as hexadecimal numbers.

An example of a Database block analysis report is shown as follows:

```
Page          1                      IMS BUFFER POOL ANALYZER 1.4.0
DATABASE BLOCK ANALYSIS REPORT             IMS IMB1 BLOCK ANALYZER

DATA SET BLOCK ANALYSIS

DATABASE DATA SET INFORMATION:
    DBDNAME ITEMMASG
    DCB NUMBER
    DDNAME IMMSTRIG
    DSNAME DSFFT.VDBSG.IMMSTR01.DATA
    DATA COMPONENT

TRACK INFORMATION:
    SUBPOOL: VSAM          4096 VDEF
    HIT RATIO:  52.3%
    NUMBER OF REQUESTS:        2,503
    NUMBER OF UNIQUE BLOCKS          956

MOST FREQUENTLY ACCESSED BLOCKS:
    BLOCK#      REQUESTS    HITS
    0000008A    00000024  00000021
    00000010    00000019  00000017
    000000F9    00000019  00000016
    000000A5    00000018  00000016
    00000107    00000016  00000014
    000000E2    00000015  00000012
    0000012A    00000015  00000012
    0000017B    00000015  00000013
    00000175    00000013  00000010
    00000103    00000013  00000011
    000000B1    00000013  00000011
    00000049    00000011  0000000F
    00000056    00000010  0000000D
    0000005B    00000010  0000000E
    00000031    00000010  0000000F
    00000143    00000010  0000000D
    00000005    0000000F  0000000D
```

Figure 14. Sample report: Database block analysis

Subpool block analysis report

The Subpool block analysis reports are written to the SPLRPT output file. Subpool block analysis reports are generated for each subpool that was used during the trace interval.

For each subpool, IMS Buffer Pool Analyzer documents subpool information including the subpool type and size, subpool ID, and hit ratio. In addition, the number of subpool requests and the number of unique blocks accessed during the trace interval are shown. The report then shows the 200 most frequently accessed blocks in the subpool, including the DBD name, DCB number or partition ID, the block number, the number of times that block number was requested, and the number of times the block was found in the buffer pool (the number of hits). Note that the block number, the number of requests, and the number of hits are all shown as hexadecimal numbers.

An example of a Subpool block analysis report is shown as follows:

```
PAGE         1                         IMS BUFFER POOL ANALYZER 1.4.0
SUBPOOL BLOCK ANALYSIS REPORT                         IMS IMB1 BLOCK ANALYZER

SUBPOOL BLOCK ANALYSIS

SUBPOOL INFORMATION:
   POOL TYPE VSAM
   BUFFER SIZE VSAR
   COMPONENTS: DATA AND INDEX
   SUBPOOL ID VDEF
   HIT RAIO: 100.0%
   NUMBER OF REQUESTS:        7.744
   NUMBER OF UNIQUE BLOCKS            28

MOST FREQUENTLY ACCESSED BLOCKS IN THIS SUBPOOL:
   DATABASE  DCB #    BLOCKER    REQUESTS     HITS

   PAINDEXG   1     00000002    000004D9   000004D9
   PAINDEXG   1     00000000    000004D9   000004D9
   PAINDEXH   1     00000002    0000046E   0000046E
   PAINDEXH   1     00000000    0000046E   0000046E
   IMINDEXG   1     00000000    00000200   00000200
   IMINDEXH   1     00000000    000001EA   000001EA
   IAINDEXH   1     00000000    0000011A   0000011A
   IAINDEXG   1     00000000    0000011A   0000011A
   CUINDEXG   1     00000000    000000DB   000000DB
   CUINDEXG   1     00000002    000000DB   000000DB
   CUINDEXH   1     00000000    000000CD   000000CD
   CUINDEXH   1     00000000    000000CD   000000CD
   PRODCNTH   1     00000002    00000094   00000094
   PRODCNTG   1     00000002    0000008E   0000008E
   PRODCNTH   1     00000003    00000022   00000022
   PRODCNTH   1     00000005    0000001F   0000001F
   PRODCNTG   1     00000000    0000001A   0000001A
   PRODCNTG   1     00000003    00000019   00000019
   PRODCNTH   1     00000004    00000018   00000018
   PRODCNTG   1     00000006    00000016   00000016
   PRODCNTG   1     00000001    00000016   00000016
```

Figure 15. Sample report: Subpool block analysis report

Chapter 5. IMS buffer pool reports

IMS Buffer Pool Analysis provides several reports to help you analyze IMS database buffer pools in order to provide statistical analysis of the impact of changes affecting buffer pools.

The IMS Buffer Pool Analyzer analyzes I/O rates and buffering information for a specific database to facilitate buffer pool changes. These I/O rates and buffering information might be required for changes to a database structure and to develop models to assist you with making informed decisions on the addition of buffers to an existing pool, or sizing requirements for a new buffer pool.

Topics:

IMS buffer pool reports overview

The summary reports that are included in the IMS buffer pool report provide recommendations for changes to your database buffer pool environment.

The IMS buffer pool reports include the following information:

- A list of databases that are not using the most efficient buffer size because an eligible subpool is not defined with the most efficient buffer size. For example, the report shows a DBDS with a block size of 2K which is currently using a 4K buffer pool, this leaves 50% of the space in each buffer unused. The report shows the amount of storage that could be saved if a buffer pool with the most efficient size was defined, based on the average number of buffers that are present in the subpool.
- Recommendations for the number of buffers to be allocated to each buffer subpool. These recommendations result from examining each subpool to identify where added storage will provide the greatest reduction in I/O, and where reducing storage results in the least increase in I/O. This technique takes storage from the subpools that use it least efficiently, and gives additional storage to the subpools where the storage provides the most benefit.
- Detailed information regarding each buffer subpool used in your environment. You can use these reports to help you understand the reasons for and the details of the recommendations sets in the summary reports. The reports include the following information:
  - Performance information about the current subpool configuration
  - Projections for subpool performance with different numbers of buffers allocated for the subpool
  - A list of each DBDS that used buffers in the subpool during the data-gathering period
  - Performance information for each DBDS

You can create reports for both the buffer pool configuration that was in use when the data-gathering process was performed, and for an updated buffer pool configuration that you define. Based on reports that you create, you can predict the impact of adding a new buffer subpool and allocating several DBDSs to that subpool. IMS Buffer Pool Analyzer can then examine your new environment and suggest optimal storage allocation for that configuration.

Based on the IMS buffer pool report, you can model a change to your DBDS environment by specifying that one or more data sets has a different block size, CI size, or access method from that which was used at the time that the data was gathered. This is helpful when designing a new buffer pool environment where a major DBDS configuration change will be made.

Include a statement in your report job that specifies the access method, either OSAM or VSAM, and the new block size or CI size to be used for each data set. IMS Buffer Pool Analyzer creates a model environment with changes to your data sets. When the block size or CI size is changed, IMS Buffer Pool Analyzer adjusts the individual buffer requests to model the new block number that is required to fulfill the buffer request.

For example, if the block size of a data set is increased from 1024 to 4096, four times the amount of data will be included in each data set buffer, and IMS Buffer Pool Analyzer will convert each buffer request from a 1K block size to a 4K block size, possibly reducing the number of I/Os that are required for consecutive blocks that are retrieved.

Generating buffer pool reports

You can generate a database buffer pool report by running a batch buffer pool report job. Sample JCL, which is shown in the following figure, is included in the SBPLSAMP library in member BPLMAIN.

```
//BPLMAIN  JOB
//*
//* This job runs the IMS Buffer Pool Analyzer buffer pool report.
//*
//  SET  SBPLLINK=BPL140.SBPLLINK    SBPLLINK DSN
//  SET  TRACE=BPL.TRACE             TRACE DSN
//  SET  ITKBSRVR=                   OPTIONAL ITKB SERVER NAME
//*
//* Any number of GTF trace data sets can be included in each
//* run of the reporting utility. Each GTF trace data set must
//* be allocated to a DD statement that begins with DDNAME TRACE. If
//* there were three GTF trace data sets, a report run could
//* include DDNAMES TRACE1, TRACE2, and TRACEXYZ.
//*
//BPL      EXEC PGM=BPLMAIN,REGION=64M,PARM='ITKBSRVR=&ITKBSRVR'
//STEPLIB  DD   DSN=&SBPLLINK,DISP=SHR
//SYSPRINT DD   SYSOUT=*
//GTFRPT   DD   SYSOUT=*
//SPLRPT   DD   SYSOUT=*
//MDLRPT   DD   SYSOUT=*
//DBDRPT   DD   SYSOUT=*
//DSNRPT   DD   SYSOUT=*
//SYSUDUMP DD   SYSOUT=*
//TRACE    DD   DSN=&TRACE,DISP=SHR
//DVSVSAMP DD   DUMMY
//SYSIN    DD   DUMMY
```

Figure 16. Sample JCL for running the batch buffer pool report job

The following symbolic parameters are defined in the sample job:

SBPLLINK

The **SBPLLINK** symbolic parameter sets the data set name of the IMS Buffer Pool Analyzer load library used in the STEPLIB DD.

TRACE

The **TRACE** symbolic parameter can be used to set the data set name of one trace data set. Since multiple data sets can be used as input to a report, other trace data sets should be added to the JCL using additional DD names that start with TRACE. Note that trace data sets cannot be concatenated to a single DD name.

For other considerations, see "Combining multiple traces into a single report" on page 19.

ITKBSRVR

The **ITKBSRVR** parameter can optionally be used to set the XCF group name for the IMS Tools Knowledge Base servers on the LPAR where the report job will run. If IMS Tools Knowledge Base is not used in your installation, or if you do not want to write IMS Buffer Pool Analyzer reports to the Knowledge Base repository, leave this symbolic parameter blank.

Following are the DD names that are included in the JCL for the Buffer Pool report:

STEPLIB

Required DD statement that must refer to a data set that contains the IMS Buffer Pool Analyzer load modules. The library does not require APF authorization for any batch report jobs.

SYSIN

Optional DD that you can use to specify reporting options. SYSIN control statements are described next.

SYSPRINT

Print output file for error messages and SYSIN control statements This is a required DD statement. The DCB information for this DD, which need not be specified, is RECFM=FBA,LRECL=133.

GTFRPT

Print output file for the GTF trace report. This is an optional DD statement; if this DD statement is not present, the GTF trace report is not produced. DCB information need not be specified, but it is RECFM=FBA,LRECL=133.DSORG=PS.

SPLRPT

Print output file for the Database Buffer Subpool reports. This is an optional DD statement; if this DD statement is not present, the database buffer subpool report is not produced. DCB information need not be specified, but it is RECFM=FBA,LRECL=133.DSORG=PS.

MDLRPT

Print output file for the model subpool reports if an alternate configuration was requested in a DFSVSAMP DD statement or a CHANGE DBD statement in the SYSIN DD. This is an optional DD statement; if this DD statement is not present, the model subpool report is not produced. DCB information need not be specified, but it is RECFM=FBA,LRECL=133.DSORG=PS.

DBDRPT

Print output file for the Database Information report. You can disable this report by specifying the SYSIN statement DBDRPT=NO. This is an optional DD statement; if this DD statement is not present, the database information report is not produced. DCB information need not be specified, but it is RECFM=FBA,LRECL=133.DSORG=PS.

DSNRPT

Print output file for the Data Set Information report. You can disable this report by specifying the SYSIN statement DSNRPT=NO. This is an optional DD statement; if this DD statement is not present, the Data Set Information report is not produced. DCB information need not be specified, but it is RECFM=FBA,LRECL=133.DSORG=PS.

SYSUDUMP

Diagnostic information. If an abend occurs, this information can be useful for troubleshooting the cause.

TRACE*xxx*

Any number of TRACE*xxx* DD names can be included in the report job. The last three characters of the DD name can be any valid characters that are allowed in JCL. Each DD name must have only one GTF trace data set. All data sets are read and combined for reporting purposes. At least one TRACE*xxx* DD statement is required.

For other considerations, see "Combining multiple traces into a single report" on page 19.

DFSVSAMP

Optional DD that you can use to specify an alternate configuration to model. When this DD is specified, this data set must contain IMS buffer pool definition statements as described in the *IMS System Definition*.

Customizing buffer pool reports using SYSIN statements

By including statements in the SYSIN DD of a buffer pool report job, you can customize the output that is produced. The SYSIN DD is optional, as are the statements that are documented in this section.

Most SYSIN statements cannot be continued to multiple lines. The exception is the BUFNO statement, which can be continued on as many lines as are necessary. To continue the statement, leave a comma as the last non-blank character on the statement line in one of columns 1 through 72.

The following statements can be specified:

BUFNO type=(*number,number,...*)

The BUFNO statement allows you to customize the subpool reports that show performance metrics for different numbers of buffers. The BUFNO statement allows you to replace the default number of buffers with numbers that you choose.

The *type=* parameter in the BUFNO statement specifies the type of subpools for which the numbers of buffers will be used. You can specify different numbers of buffers for OSAM, VSAM, VSAM with hiperspace, and DEDB subpools. You can specify the following options in each BUFNO statement:

- OSAM
- VSAM
- HIPERSPACE= (or HS=)
- DEDB

In addition to the *type=* parameter, you can specify up to 45 different numbers of buffers. Each number can range from 3 to 32768 for OSAM and VSAM, or from 3 to 999999 for Hiperspace or DEDB storage formats. The default number of buffers are as follows:

```
BUFNO OSAM=(4,8,16,32,64,128,192,256,384,512,768,1024,1536,
            2048,3072,4096,6144,8192,12288,16384,20480,24576,
            28672,32768)
```

```
BUFNO VSAM=(4,8,16,32,64,128,192,256,384,512,768,1024,1536,
            2048,3072,4096,6144,8192,12288,16384,20480,24576,
            28672,32768)
```

```
BUFNO HIPERSPACE=(4,8,16,32,64,128,192,256,384,512,768,1024,
                  1536,2048,3072,4096,6144,8192,12288,16384,
                  20480,24576,28672,32768,40960,49152,57344,
                  65535,81920,98302,131072,163840,196608,262144)
```

```
BUFNO DEDB=(4,8,16,32,64,128,192,256,384,512,768,1024,1536,
            2048,3072,4096,6144,8192,12288,16384,20480,24576,
            28672,32768)
```

IMS Buffer Pool Analyzer automatically adds entries for the actual number of buffers in the subpool, and if DFSVSAMP input is included, the number of buffers that are specified for that subpool.

You can specify all four types of BUFNO statements in the SYSIN DD, or specify only those types that you choose. For example, you could include only OSAM= and VSAM= statements and allow hiperspace and DEDB values to default to the values that are shown above.

BUFNO statements are the only SYSIN statements that can be continued. Only columns 1 through 72 are treated as data. To continue a BUFNO statement, ensure that the last non-blank characters within columns 1 through 72 is a comma. A complete BUFNO statement is identified by a closed parenthesis.

Tip:

- IMS Buffer Pool Analyzer uses only the values that you specify when it evaluates whether increasing or decreasing the number of buffers makes a subpool more or less efficient. The accuracy of the summary subpool reports depends on whether there are enough valid entries and sizes, and whether they span a sufficiently large number of subpools.
- The largest number for each type determines the number of buffer slots that are tracked during processing of the Buffer pool report job. Having a very large number can significantly increase the job's processing time.

The following samples show valid BUFNO statement examples:

```
//SYSIN  DD *
BUFNO OSAM=(3,32,100,200,300,400,500,600,700,800,900,1000,2000,
            3000,4000,5000,6000,7000,8000,9000,10000,20000,
            30000)
BUFNO VSAM=(500,400,300,200,100,4,600,700,800,900,1000,2000,
```

```
               3000,4000,5000,6000,7000,8000,9000,10000,20000,
               30000)
 BUFNO DEDB = ( 5, 100, 200, 300,400,500,600,700,800,900,1000,
           2000,3000,4000,5000,6000,7000,8000,9000,
           10000,20000,30000,40000,50000
               30000)
```

CHANGE DBD statements

The following two CHANGE statements can be used:

```
CHANGE DBD=(dbdname,dsid),AM=OSAM,BLKSIZE=size

CHANGE DBD=(dbdname,dsid),AM=VSAM,DATACISIZE=size,INDEXCISIZE=size
```

The CHANGE statements allow you to model the impact of changing the access method, and block size or CI size of a DBDS. When a CHANGE statement is included in the SYSIN statements, IMS Buffer Pool Analyzer produces both a subpool report (in the SPLRPT DD) and a model subpool report in the MDLRPT DD. The SPLRPT subpool report reflects the actual DBDS characteristics. The MDLRPT DD subpool report reflects the impact of changing the DBDS access method, block size, or CI size.

The CHANGE statement keywords are as follows:

DBD

This keyword specifies the database name and, optionally, the data set number. For HALDB, specify the partition name and, optionally, the partition data set number (typically A through J, L, and X).

For example, DBD=DI21PART will cause all data sets that are associated with the DI21PART database to be changed. Specifying DBD=(DI21PART,1) will cause only the first data set to be changed.

AM

This keyword specifies the access method for the new database data sets. You can specify either VSAM or OSAM.

BLKSIZE

This keyword is required for AM=OSAM and is invalid for AM=VSAM. It specifies the new block size that the specified data sets will have in the model reports MDLRPT DD.

DATACISIZE

This keyword is required for AM=VSAM and is invalid for AM=OSAM. It specifies the new data CI size that the specified data sets will have in the model reports MDLRPT DD.

INDEXCISIZE

This keyword is optional for AM=VSAM and is invalid for AM=OSAM. It specifies the new index CI size that the specified data sets will have in the model reports MDLRPT DD.

DBDRPT=YES | NO

The DBDRPT statement allows you to include or exclude the creation of the Buffer pool database report. Include the creation of the Buffer pool database report in the DBDRPT statement by specifying YES. Or exclude the creation of the Buffer pool database report in the DBDRPT statement by specifying NO. The default is DBDRPT=YES. If you specify DBDRPT=NO, the DBDRPT DD statement in the Buffer pool report JCL is not required.

DSNRPT=YES | NO

The DSNRPT statement allows you to include or exclude the creation of the Buffer pool data set report. Include the creation of the Buffer pool data set report in the DSNRPT statement by specifying YES. Or exclude the creation of the Buffer pool data set report by specifying NO. The default is DSNRPT=YES. If you specify DSNRPT=NO, the DSNRPT DD statement in the Buffer pool report JCL is not required.

VARYPCT=*nnn*

The VARYPCT statement allows you to specify the criteria for choosing subpools that contain databases with dispersed I/O hit ratios.

The value specified in the VARYPCT statement indicates the percentage of buffer requests whose I/O hit ratios should be concentrated in at least one of 20% ranges (for example, between 20% and 40%). The value can be between 0 and 100, and the default value is 80.

Suppose you specified VARYPCT=70. IMS Buffer Pool Analyzer checks each subpool and examines whether the number of buffer requests that have similar I/O hit ratios (hit ratios within every 20% range; for example, 40 - 60% or 30 - 50%) accounts for at least 70% of the total buffer requests. If this condition is not satisfied, that is, if I/O hit ratios of buffer requests are so dispersed that none of 20% ranges of I/O hit ratios constitutes at least 70% of the total buffer requests, the subpool will be reported in a varied hit ratio report.

Therefore, the higher the value specified, the more likely that a subpool will be reported in the varied hit ratio report.

Buffer pool reports print DD names

IMS Buffer Pool Analyzer's Buffer pool report produces a multitude of reports, each of which is written to a different print DD.

The following DD names are used for buffer pool report output. The description of each DD name is presented in the following list.

SYSPRINT

The SYSPRINT DD includes diagnostic messages that describe any error conditions that were encountered during processing, along with a list of the statements that were supplied in the SYSIN DD to control the options that were used to generate the report.

GTFRPT

The GTFRPT DD includes the GTF trace report that describes the GTF trace data sets that are supplied as input to this buffer pool report.

SPLRPT

The SPLRPT DD contains the database buffer subpool reports. These reports document the current IMS buffer pool environment. If a DFSVSAMP DD was used to specify an alternate environment, the MDLRPT DD will contain reports documenting the alternate environment. The database subpool reports include summary subpool information, recommendations for improving buffer pool performance, and detailed information about each subpool in the configuration.

MDLRPT

The MDLRPT DD contains database subpool reports for the alternate configuration that was requested in the DFSVSAMP DD. The database subpool reports include summary subpool information, recommendations for improving buffer pool performance, and detailed information about each subpool in the configuration.

DBDRPT

The DBDRPT DD contains the database report. This optional report can be suppressed by including DBDRPT=NO in the SYSIN statements. When it is produced, the database report provides information about the databases that were used by IMS during the trace intervals. This includes some statistical information such as hit ratio and activity rates.

Note: This information is also provided in the subpool reports; each subpool detailed report includes a list of all the databases that were used in that subpool.

DSNRPT

The DSNRPT DD contains the data set reports for OSAM and VSAM data sets. These reports provide a listing of the databases that were used during the trace interval, their DD names and data set names. This optional report can be suppressed by including DSNRPT=NO in the SYSIN.

Modeling an alternate buffer pool configuration

IMS Buffer Pool Analyzer allows you to make buffer pool configuration alterations and to produce the subpool reports that are based on these alternate configurations.

You can add, remove, or change subpool definitions, change the association of DBDSs with subpools, and change the block or CI size of a data set.

Changing the access method or block size of a database

IMS Buffer Pool Analyzer allows you to model the impact of a change to the access method, either OSAM or VSAM, and the block size or CI size of a DBDS. If you plan to make a database change, IMS Buffer Pool Analyzer allows you to review the impact of the change on your buffer pool configuration and to plan changes to your configuration before the actual database change is implemented.

To model the impact of a database change, supply one or more **CHANGE DBD** statements in the SYSIN DD of the BPLMAIN report job. IMS Buffer Pool Analyzer will produce reports for the trace data that was supplied that documents the performance of the old database configuration in the SPLRPT DD and the performance of the new database configuration in the MDLRPT DD. This allows you to compare buffer pool performance before and after the change, and also to review suggestions for changes to your buffer pool configuration that would affect changes you choose to make in your database configuration.

You can also make changes to your buffer pool configuration, as described in the next section, which IMS Buffer Pool Analyzer combines with changes to your database configuration to produce the model subpool reports.

Changing the IMS buffer pool configuration

IMS Buffer Pool Analyzer also allows you to specify a new buffer pool configuration that will be used along with existing trace data to determine the impact of a buffer pool configuration change. Specify an updated buffer pool configuration by adding a DFSVSAMP DD statement to the BPLMAIN report JCL, and specify IMS buffer pool definition statements using this DD name. IMS Buffer Pool Analyzer supports all of the valid control statements that can be specified in DFSVSM*xx* as described in the *IMS Installation*. Statements that are unrelated to buffer pool definitions and database assignments to buffer pools are ignored. Ignoring those statements and assignments allows you to copy all existing statements from your current DFSVSAMP definitions for the IMS control region or batch job, make changes to the buffer pool definitions, validate performance using the buffer pool report job, and copy the updated DFSVSAMP control statements back for use by the IMS control region or batch job.

Alternate configuration reports

IMS Buffer Pool Analyzer produces a second set of subpool statistics reports for the new configuration. All reports that are written to the SPLRPT DD for the configuration that was in place when the trace was taken, are produced for the alternate configuration and are written to the MDLRPT DD. In addition to the subpool statistics reports, IMS Buffer Pool Analyzer also lists the DFSVSAMP statements that were used to create the model configuration. The DFSVSAMP information is listed before the subpool statistics reports in the MDLRPT. An example of the DFSVSAMP statement information is shown in the following figure:

```
PAGE    1                                 IMS BUFFER POOL ANALYZER 1.4.0
DFSVSAMP MODEL SUBPOOL INFORMATION                       BUFFER POOL REPORT

USER DFSVSAMP SPECIFICATION
       VSRBF=4096,2000
       IOBF=(8192,500,Y,Y)
       SBONLINE,MAXSB=10
       OPTIONS,BGWRT=YES,INSERT=SKP,DUMP=YES,DUMPIO=YES
       OPTIONS,VSAMFIX=(BFR,IOB),VSAMPLS=LOCL
       OPTIONS,DL/I=ON,LOCK=ON,DISP=ON,SCHD=ON,DLOG=ON,LATC=ON,SUBS=ON
       OPTIONS,STRG=ON
       OLDSDEF OLDS=(00,01,02,03,04,05,99),BUFNO=005,MODE=SINGLE
       WADSDEF WADS=(0,1,8,9)
```

Figure 17. User DFSVSAMP model subpool specification information

Following this report, all of the subpool statistics reports are written to the MDLRPT DD. This report listing includes the subpool summary, all of the IMS Buffer Pool Analyzer recommendation reports, and subpool detail reports. All of these reports reflect the changes in the configuration, including any CHANGE DBD statements along with updates to the IMS buffer pool definitions.

Viewing IMS Buffer Pool Analyzer reports

IMS Buffer Pool Analyzer produces several reports. The reports are written to several SYSOUT files that can also be archived and viewed using IMS Tools Knowledge Base, if it is installed and available at your installation.

When accessing IMS Buffer Pool Analyzer reports through IMS Tools Knowledge Base, you will select a report job and be presented with a selection menu for the reports that are available for that report job. The reports will be named the same as the DD names that are documented in the preceding section. Some report jobs might have fewer reports, because there might not be any model reports, or the DBD or DSN reports could have been disabled with control cards.

```
SERVER: ITKBSRVR              Available Reports - Job              Row 1 from 5
Command ===>                                                  Scroll ===> PAGE

Select a row action or press END to exit.

Row actions: S - View J - Job T - Step H - History P - Print D - Delete
             I - Info A - Archive
Act Product        Report                      DBD      Area/Part DD      Set
    IMS BPA        DBDRPT
    IMS BPA        DSNRPT
    IMS BPA        GTFRPT
    IMS BPA        MDLRPT
    IMS BPA        SPLRPT
****************************** Bottom of data ********************************
```

Figure 18. A typical list of available reports for a job

Viewing a report in IMS Tools Knowledge Base looks exactly the same as viewing the reports that are directed to SYSOUT files. All the messages and report fields are exactly as they were described in this topic.

GTF trace report

The GTF trace report shows timestamps and the number of records for each GTF trace data set that is included in the report job.

The following example shows a GTF trace report that includes two input GTF trace data sets. For each trace data set, the start and end times of the trace data and the elapsed time of each trace are shown. The number of records of each type are also shown, including breakdowns for OSAM, VSAM, and DEDB record, by type, of GTF trace record.

In addition, the GTF trace report provides information on lost records and lost blocks. Lost records occur when GTF buffers fill, and no buffer space is available for additional trace records. Lost blocks typically

occur during write errors to the trace data set or during periods of very high trace activity. After trace buffers become available, GTF writes an error record that provides information on the number of trace records that were discarded for lack of buffer space. The report shows both the number of GTF error records and the total number of trace records that were discarded.

```
PAGE     1   IMS BUFFER POOL ANALYZER 1.4.0    DATE: 08/30/2021
GTF TRACE REPORT                 IMS IMB1             TIME: 19:35:30

GTF TRACE DATA FOR INPUT DATASET DDNAME=TRACE   DSN=P390M.BPL.TRACE.SVL3

                NUMBER                                       ELAPSED
GTF RECORD TYPE RECORDS  START TIME        END TIME          SECONDS
--------------- -------- ---------------   ----------------- ---------
CONTROL           6,001  2021.074 20:03:12 2021.074 20:04:43   91.0
BFR POOL ANAL     1,909  2021.074 20:03:14 2021.074 20:03:14    0.0
OSAM             1909.5K 2021.074 20:03:12 2021.074 20:04:43   91.0
   BLOCK LOCATE  363.4K
   BYTE  LOCATE 1541.2K
   CREATE BLOCK  4,810

VSAM              329.8K 2021.074 20:03:12  2021.074 20:04:43   91.0
    GET           329.8K
UNKNOWN TYPE           0
LOST RECORDS           0
LOST BLOCKS            0
TOTAL            2247.2K

GTF TRACE DATA FOR INPUT DATASET DDNAME=TRACE2 DSN=P390M.BPL.TRACE

                NUMBER                                       ELAPSED
GTF RECORD TYPE RECORDS  START TIME        END TIME          SECONDS
--------------- -------- ---------------   ----------------- ---------
CONTROL              2   2021.198 16:42:14 2021.198 16:45:41  206.4
BFR POOL ANAL      141   2021.198 16:42:16 2021.198 16:45:41  205.2
OSAM                 0
VSAM                 0
UNKNOWN TYPE         0
LOST RECORDS         0
LOST BLOCKS         0
TOTAL              145
```

Figure 19. GTF trace report

The reference information that is provided in the GTF trace report includes the following fields:

CONTROL
A GTF record type that shows the number of control records that were produced by GTF that are not used by IMS Buffer Pool Analyzer.

BFR POOL ANAL
These are records are generated by IMS Buffer Pool Analyzer. Records might include information on subpool definitions, DBDS information, control records, or DEDB lookaside buffer statistics.

OSAM
Shows information about the OSAM trace records that are produced by IMS. OSAM GTF trace record subtypes are also listed as appropriate. The OSAM subtype can be an OSAM request type, **DIFFERNT IMS**, or **UNKNOWN**.

OSAM request type
Indicates the number of records for each OSAM request type. The hexadecimal OSAM-related buffer handler function code that is set in the record is converted to a meaningful text string for this report. See the following table for possible values:

Hexadecimal value	Subtype
E1	BLOCK LOCATE
E2	BYTE LOCATE

Hexadecimal value	Subtype
EB	BYTLOC+MARK (This subtype represents "Byte locate and mark the buffer as altered")
E4	CREATE BLOCK

DIFFERNT IMS
> The number of OSAM records from other subsystems.

UNKNOWN
> The number of records other than those listed under one of OSAM request types or **DIFFERNT IMS**.

VSAM
> Shows information about the VSAM trace records that are produced by VSAM. VSAM GTF trace record subtypes are also listed as appropriate. The VSAM subtype can be a VSAM request type, **NOT IMS DB**, or **UNKNOWN**.

VSAM request type
> Indicates the number of records for each VSAM request type. The hexadecimal RPLREQ (RPL request type) value that is set in the record is converted to a meaningful text string for this report. See the following table for possible values:

Hexadecimal value	Subtype
00	GET
01	PUT
02	CHECK
03	POINT
04	ENDREQ
05	ERASE
06	VERIFY
07	IMPORT/JESFM ("JESFM" stands for "JES FORMAT")
08	DATA PREFMT ("PREFMT" stands for "PREFORMAT")
09	INDEX PREFMT ("PREFMT" stands for "PREFORMAT")
0A	FORCE I/O
0B	RPLREQ=0B
0C	RPLREQ=0C
0D	RPLREQ=0D
0E	RPLREQ=0E
0F	RPLREQ=0F
10	CNVTAD
11	MNTACQ
12	ACQRANGE
13	TERMRPL
14	VERIFY REFR ("REFR" stands for "REFRESH")
15	LKADD

Hexadecimal value	Subtype
16	LKREQ
17	LKCD
18	INQRC
19	RETLK
1A	RECOV REQ
1B	QUIESE REQ
1C	EADD REQ

NOT IMS DB
The number of VSAM buffer activities (GTF F61 records) that were not VSAM traces for the IMS full-function database.

UNKNOWN
The number of records other than those listed under one of VSAM request types or **NOT IMS DB**.

UNKNOWN TYPE
GTF trace records with an event ID that does not match any record types that IMS Buffer Pool Analyzer processes (440 - Buffer Pool Analyzer records, F4F - IMS OSAM records, and F61 - VSAM records). There should be no records in this category. If the report indicates the presence of records, investigate the event ID that is associated with the records, determine the source of the record type, and disable recording the records. For assistance, contact IBM Software Support.

LOST RECORDS
GTF logs the number of records that were discarded because of insufficient buffer space. The number of records shows how many occurrences of lost records were recorded. When present, the actual number of records that were discarded is also listed.

LOST BLOCKS
GTF logs the number of blocks that were discarded because of insufficient buffer space. The number of records shows how many occurrences of lost blocks were recorded. When present, the actual number of blocks that were discarded is also listed.

Reducing GTF lost records

If you notice a significant number of lost GTF trace records or blocks in a GTF trace report generated by a report job, you can make changes to the BPLGTF PROC that can help reduce or eliminate the lost information.

Increasing the **BLOK**, **NCP**, and **SIZE** parameters can help because doing so increases the amount of space that is available to GTF for records before they are written to the trace data set. Remember that the BLOK= PROC symbolic value specifies the amount of ECSA that GTF will use for GTF record storage. Use care when increasing this value because shortages of ECSA can cause significant problems for z/OS.

The SIZE= symbolic value specifies the amount of private storage in the BPLGTF address space that will be used for output buffers. Increasing this value will not impact other jobs that are running on the system. You can also increase the NCP= symbolic value, which will allow more concurrent I/O requests for the trace data set.

If your trace data set is allocated to DASD, try moving it to tape. While tape can be a slower medium, buffering capabilities in new hardware and chained writes that are issued by GTF can actually improve the performance of the trace data set, especially if the data set is allocated on an active volume.

Data set information reports

The Data set information report shows which data sets are associated with the performance reports.

There are two reports: one for OSAM DBDSs and one for VSAM DBDS components.

Data set information reports are optional. These reports can be suppressed by supplying a statement in the SYSIN DD with DSNRPT=NO.

Data set information reports provide information about data set names that could be referenced by the batch job or subsystem. Only data sets that could be used by the job are presented:

- For batch jobs, this includes every DBD in the PSB that has a database data set that is either present in the job's JCL or available for dynamic allocation by DFSMDA members or from information in the RECON data sets for Fast Path and HALDB databases.
- For online systems, the list includes data set names for every DBDS that has a valid ACBLIB member; including data set names that are available from the JCL or DFSMDA members or those that are available from RECON definitions.

The following data set information report is an example. Information about block or CI size and buffer size is present only for data sets that were active during the trace interval. Note that three types of reports can be generated – one for OSAM data sets, one for VSAM data sets, and one for DEDB DBDSs.

```
PAGE  1    IMS BUFFER POOL ANALYZER 1.4.0  (H28A140)  DATE: 08/30/2021
OSAM DATASET INFORMATION          IMS IMB1                   TIME:  19:35:30

         DCB  DB    MASTER    SUB-  BUFFER  BLOCK
DBDNAME  NUM  TYPE  HALDBD    POOL  SIZE    SIZE    DDNAME    DSNAME
_____ ___  ____  _____  ____  _____  _____  _____   _____
PARTSDBC  1   DL/I                                   PARTSDBC  DSFFT.DBSC.PARTSDD
PARTSDBD  1   DL/I                                   PARTSDBD  DSFFT.DBSD.PARTSDD
PARTSDBE  1   DL/I                                   PARTSDBE  DSFFT.DBSE.PARTSDD
PARTSDBF  1   DL/I                                   PARTSDBF  DSFFT.DBSF.PARTSDD
PDB1001   A   PART  CUSTOMRA  O08K  8,192   8,192   PDB1001A  BPA1.DBD1.A00001
PDB1001   B   PART  CUSTOMRA                         PDB1001B  BPA1.DBD1.B00001
PDB1001   C   PART  CUSTOMRA                         PDB1001C  BPA1.DBD1.C00001
PDB1001   D   PART  CUSTOMRA                         PDB1001D  BPA1.DBD1.D00001
PDB1001   E   PART  CUSTOMRA                         PDB1001E  BPA1.DBD1.E00001
PDB1002   A   PART  CUSTOMRA  O08K  8,192   8,192   PDB1002A  BPA1.DBD1.A00002
PDB1002   B   PART  CUSTOMRA                         PDB1002B  BPA1.DBD1.B00002
PDB1002   C   PART  CUSTOMRA                         PDB1002C  BPA1.DBD1.C00002
PDB1002   D   PART  CUSTOMRA                         PDB1002D  BPA1.DBD1.D00002
PDB1002   E   PART  CUSTOMRA                         PDB1002E  BPA1.DBD1.E00002

PAGE    21            IMS BUFFER POOL ANALYZER 1.4.0  (H28A140)  DATE: 08/30/2021
VSAM DATASET INFORMATION            IMS IMB1                   TIME:  19:35:30

         DCB  VSAM   DB    MASTER    SUB-  BUFFER  CI
DBDNAME  NUM  COMP   TYPE  HALDB     POOL  SIZE    SIZE  DDNAME    DSNAME
_____  ___  _____  ____  _____    ____  ____    ____  _____    _____
PDB1046   X   INDEX  PART  CUSTOMRA        4,096   4,096 PDB1046X  BPA1.DBD1.X00046X
PDB1046   X   DATA   PART  CUSTOMRA        4,096   4,096 PDB1046X  BPA1.DBD1.X00046D
PDB1047   L   INDEX  PART  CUSTOMRA                      PDB1047L  BPA1.DBD1.L00047X
PDB1047   L   DATA   PART  CUSTOMRA                      PDB1047L  BPA1.DBD1.L00047
PDB1047   X   INDEX  PART  CUSTOMRA        4,096   4,096 PDB1047X  BPA1.DBD1.X00047X
PDB1047   X   DATA   PART  CUSTOMRA        4,096   4,096 PDB1047X  BPA1.DBD1.X00047D

PAGE    43            IMS BUFFER POOL ANALYZER 1.4.0
DEDB DATASET INFORMATION            IMS IMB1 BUFFER POOL REPORT

                       LOOK           BUFFER  CI
DBDNAME   AREA     VSO ASID  SUB POOL SIZE    SIZE   DDNAME    DSNAME
_____  _____ ___ ____  _____ _____  _____ _____   _____
ACCTA     CCAREA01 VSO LKAS  FPMAS1   4,096   4,096  CCARDD1P  BPA1.CCARA01P
ACCTA     CCAREA01 VSO LKAS                          CCARDD1S  BPA1.CCARA01S
ACCTA     CCAREA02 VSO LKAS  FPMAS1   4,096   4,096  CCARDD2P  BPA1.CCARA02P
ACCTA     CCAREA02 VSO LKAS                          CCARDD2S  BPA1.CCARA02S
ACCTA     CCAREA03                                   CCARDD3P  BPA1.CCARA03P
ACCTA     CCAREA03                                   CCARDD3S  BPA1.CCARA03S
STOREDBA  STOREA01 VSO LKAS  FP2STORE  512     512   STOREA1P  BPA1.STORA01P
```

Figure 20. Sample report: Data set information

The reference information that is provided in the OSAM and VSAM data set information reports includes the following fields:

DBDNAME
    The database name.

AREA

The DEDB AREA name.

DCB NUM

The data set number of this data set within the database or the HALDB partition.

VSAM COMP

This information is present only in the VSAM data set information report, this field shows the VSAM data set component. It can be either DATA or INDEX.

DB TYPE

This field indicates the type of database, either: DL/I full function database, PART that is a partition of a HALDB, or DEDB that is a Fast Path data entry database.

MASTER HALDB

This field is blank unless the database is a HALDB partition. For partitioned databases, this field shows the master HALDB database name of which this partition is a member.

VSO

The VSO field is shown only in the DEDB section of the report. The VSO column indicates whether the area is using the DEDB Virtual Storage Option (VSO) or not. The VSO column entries are as follows: If VSO is shown in the VSO column, the area is using VSO. If there are blanks in the VSO column, the area is not using VSO.

LOOK ASID

The Lookaside column is shown only in the DEDB section of the report. This column indicates whether the area is using the DEDB Lookaside option for database buffering. The LOOK ASID column entries are as follows: If LKAS is shown in the LOOK ASID column, the area is using the Lookaside option. If there are blanks in the LOOK ASID column, the area is not using the Lookaside option.

SUBPOOL

The subpool indicates the name of the subpool that was used by this DBDS during the trace interval. If the database was modified in the DFSVSAMP statements to use a different subpool, the new subpool name is not indicated in this report.

BUFFER SIZE

The BUFFER SIZE shows the size of the buffers that were used by this database during the trace interval. This does not reflect changes to the assigned buffer pool that might have occurred because of a **CHANGE DBD** command or changes in the buffer pool sizes that were made in DFSVSAMP statements. Also, if the DBDS was not used during the trace interval, this field will be blank.

BLOCK SIZE / CI SIZE

This field shows the size data set block size for OSAM or the CI size for VSAM. This does not reflect updates to the database that were specified in a CHANGE statement that was included in the SYSIN statements.

DDNAME

This field shows the DD name that was used by this DBDS.

DSNAME

This field shows the data set name of the DBDS. For VSAM data sets, this is the data set name of the data or index component.

Database statistics reports

Database statistics reports are listed by database access method, either OSAM or VSAM. The report lists detailed information about the DBDS along with I/O activity rates and hit ratios by DBD name and DCB number.

Database statistics reports are optional. You can suppress them by supplying a statement in the SYSIN DD specifying DBDRPT=NO.

Database statistics reports show actual numbers that are based on the GTF data that was processed.

The following figure shows an example of portions of the OSAM, VSAM, and DEDB database statistics reports.

```
PAGE  1    IMS BUFFER POOL ANALYZER 1.4.0  (H28A140)   DATE: 08/30/2021
OSAM DATA SET INFORMATION          IMS IMB1                    TIME:  19:35:30

         DCB   MASTER   SUB-  BUFFER  BLOCK   HIT   REQUESTS   READS
DBDNAME  NUM   HALDBD   POOL  SIZE    SIZE   RATIO PER SECOND PER SECOND
_____  ___   _____  ____  _____  _____ _____ _____ _____
PDB1001   A    CUSTOMRA 008K  8,192   8,192  93.3%    244.90    16.23
PDB1002   A    CUSTOMRA 008K  8,192   8,192  93.2%    213.47    14.32
PDB1003   A    CUSTOMRA 008K  8,192   8,192  93.3%    240.78    15.90
PDB1004   A    CUSTOMRA 008K  8,192   8,192  93.6%    226.09    14.32
PDB1005   A    CUSTOMRA 008K  8,192   8,192  93.8%    242.64    14.94
PDB1006   A    CUSTOMRA 008K  8,192   8,192  93.6%    252.02    15.91
PDB1007   A    CUSTOMRA 008K  8,192   8,192  92.9%    217.40    15.27

**TOTALS                                     92.9% 20,981.40  1,481.44


PAGE    21          IMS BUFFER POOL ANALYZER 1.4.0  (H28A140)  DATE: 08/30/2021
VSAM DATA SET INFORMATION                 IMS IMB1                 TIME:  19:35:30

         DCB VSAM  MASTER    SUB-  BUFFER CI    HIT   REQUESTS   READS
DBDNAME  NUM COMP  HALDB     POOL  SIZE   SIZE  RATIO PER SECOND PER SECOND
_____  ___ ____  _____   ____  _____ ____  _____ _____ _____
PDB1046   X  INDEX CUSTOMRA  ____  8,192  4,096 100.0%   13.93     0.00
PDB1046   X  DATA  CUSTOMRA        8,192  4,096 100.0%   13.93     0.00
PDB1047   X  INDEX CUSTOMRA        8,192  4,096 100.0%   13.67     0.00
PDB1047   X  DATA  CUSTOMRA        8,192  4,096 100.0%   13.67     0.00
PDB1048   X  INDEX CUSTOMRA        8,192  4,096 100.0%   14.26     0.00
PDB1048   X  DATA  CUSTOMRA  008K  8,192  4,096 100.0%   14.26     0.00
PDB1049   X  INDEX CUSTOMRA  008K  8,192  4,096 100.0%   14.53     0.00
PDB1049   X  DATA  CUSTOMRA  008K  8,192  4,096 100.0%   14.53     0.00

**TOTALS                                        100.0% 3,624.55    0.00

PAGE    4                            IMS BUFFER POOL ANALYZER 1.4.0
DEDB DATABASE STATISTICS                   IMS IMB1 BUFFER POOL REPORT

         VSAM                      BUFFER  CI    HIT   REQUESTS   READS
DBDNAME  COMP    AREA    SUBPOOL   SIZE    SIZE  RATIO PER SECOND PER SECOND
_____  ____  _____   _____  _____  ____  _____ _____ _____
ACCTA    DATA  CCAREA01  FPMAS1    4,096   4,096 14.4%   260.34   222.75
ACCTA    DATA  CCAREA02  FPMAS1    4,096   4,096  0.9%    15.29    15.15
STOREDBA DATA  STOREA01  FP2STORE    512     512 80.2%   306.38    60.52
**TOTALS                                         48.7%   582.01   298.42
```

Figure 21. Sample report: Database statistics

The reference information that is provided in the OSAM and VSAM database information reports includes the following fields:

DBDNAME
The database name.

AREA
The DEDB AREA name.

DCB NUM
The data set number of this data set within the database or the HALDB partition.

VSAM COMP
Present only in the VSAM data set information report, this field shows the VSAM data set component. It can be either DATA or INDEX.

MASTER HALDB
This field is blank unless the database is a HALDB partition. For partitioned databases, this field shows the master HALDB database name of which this partition is a member.

SUBPOOL
Indicates the name of the subpool that was used by this DBDS during the trace interval. If the database was modified to use a different subpool in the DFSVSAMP statements, the new subpool name is not indicated in this report.

BUFFER SIZE

The size of the buffers that were used by this database during the trace interval. This does not reflect changes to the assigned buffer pool that might have occurred because of a **CHANGE DBD** command or changes in the buffer pool sizes that were made in DFSVSAMP statements. Also, if the DBDS was not used during the trace interval, this field will be blank.

BLOCK SIZE/CI SIZE

The data set block size for OSAM, or CI size for VSAM. This does not reflect updates to the database that were specified in a CHANGE statement that was included in the SYSIN statements.

HIT RATIO

Shows the percentage of buffer requests that were satisfied from buffers in the subpool.

REQUESTS PER SECOND

Shows the average number of database buffer requests per second that were made during the trace interval.

READS PER SECOND

Shows the average number of database I/Os per second that were required to fulfill buffer requests during the trace interval.

Subpool statistics reports

IMS Buffer Pool Analyzer can produce either one or two sets of subpool statistics. The first set is written to DD name SPLRPT. The second set is written only if you requested a change to the subpool configuration.

The reports that are written to SPLRPT document the IMS buffer pool configuration that was in place at the time that trace data was gathered. The second set of subpool statistics are written if you requested a change to the subpool configuration through the DFSVSAMP DD statement, or a change to database access method or block size through a CHANGE statement in the SYSIN DD. The same reports are produced for both the actual configuration that is written to the SPLRPT DD and the model configuration that is written to the MDLRPT DD.

There are several types of reports that can be created, all of which are described in the following sections. Some of the reports, such as the Inefficient buffer size report and the Varied hit ratio report, are exception reports. These exception reports will not appear in the output if no exceptions were encountered.

Subpool summary

The first report in the Subpool Statistics reports is the subpool summary. It shows a summary of all of the defined subpools.

The following figure shows an example of a Subpool summary report (some lines have been removed):

```
DATABASE SUBPOOL SUMMARY
                                    REQUESTS   READS
           BUFFER             NUMBER   PER       PER      HIT    BUFFER    MARGINAL
TYPE        SIZE   SUBPOOL    BUFFERS  SECOND   SECOND   RATIO    LIFE    REDUCTION
────        ────   ───────    ───────  ──────   ──────   ─────   ──────   ─────────
OSAM        4,096  ───────     10,000   1,137     276    75.7%      36           1
OSAM        8,192             20,000    2,644     481    81.8%      41           0
OSAM       10,240               300         1       0    79.6%    1062           0
OSAM       32,768                10     ────────── NO SUBPOOL ACTIVITY ─────────
OSAM       *TOTAL             79,210    33,916   1,576   95.3%

VSAM-I      2,048  DFLT          50     ────────── NO SUBPOOL ACTIVITY ─────────
VSAM-D      2,048  DFLT        3,000        46      14    68.7%     208           0
VSAM-I      4,096  DFLT        4,000     8,837     984    88.8%       4         124
VSAM-D      4,096  DFLT        5,000     5,701   1,642    71.2%       3          21
VSAM-I      8,192  DFLT          250        28      13    55.1%      19           1
VSAM-D      8,192  DFLT        4,000       368       0    99.9%  35,434           0
VSAM-D     12,288  DFLT          50     ────────── NO SUBPOOL ACTIVITY ─────────
VSAM-D     32,768  DFLT           5     ────────── NO SUBPOOL ACTIVITY ─────────
VSAM        8,192  V000        2,500     1,769      51    97.1%      49          13
VSAM-I      4,096  V001          100     4,881       0    99.9%     734      12,641
VSAM-D      4,096  V001        1,000     1,628      32    98.0%      31          38
VSAM       *TOTAL            29,705    26,199   3,416    86.9%

*GRAND    TOTAL             108,915    60,112   4,992    91.6%
```

Figure 22. Sample report: Database subpool summary

The Subpool summary report documents the buffer pool environment. The columns of the summary report show the configuration: subpool type, size, name, and number of buffers, and performance information for each subpool. The report also provides subtotals by access method and grand totals for all access methods.

The Subpool summary report provides a comparison of the extent of buffer pool activity and how hit ratios, buffer life, and marginal reduction compare. You can use this information to make decisions about which subpools are over-allocated, and which subpools could benefit from additional storage. This information can also help you understand why IMS Buffer Pool Analyzer makes certain tuning suggestions.

Inefficient buffer size report

The inefficient buffer size report appears only if at least one database that was used during the trace interval does not have buffers of the most efficient size.

This report shows which DBDSs are using buffer sizes that are larger than are required because a smaller buffer size was not available. An example is shown in the following figure:

```
THE BUFFER SIZE USED BY THE FOLLOWING DATABASES EXCEEDS THE MOST EFFICIENT SIZE
                                       CISIZE      USE       BEST
                                         OR      BUFFER    BUFFER
DBDNAME     DCB    TYPE    SUBPOOL     BLKSIZE     SIZE      SIZE     STORAGE
───────     ───    ────    ───────     ───────   ──────    ──────    ───────
TH            1    OSAM    ───────      8,970    16,384    10,240       768K
VGGDGMP3      1    VSAM-I              10,240    32,768    12,288         3K
VGOMAPP1      1    VSAM-I              12,288    32,768    12,288       157K
*TOTAL                                                                 928K
```

Figure 23. Sample report: Inefficient buffer size report

The fields in the Inefficient buffer size report are as follows:

DBDNAME
   The database name.

DCB
   The data set number of this data set within the database or the HALDB partition.

TYPE
   Database type: OSAM, VSAM-I (VSAM index component), VSAM-D (VSAM data component), or DEDB.

SUBPOOL
> The subpool name, if any.

CI SIZE OR BLKSIZE
> The actual DBDS CI size or block size.

USED BUFFER SIZE
> The buffer size IMS is using for this DBDS.

BEST BUFFER SIZE
> The optimal buffer size for the actual CI size or block size.

STORAGE
> The average amount of storage that could be saved if a subpool of the "best buffer size" was made available for this DBDS. The amount of storage is calculated based on the difference between the used buffer size and the best buffer size, and the average number of buffers occupied by this DBDS.

If the report is written, consider adding a new buffer subpool that allows the DBDS to use the buffer size that is shown as the best buffer size. In some cases, it might be counterproductive to create a new buffer subpool for an infrequently-accessed data set.

In the preceding report example, the second DBDS, VGGDGMP3, would require a minimum of 3 12K buffers (36K), but the amount of storage that could be saved is only 3K.

Varied hit ratio report

The varied hit ratio report appears only if there are varied hit ratio subpools to report.

IMS Buffer Pool Analyzer determines whether a subpool is a varied hit ratio subpool by referring to the VARYPCT statement. For more information about the VARYPCT statement and the condition for reporting, see "Customizing buffer pool reports using SYSIN statements" on page 29.

The following figure shows an example of a varied hit ratio report:

```
                                        HALDB
BUFFER                                  MASTER    HIT
 SIZE         TYPE  SUBPOOL  DATABASE  DCB OR AREA  RATIO
_____       _____ _____  _____  ___ _____  _____
8.192        OSAM  _____  PCZ7AMP    A  PCZAMP   91.5%
                            NIZXCI2    3           84.9%
                            PCZ5AMP    A  PCZAMP   96.7%
                            NIZXCI1    3           71.7%
                            PG15IDP    A  PGZIDP   44.5%
                            NIZICR     2           20.4%
                            NIZICR     3           25.4%
                            PGZTDA     1           88.6%
                            NIZXCI1    2           72.2%
                            NIZXCI2    2           84.9%
```

Figure 24. Sample report: Varied hit ratio report

The fields in the varied hit ratio report are as follows:

BUFFERSIZE
> The subpool buffer size.

TYPE
> Subpool type: OSAM, VSAM-I (VSAM index component), VSAM-D (VSAM data component), or DEDB.

SUBPOOL
> The subpool name, if any.

DATABASE
> The names of the 10 busiest database data sets in the subpool. The DBDSs are listed in the order of the number of buffer requests.

DCB
> The data set number of this data set within the database or the HALDB partition.

HALDB MASTER OR AREA

For a HALDB partition, the name of the HALDB master database; or, for a DEDB, the area name. For a full-function database, this field is left blank.

HIT RATIO

The hit ratio for this DBDS.

This report is generated if this subpool has varied hit ratios. The report shows up to 10 of the busiest DBDSs that use this subpool.

This report might show possible conflicts within a buffer pool. In the preceding report example, the busiest DBDS (PCZ7AMP partition A) has a 91.5% hit ratio, and there are two DBDSs in the 10 busiest databases in that subpool that have hit ratios of 25.4% or less.

For more information about subpools, see "Subpool detail report" on page 45.

Having lower hit ratio databases in the same subpool as much higher hit ratio databases can inhibit the ability to optimally tune your buffer subpool. To tune a subpool effectively, allocate more buffers to very high hit ratio databases and fewer buffers to lower hit ratio databases.

If this report is produced, consider moving either the higher or lower hit ratio databases to a new subpool. You can use the modeling feature of IMS Buffer Pool Analyzer to see the impact of adding a new subpool and assigning some of the DBDSs to the new subpool.

Do this by adding a DFSVSAMP DD statement to the report JCL, and adding the new subpool and DBD statements to your existing DFSVSAMP buffer pool definitions. The MDLRPT DD output will include suggestions about how to allocate buffers to the old and new subpool to maximize the performance of the split subpool.

Subpool buffer change recommendations report

IMS Buffer Pool Analyzer examines each buffer subpool to determine the impact of raising or lowering the number of buffers in each subpool.

It produces a report showing changes that could be made to the number of buffers in the traced configuration, and the resulting change in storage used by IMS and the number of database reads per second.

The following figure shows an example of a Subpool buffer change recommendations report (some lines have been removed):

```
SUBPOOL BUFFER CHANGE RECOMMENDATIONS:

BUFFER    POOL    SUBPOOL   CURRENT    NEW     ----- CHANGE IN -----
SIZE      TYPE     NAME     BUFFERS   BUFFERS   STORAGE  READS/SEC
_____    ____    _____   _____   _____  _____  _____
 4,096    OSAM              10,000    20,000   +40,000K   -11.81
 8,192    OSAM              20,000    20,000       +0K    +0.00
10,240    OSAM                 300       100    -2,000K    +0.00 100 BUFFERS HAVE
                                                                    SAME I/O
32,768    OSAM                  10         0     -320K     +0.00 SUBPOOL NOT USED
 2,048    OSAM-I DFLT          50         0     -100K     +0.00 SUBPOOL NOT USED
 2,048    VSAM-D DFLT        3,000     2,000    -2,000K    +0.51
 4,096    VSAM-I DFLT        4,000     4,000       +0K    +0.00
 4,096    VSAM-D DFLT        5,000     5,000       +0K    +0.00
 8,192    VSAM-I DFLT          250       250       +0K    +0.00
 8,192    VSAM-D DFLT        4,000       300   -29,600K    +0.13 400 BUFFERS HAVE
                                                                    SAME I/O
12,288    VSAM-D DFLT          50         0     -600K     +0.00 SUBPOOL NOT USED
32,768    VSAM-D DFLT           5         0     -160K     +0.00 SUBPOOL NOT USED
 8,192    VSAM   V000        2,500     2,500       +0K    +0.00
 4,096    VSAM-I V001          100       100       +0K    +0.00
 4,096    VSAM-D V001        1,000     1,000       +0K    +0.00
*TOTAL                                          -7,820K   -58.37
```

Figure 25. Sample report: Subpool buffer change recommendations

The Subpool buffer change recommendations report shows IMS Buffer Pool Analyzer's recommendations to improve buffer pool performance. For each subpool, the following fields are displayed:

BUFFER SIZE
The subpool buffer size.

POOL TYPE
Subpool type: OSAM, VSAM-I (VSAM index subpool), VSAM-D (VSAM data subpool), or DEDB.

SUBPOOL NAME
The subpool name, if any.

CURRENT BUFFERS
The number of buffers in the pool at the time the trace was run. Or, the number of buffers requested in the DFSVSAMP DD for a model report in the MDLRPT DD.

NEW BUFFERS
IMS Buffer Pool Analyzer's recommendation for the number of buffers that should be allocated to this subpool.

CHANGE IN STORAGE
The difference between the current amount of storage that is used by the subpool and the amount of storage that the IMS Buffer Pool Analyzer recommends.

CHANGE IN READS/SEC
The difference between the current number of reads-per-second for this subpool and the number of reads-per-second projected for IMS Buffer Pool Analyzer's recommended number of buffers.

IMS Buffer Pool Analyzer examines each subpool and adds or removes buffers while working to keep approximately the same amount of storage in the total buffer pool. In the example above, IMS Buffer Pool Analyzer was able to reduce the total storage that was allocated to the buffer pool by 7820K. This reduction occurred at the same time that the number of reads were reduced by 58 per second, as shown on the *TOTAL line.

Subpool detail report

IMS Buffer Pool Analyzer provides detailed information about performance for each subpool that is defined.

The subpool detail report includes the following four sections:

- Actual subpool configuration
- Actual subpool performance
- Model of varying buffer pool sizes for actual subpool configuration
- Databases with activity in this subpool

The following figure shows an example subpool detail report:

```
PAGE   160                          IMS BUFFER POOL ANALYZER 1.4.0
VSAM 4,096 SUBPOOL DFLT DETAIL              IMS ZICA BUFFER POOL REPORT

 ACTUAL SUBPOOL CONFIGURATION:
  IMSID/JOBNAME:            IMS ZICA
  BUFFER SIZE:             4,096 VSAM INDEX
  BUFFER POOL ID:            DFLT
  NUMBER BUFFERS:          4,000

 ACTUAL SUBPOOL PERFORMANCE:
  BUFFER REQUESTS:          2,661,697   8,837.27 PER SECOND
  BUFFER POOL HITS:         2,365,445   7,853.66 PER SECOND  88.8%
  BUFFER READS:               296,252     983.61 PER SECOND  11.1%
  HIT RATIO:                   88.8%
  BUFFER LIFE:                  4.06 SECONDS
  MARGINAL REDUCTION:           124

 MODEL OF VARYING BUFFER POOL SIZES FOR ACTUAL SUBPOOL CONFIGURATION:

                        ----------------- PROJECTIONS  ------------------
        NUMBER      POOL      HIT      READS PER      BUFFER         MARGINAL
        BUFFERS     SIZE     RATIO      SECOND         LIFE          REDUCTION

        ------     ------    -----     --------     --------       --------
             4       16K     18.2%     7,221.91       0.00            41,512
           100      400K     68.3%     2,793.87       0.03             4,281
           200      800K     73.7%     2,318.18       0.08             1,867
           300    1,200K     76.1%     2,110.66       0.14             1,230
           400    1,600K     77.6%     1,973.95       0.20               922
           500    2,000K     78.8%     1,871.40       0.26               720
           600    2,400K     79.7%     1,791.32       0.33               590
           700    2,800K     80.4%     1,725.76       0.40               502
           800    3,200K     81.1%     1,669.96       0.47               429
           900    3,600K     81.6%     1,622.20       0.55               378
         1,000    4,000K     82.1%     1,580.19       0.63               254
         2,000    8,000K     85.3%     1,297.94       1.54               158
         3,000   12,000K     87.3%     1,122.22       2.67               124
         4,000   16,000K     88.8%       983.57       4.06                 0
         5,000   20,000K     88.8%       983.48       5.08                 0
         6,000   24,000K     88.8%       983.48       6.10                 0
         7,000   28,000K     88.8%       983.48       7.11                 0
         8,000   32,000K     88.8%       983.48       8.13                 0
         9,000   36,000K     88.8%       983.48       9.15                 0
        10,000   40,000K     88.8%       983.48      10.16                 0
        20,000   80,000K     88.8%       983.48      20.33                 0
        30,000  120,000K     88.8%       983.48      30.50                 0
        40,000  160,000K     88.8%       983.48      40.67                 0

 DATABASES WITH ACTIVITY IN THIS SUBPOOL:
            DCB  VSAM   HALDB     CI    HIT    REQUESTS      READS
   DBDNAME  NUM  COMP   MASTER   SIZE  RATIO  PER SECOND  PER SECOND

   _____   ___  _____  _____  _____ _____  _____  _____
   PG15IDF   X   INDEX  PGZIDF  4,096 79.2%      102.28       21.22
   NIZXCI1X  3   INDEX          4,096 69.2%       58.92       18.13
   NIZXCI2X  3   INDEX          4,096 84.6%       84.59       12.94
   NIZXFDRX  1   INDEX          4,096 57.5%       29.99       12.73
   PGZPMTX   1   INDEX          4,096 97.2%      463.02       12.52
```

Figure 26. Sample report: Subpool detail

```
VSAM 4,096 SUBPOOL VDEF DETAIL                    IMS IMB1 BUFFER POOL REPORT

   ACTUAL SUBPOOL CONFIGURATION:
      IMSID/JOBNAME:                 IMS IMB1
      BUFFER SIZE:                     4,096  VSAM    BOTH
      BUFFER POOL ID:                   VDEF
      NUMBER BUFFERS:                  1,600
      NUMBER HIPERSPACE BUFFERS:         400

   ACTUAL SUBPOOL PERFORMANCE (BASE):
      BUFFER REQUESTS:               141,446         156.74 PER SECOND
      BUFFER POOL HITS:              140,949         156.19 PER SECOND      99.6%
      HIT RATIO:                       99.6%
      BUFFER LIFE:                   2905.17  SECONDS
      MARGINAL REDUCTION:                  8

   ACTUAL SUBPOOL PERFORMANCE (HIPERSPACE):
      BUFFER REQUESTS:                   497           0.55 PER SECOND
      HIPERSPACE HITS:                     0           0.00 PER SECOND       0.0%
      HIT RATIO:                        0.0%
      BUFFER LIFE:                   3631.46  SECONDS
      MARGINAL REDUCTION:                  0

   ACTUAL SUBPOOL PERFORMANCE (BASE+HIPERSPACE):
      BUFFER REQUESTS:               141,446         156.74 PER SECOND
      BUFFER POOL HITS:              140,949         156.19 PER SECOND      99.6%
      HIPERSPACE HITS:                     0           0.00 PER SECOND       0.0%
      DASD READS:                        497           0.55 PER SECOND       0.3%
      HIT RATIO:                       99.6%
      BUFFER LIFE:                   3631.46  SECONDS
      MARGINAL REDUCTION:                  0
```

Figure 27. Sample report: Subpool detail with hiperspace buffers

The first section of the subpool detail report provides information about the actual subpool configuration. In the preceding example, the actual subpool configuration is displayed, which indicates that this is the IMS configuration that is used during the trace and that is written to the SPLRPT DD data set. The subpool configuration for a model environment that is written to the MDLRPT DD data set would display the following message SUBPOOL CONFIGURATION FROM DFSVSAMP. The subpool configuration information identifies the IMSID or DL/I batch job name that was traced and the buffer size, buffer pool ID, and number of buffers in the subpool.

The second section of the subpool detail report provides information about the actual subpool performance. The preceding example displays the actual subpool performance that is based on the trace data. If a model configuration is in place, the report message in the MDLRPT DD data set would display DFSVSAMP SUBPOOL CONFIGURATION PERFORMANCE, and the performance metrics would reflect the subpool configuration as it was requested in DFSVSAMP, along with any CHANGE DBD statements that were specified on the SYSIN DD statement.

The performance metrics that are listed in the actual subpool performance section show the number of buffer requests, buffer pool hits, DASD buffer reads, the per-second rates of these metrics, and the percentages of requests that resulted in buffer pool hits, DASD reads, then the hit ratio, buffer life, and marginal reduction value.

The third section of the subpool detail report shows projections for subpool performance metrics for different amounts of storage that are allocated to the subpool, and it displays the resulting performance for different numbers of buffers that might be allocated to the subpool. The number of buffers on each line of the report is determined based on the default numbers of buffers for a subpool. Or, you can specify the number of buffers on each line using **BUFNO** statements in the SYSIN DD statement for the report job.

IMS Buffer Pool Analyzer uses the information in performance projections to review subpool performance and determine whether performance can be improved by increasing or decreasing the number of buffers that are allocated to each subpool.

The fourth section of the subpool detail report lists all of the database data sets that had activity in the subpool, and performance metrics are shown for each data set. The data sets are presented in order of the highest through the lowest number of reads-per-second, which means that the highest I/O data sets are displayed first in each subpool's list of active databases.

Chapter 6. Tuning database buffer pools

Tuning is the process of balancing the distribution of resources to maximize throughput.

Topics:

Function and effect of I/O, storage, and CPU resources

When tuning IMS database buffer pools, you should understand the function and effect of three main resources: I/O, storage, and CPU.

Tuning resources helps to increase speed and performance. The following descriptions point out the functions of various resources.

I/O

DASD I/O
DASD I/O is usually the target of reduction, because it is very slow. One thing to consider is that I/O is also a CPU-intensive operation. In CPU-constrained environments, reducing I/O can relieve the CPU resource.

Paging I/O
Paging occurs when MVS determines that there is not enough real storage to meet the needs of all tasks on an LPAR. If increasing IMS buffer storage to reduce DASD I/O causes paging I/O, the net result can be worse performance than would result with increased DASD I/O.

Storage

Central Storage (pre zSeries/Architecture)
This is the traditional storage that is present in the mainframe environment. Increasing the number of buffers in a pool increases IMS subsystem's use of central storage space.

Expanded Storage (pre zSeries/Architecture)
Expanded storage can be used for VSAM hiperspace buffers or for MVS paging. Excessive use of expanded storage can increase MVS paging I/O, and worsen IMS and system performance.

64-bit Storage (zSeries/Architecture)
In the zSeries/Architecture, expanded storage is combined with central storage. Because the combined total of real storage can exceed 2G, storage above this line is referred to as 64-bit real storage.

CPU
The CPU resource can increase if there are not enough buffers, causing excessive DASD I/O, or if buffer pools are so large that IMS has to search excessively on each buffer request.

Restriction: If your z/Series processor is not running in compatibility mode, 64-bit storage is available only for OSAM database buffers.

Be aware that database I/O avoidance should not be the only consideration when tuning IMS database buffer pools.

It is important to review available real storage before making any changes to the database buffer pools. The RMFMON paging activity report (identified as SPAG on the RMFMON main menu) can provide short-term indications of the number of free frames of central storage (AFC is the number of 4K frames that are available) and expanded storage (ESF AVL is the number of expanded storage frames available).

If you plan to increase the size of any buffer pools, make sure that paging rates will not be adversely affected before implementing the change. Before making such changes, review RMF data or discuss the situation with your system performance team.

Buffer pool and subpool organization

IMS allows for separating buffer pools by data set, by group of data sets or, for HALDB, by partition.

You can then size OSAM and VSAM buffer pools to reduce the number of database I/Os required to complete a transaction or batch workload.

OSAM buffer pools

OSAM database buffer pools can direct specific DBDSs to specific buffer pools. There is a default buffer pool, typically with several subpools containing different size buffers. For example, the default buffer pool can have one subpool containing 1024-byte buffers, another subpool containing 8,192-byte buffers, and a third subpool containing buffers with 32,768 bytes each.

VSAM buffer pools

VSAM provides for 15 buffer pools, with *subpools* within each pool. A subpool is a group of buffers of a specific size and component type. VSAM DBDSs can have two types of components: index components and data components. You can have index and data buffers in the same subpool, or you can separate index buffers from data buffers by using two different subpools.

When a DBDS is opened, one specific subpool is assigned for the buffers of that DBDS or possibly, two subpools for a KSDS DBDS, because it contains both a data component and an index component. During a database open, IMS looks at definitions in the DFSVSAMP statements to determine whether this DBDS should be assigned to the default buffer pool or a user-specified buffer pool. Once the buffer pool is identified, the "best fit" subpool within that buffer pool is assigned for use by the buffers of that DBDS. For a VSAM KSDS, there are two subpools assigned, one for the data component and another for the index component.

One important detail for VSAM data and index pools for a given KSDS with both data and index components: both the data component and the index component must reside in the same buffer pool, although they can be in different subpools within that pool.

Buffer pool definitions - DFSVSAMP statements

Buffer pool specifications are defined in DD statements, which are referred to as *DFSVSAMP statements*.

IMS buffer pool specifications are defined in statements that are read from the following sources:

- DFSVSAMP DD for a batch DLI or DBB job
- DFSVSM*xx* member of IMS PROCLIB for a DB/DC or DBCTL subsystem where, *xx* is the VSPEC= value in the DFSPB*xxx* member of PROCLIB

These statements are referred to as DFSVSAMP statements.

The following example shows how an OSAM buffer pool with three subpools could be defined. The example defines a 512-byte subpool with 4 buffers, a 2K subpool with 5 buffers, and a 4K subpool with 12 buffers.

```
IOBF=(512,4,Y,Y)
IOBF=(2048,5,Y,Y)
IOBF=(4096,12,N,N)
```

Figure 28. Example of an OSAM buffer pool with three subpools

Additional parameters are required when defining subpools, but for this discussion, only the buffer pools, subpools, and number of buffers are documented. All DFSVSAMP statements and parameters are documented in *IMS System Definition*.

IMS defines default OSAM subpools as the first subpool with that buffer size. Any DBDS can use that subpool if it is not specifically assigned to another subpool name. If you define two 4K OSAM subpools named SP01 and SP02, all 4K OSAM DBDSs will use SP01 unless you direct IMS to use SP02, by including a DBD= statement in the DFSVSAMP statements.

The two statements in the following example define a pool named SP01 with two subpools, one with 100 4096-byte buffers and the other with 200 8192-byte buffers.

```
IOBF=(4096,100,N,N,SP01)
IOBF=(8192,200,N,N,SP01)
```

Figure 29. Sample OSAM buffer pool with two subpools

DFSVSAMP statements for VSAM pool and subpool definitions are similar. If multiple VSAM pools are to be defined, a POOLID statement defines the name of a VSAM pool, and VSRBF statements that follow the POOLID define each subpool within that pool. The following example defines a default subpool named DFLT with the same subpool configuration as in the preceding OSAM example and a second pool named SP01 with the same configuration as in the OSAM example.

```
POOLID=DFLT
VSRBF=512,4
VSRBF=2048,5
VSRBF=4096,12
POOLID=SP01
VSRBF=4096,100
VSRBF=8192,200
```

Figure 30. Sample default subpool

VSAM subpools have the additional option of specifying that a particular subpool is to be dedicated to either index or data buffers. This option is specified by adding a D for data or an I for index following the number of buffers on the VSRBF= statement. In a given VSAM pool, if any subpool is dedicated to data or index buffers, then all subpools in that pool are also dedicated. If the D or I option is not specified on a VSRBF= statement in such a VSAM pool, then data is assumed.

In the following example, VSAM pool SP02 has 10 1K data buffers, 20 2K index buffers, and 40 4K data buffers. The 4K buffers are data buffers because data is the default when neither data nor index is specified on a subpool definition in a pool that has separate data and index pools.

```
POOLID=SP02
VSRBF=1024,10,D
VSRBF=2048,20,I
VSRBF=4096,40
```

Figure 31. Sample VSAM buffer pool

VSAM allows for 15 buffers pools, each of which can have multiple subpools with different CI sizes or DATA/INDEX types. In a DLI/DBB batch job, subpool 15 is reserved for RECON data sets, so batch IMS has a limit of 14 buffer pools. Within each VSAM buffer pool, there can be as many as 22 subpools. There are 11 CI sizes (512, 1K, 2K, 4K, and multiples of 4K through 32K), each of which can have a data and an index specification.

In addition, VSAM subpools that are 4K in size or more can have hiperspace buffers. Hiperspace buffers, in a pre zSeries/Architecture environment, provide additional buffers that reside in expanded storage instead of central storage. In central storage-constrained environments, hiperspace buffers can provide additional buffering without consuming additional central storage. Hiperspace buffers are specified on the VSRBF= statement. See *IMS System Definition* for details on specifying hiperspace buffers.

When creating a new buffer pool, DBDSs must be assigned to the new buffer pool explicitly on the DBD= statement. The statement in the following example assigns data set number 1 of database DI21PART to subpool SP01.

```
DBD=DI21PART(1,SP01)
```

Figure 32. Assignment of DBDS to new buffer pool

Tuning process overview

During the tuning process, you can review resources, analyze I/O information, balance subpools, and validate the results.

This topic describes the process of tuning buffer pools, which involves the following activities:

1. Reviewing storage resources and paging activity
2. Gathering BPLTRACE database I/O information
3. Reviewing buffer pool sizes
4. Balancing subpools
5. Reviewing subpool buffer allocations
6. Validating the tuning results

Reviewing storage resources and paging activity

Before increasing any IMS buffer pool, you should understand not only the implications of real storage and paging activity, but also the implications of virtual storage.

About this task

Review real storage and paging rates to ensure that the addition of buffers will not adversely impact either IMS or the other tasks on that MVS LPAR.

Virtual storage should also be a consideration when increasing database buffer pools. All VSAM database buffers are located in 31-bit virtual storage (above the 16M line) except for hiperspace buffers, which are located in a separate data space. When increasing database buffer pool sizes, make sure that sufficient virtual storage is available in the batch job's address space or in the IMS DLI address space.

Virtual storage usage and limits can be found in the JES messages for the job immediately following the condition code for the job. Look for an MVS system message, such as the IEF374I message shown in the following example:

```
IEF374I STEP/G/STOP  2021221.0348 CPU    0MIN 00.56SEC SRB
0MIN 00.10SEC VIRT   852K SYS   304K EXT    1528K SYS    9876K
```

Figure 33. Message IEF374I regarding virtual storage use and limits

The virtual storage utilization information is at the end of the message. This example shows that the job used 852K of private area below the 16M line and 304K of local system queue area (LSQA). It also used 1528K of private area above the 16M line (extended private area) and 9876K of ELSQA (extended LSQA). If you compare the extended private area used with the available extended private area (available either from the **REGION=** JCL specification or the default extended private area size), you can ensure that increasing the database buffer pool space does not cause problems because of insufficient virtual storage.

Gathering BPLTRACE database I/O information

Analysis of trace data is critical to the steps in the tuning process. For an online subsystem, gathering multiple traces at different times helps provide a better sample and ensures that performance is reasonable at all times.

About this task

Sample several different workload types if appropriate, for example: prime time workload, nighttime batch workload, and so on. The length of time to perform trace activity depends on the amount of space in the GTF trace data set, and on the number of buffers in the buffer pools. In order to provide reasonable projections, the number of OSAM and VSAM records that are captured in the traces should be at least five

times the total number of buffers in your buffer pools. The more trace data gathered, the more accurate the projections will be.

For a batch job, if the job has consistent database call patterns, several times the number of buffers in the buffer pools should provide reasonable estimates. If the batch job first reads large amounts of data and then does updates, then either one long trace interval or several shorter trace intervals should provide reasonable projections.

Reviewing buffer pool sizes

Review buffer pool sizes to ensure that subpools with appropriate buffer sizes exist.

About this task

For example, if there is a VSAM data set with a 2K CI size but the VSAM buffer pool assigned to that data set only has 1K, 4K, and 8K buffers, then half the storage that is used by buffers of the 2K data set will be wasted because the data set will use the 4K buffers.

IMS Buffer Pool Analyzer will identify any DBDSs that do not have the most efficient buffer size available. If the inefficient buffer size report is not present in the subpool report SPLRPT DD or model report MDLRPT DD, then no databases were identified with inappropriate buffer sizes. If the report is present, you can use the report to identify potential savings from adding new buffer subpools with more efficient buffer sizes. You can then use IMS Buffer Pool Analyzer's modeling feature to define new subpools, and rerun the report to see the impact of the change.

After reviewing the inefficient buffer size report, if you choose to add new subpools, you can use the DFSVSAMP input to allow IMS Buffer Pool Analyzer to determine how many buffers to assign to the old and new subpools. The MDLRPT DD reports will show recommendations for how many buffers to assign to each subpool in your buffer pool configuration.

Balancing subpools

When reviewing the existing buffer pool configuration, also review the busiest data sets in each subpool.

About this task

A database subpool can be overwhelmed by a large DBDS. A single BMP could easily scan a data set and replace every other database's buffers. Therefore, if the hit ratio for any of the busy data sets is lower than the hit ratio for other data sets in the subpool, consider moving the data set with the lower hit ratio to its own subpool or a subpool that is shared with other data sets with similar activity.

Among the most obvious examples of balancing subpools is separating VSAM data and index buffers into separate subpools. Presumably, VSAM index components will have a very high hit ratio, which means that they should not be in a subpool that is shared with a data set that could replace all the index buffers on a regular basis.

IMS Buffer Pool Analyzer will document any subpools in your configuration that have data sets with varied hit ratios. Review this report to determine whether you could make changes in your subpool configuration that might provide more efficient storage utilization. You can use IMS Buffer Pool Analyzer's modeling feature to split a subpool and reassign high or low hit ratio databases into different subpools.

If you confirm that there are problems with unbalanced subpools, use the DFSVSAMP input to the buffer analysis job to add new subpools and reassign DBDSs to different subpools. The MDLRPT DD will include recommendations for the number of buffers to allocate to all impacted subpools.

Reviewing subpool buffer allocations

Because I/O rates and storage are the primary drivers of a general performance tuning strategy, hit ratios and buffer life should not drive buffer pool allocations, but should be used as an ongoing monitor of performance.

About this task

To determine the appropriate number of buffers for each subpool, keep the following points in mind.

Note: There can be application-specific requirements for performance that override this general tuning strategy.

- The number of buffers for each subpool should be large enough to hold the buffers that are required by any concurrently running applications. This should help you avoid writing buffers to make room for additional reads. Basically, be careful to not make any buffer pool too small.
- Avoid making buffer subpools too large. Wasting storage causes performance degradation of, not only IMS, but of the entire MVS LPAR.
- OSAM sequential buffering can significantly improve some applications.
- OSAM buffer caching in a coupling facility can improve performance but, unless storage in your environment is severely constrained, is not a substitute for a well-tuned buffer pool.

A number of strategies have been used to tune a buffer pool. Buffer pools can be tuned to try to achieve a specific hit ratio in each subpool (for example, 80% for data subpools and 90% for index subpools) or to try to achieve a specific buffer life in each subpool (for example, 30 seconds to 5 minutes in each subpool). These strategies fail to consider that different databases have different access patterns. While it might be possible to tune buffer pools to achieve these objectives, it might be possible to change buffer pool allocations to achieve significantly lower I/O rates using significantly less storage.

The strategy presented here for tuning buffer pool allocations strives to achieve the lowest I/O rate with the least use of storage. IMS Buffer Pool Analyzer now provides a report that uses the following methodology to produce buffer pool recommendations. You can review IMS Buffer Pool Analyzer's recommendations to determine if they make a sensible alternative configuration in your environment. However, note that Buffer Pool Analyzer does not make recommendations for the addition or removal of subpools (as described later). Therefore, you should still understand the process that follows, and review whether to add new subpools or consolidate existing subpools.

The Marginal reduction statistics section in the Database subpool statistics report provides a way to compare subpool efficiency as the number of buffers is increased. Each line of the report shows a different number of buffers and information about subpool performance using that number of buffers. The marginal reduction number shows how many I/Os are reduced from the prior line's buffer number per kilobyte of storage that was added to the pool. This provides a way to identify which subpool would benefit the most from adding more buffers.

```
                  ----------------- PROJECTIONS ------------------
NUMBER      POOL     HIT    READS PER   HITS PER    BUFFER    MARGINAL
BUFFERS   SIZE (K)   RATIO   SECOND     SECOND      LIFE     REDUCTION
-------   --------  ------  ---------  ---------   --------  --------
      4        16    81.3%     188.6      822.5       0.02
      8        32    85.3%     148.5      862.6       0.05     9,020
     16        64    88.4%     117.2      893.9       0.13     3,517
     32       128    91.2%      88.9      922.2       0.35     1,593
     64       256    93.1%      69.0      942.0       0.92       557
    128       512    94.9%      51.1      960.0       2.50       252
    192       768    95.9%      40.5      970.6       4.73       149
    256     1,024    96.7%      33.0      978.1       7.74       104
    384     1,536    97.4%      26.1      985.0      14.69        48
    512     2,048    97.6%      24.0      987.0      21.24        14
    768     3,072    97.6%      23.4      987.7      32.73         2
  1,024     4,096    97.6%      23.4      987.7      43.64         0
  1,536     6,144    97.6%      23.4      987.7      65.46         0
  2,048     8,192    97.6%      23.4      987.7      87.28         0
```

Figure 34. Sample report: Marginal reduction statistics

Procedure

The process of determining how many buffers should be allocated requires that you perform the following tasks:

1. Determine, for each subpool, the marginal reduction number that applies to the subpool's current number of buffers. If the current number of buffers in the pool is not listed, estimate the number by using the marginal reduction number from the two closest lines.

   In the marginal reduction statistics example, if the current number of buffers in the subpool is 64, the marginal reduction number would be 557. If the number of buffers defined is 100, you could estimate the marginal reduction number as 400 (approximately the average of 557 and 252, the marginal reduction numbers for 64 and 128 buffers, respectively).

2. For any subpools with a marginal reduction of 0, find the lowest number of buffers in that subpool that has a marginal reduction of 0. Note this number of buffers as the new recommended number of buffers.

   In the marginal reduction statistics example, if the number of buffers currently defined in the subpool is 2048, then 1024 buffers is the least number of buffers with a marginal reduction number of 0. Because these are 4096-byte buffers, note that 4M of storage is freed for use in increasing other subpools.

   a. If there are fewer buffers remaining in the pool than the number of dependent regions or the **PST=** parameter in DFSPB*xxx*, use the number of dependent regions for the number of buffers.

   b. If the number of buffers in the pool is less than 10, consider merging this subpool into the next larger subpool. If either this pool or the next larger pool is very busy (either in I/O rate or request rate), leave the subpool as it is. For subpools with a small number of buffers, consider adding buffers to allow for growth. For example, if an index has a CI split and is the only data set in a buffer pool that is tuned exactly to the number of index CIs, the CI split will cause a very large number of I/Os to occur.

3. Compare the marginal reduction numbers for each subpool. Those with a large marginal reduction number benefit the most from the addition of buffers. Those subpools with low marginal reduction numbers can have the number of buffers reduced with the least impact. Adjust the number of buffers in each subpool to make the marginal reduction for each subpool closer. Depending on whether additional real storage is available, it might be necessary to reduce the size of some subpools in order to increase the size of other subpools.

Validating the tuning results

To ensure that you have achieved the tuning results that you were seeking, you can validate the results using the process that is described in this section.

Procedure

To validate results of your tuning efforts, follow these steps:

1. Run the IMS Buffer Pool Analyzer reports using the number of buffers you determined with the preceding reviewing subpool buffer allocations process.

2. Compare the total buffer pool size in the original configuration with the total buffer pool size in the new configuration.

   Because this is the amount of storage that will be added to the virtual storage that is used by the address space, ensure that the job or started task has enough virtual storage to accommodate the addition. Also, ensure that the additional use of real storage will not adversely impact MVS performance and paging.

3. Verify that the total I/O rates for the new configuration are lower than those for the original configuration.

4. Run additional trace sessions to validate the new configuration.

   If you take several traces at different times, run all the traces into a single report, along with running the traces individually. Ensure that the changes to improve the average workload do not adversely impact any particular workload, such as nighttime BMP work or other work.

Implementing the updated buffer specifications

To implement buffer pool specifications, bring the IMS subsystem down and then back up. Ensure that any changes that you make to affect one IMS or batch job are not unintentionally implemented in other jobs.

Chapter 7. Coupling facility IMS OSAM cache reports

In addition to database buffer pools, IMS provides the option of using storage in a coupling facility to cache OSAM database blocks.

Topics:

Coupling facility IMS OSAM cache overview

You can efficiently store your database blocks using the IMS OSAM caching method.

The IMS OSAM cache was designed for use in a data sharing environment, and the IMS OSAM cache works most efficiently for small databases that are updated frequently. The IMS OSAM cache is not a substitute for the typical database buffer pools that are allocated in storage of the IMS subsystem.

Coupling facility IMS OSAM caching requires that IMS resource lock manager (IRLM) is used for block-level database sharing, and that a structure is allocated in a coupling facility to cache OSAM blocks.

Although retrieving a database block from the cache in a coupling facility structure is not nearly as efficient as obtaining a block from the database buffer pool, this approach is typically more efficient than reading the block from DASD.

Important: Using the IMS OSAM cache results in some additional overhead, however. The database blocks must be written to the coupling facility structure so that they can be retrieved later. This process introduces additional I/O activity to keep the coupling facility structure up-to-date and populated with the most current blocks.

Implementing the IMS OSAM cache

You need to set up the IMS OSAM cache structure size and configure IMS for an *n*-way block-level sharing environment.

Procedure

To implement the IMS OSAM cache structure size, perform the following steps:

1. Ensure that the IRLM lock table structure is used as your IMS Resource Lock Manager.
2. Ensure that the PROCLIB member, DFSVSM*xx* includes the following structures and ratios:

   - A CFNAMES statement that specifies names for the IRLM lock table structure
   - The VSAM cross invalidation (XI) structure
   - The OSAM cache structure
   - The CFNAMES statement must specify a directory ratio
   - The CFNAMES statement must specify an element ratio for the OSAM cache structure

3. Update at least one coupling facility to define the structures that are named in the CFNAMES statement.

   You update your coupling facilities by updating the coupling facility resource management (CFRM) policy. The CFRM policy specifies the names and sizes of the structures that are included in the coupling facilities.

   For more information about the process and statements that are associated with coupling facility structure definitions, see *z/OS MVS Setting Up a Sysplex*.

4. When you have updated your IMS configuration and coupling facility to enable IMS OSAM caching, determine and note the following three pieces of information:

   - Which subpools should be cached
   - How large the OSAM cache structure should be
   - What the directory ratio and element ratio should be

IMS OSAM cache structure storage sizing

The structure of the IMS OSAM cache storage is divided into three parts: directory space, element space, and structure overhead.

Directory space
Directory space storage is used to track information about specific database blocks, including which subsystems have each block in a local buffer pool and whether the block was updated. Each directory entry takes 200 - 350 bytes of structure storage, depending on the level of software that is running in the coupling facility. You must have created directory entries, for all blocks that are cached in the structure, and also for a directory entry for each OSAM buffer in each IMS subsystem (DL/I batch or control region) that is active in the data sharing environment.

Element space
Element space is storage that is used for database blocks. For an IMS cache, each element is 2048 bytes. Each database block in the cache will occupy from 1 - 16 elements because the maximum OSAM block size is 32 KB.

Structure overhead
Structure overhead storage is a fixed amount of storage that is required by the coupling facility software to manage the structure. The size of the overhead varies depending on the level of software that is in the coupling facility. Structure overhead can range from under 100 KB to nearly 4 MB.

IBM provides a web-based tool called IBM System z® Coupling Facility Structure Sizer Tool (CFSizer), which can assist you in determining the structure size that is required for an IMS OSAM cache structure for a specified number of buffers, blocks to cache, and average block size. The CFSizer tool provides storage estimates that are based on the most current software level for a coupling facility. You can access this CFSizer tool at: https://www.ibm.com/support/pages/cfsizer.

Specifying directory and element ratios

Specify the amounts of storage for each of the three parts of the IMS OSAM cache structure: directory space, element space, and structure overhead.

Procedure

Specify the amount of storage to allocate to directory space and the amount of storage to allocate to element space by performing the following steps:

1. Specify the amount of storage to allocate to directory space and the amount of storage to allocate to element space by specifying a directory ratio and an element ratio.

   The coupling facility software uses these ratios to determine how many bytes are allocated to directory space and element space.

2. Specify the directory ratio and element ratio in the CFNAMES statement in the IMS PROCLIB member DFSVSM*xx*.

Specify each ratio as an integer from 1 - 999. For example, if you specify a directory ratio of 1 and an element ratio of 4, four elements are allocated in the structure for each directory entry.

These settings are ideal if you have only 8 KB OSAM buffers and all buffers are cached because each 8 KB OSAM buffer requires four 2 KB elements.

3. In your environment, you will probably not cache every buffer pool. You might have more than one buffer size that you are caching. Each of these factors changes the ideal ratio settings. If you know the following pieces of information, you can use the CFSizer tool to estimate the amount of storage that is required for the current CF level:

- The maximum number of buffers that are defined in your IMSplex (including control regions and DL/I batch jobs)
- The number of blocks that you want to cache
- The average size of the blocks

Results

The CFSizer tool provides a method to determine an appropriate structure size, but the tool does not provide the directory and element ratios that are required to properly use the structure. IMS Buffer Pool Analyzer provides all of the information that is required to properly size and allow IMS to use the coupling facility structure.

When IMS Buffer Pool Analyzer calculates estimates for storage sizes, it uses storage amounts that are based on the coupling facility software level that was active when your trace data was created.

If you upgrade your CF level after trace data is created, either recreate trace data or use the CFSizer tool to estimate storage requirements.

The CFSizer tool is available at: https://www.ibm.com/support/pages/cfsizer.

Using IMS Buffer Pool Analyzer for an IMS OSAM cache

IMS Buffer Pool Analyzer can help you determine the appropriate choices for an IMS OSAM cache configuration.

IMS Buffer Pool Analyzer can display the number of buffers that are associated with a particular structure size, and what directory ratio and element ratio are the most efficient for each structure size. IMS Buffer Pool Analyzer can also help you make decisions about which subpools to cache by allowing you to choose a subpool to cache and by modeling the resulting configuration. You can also create a new subpool, assign specific databases to the new subpool, and model the cache for just the new subpool.

The IMS Buffer Pool Analyzer coupling facility report provides information about IMS OSAM cache configuration and performance. Generating the IMS Buffer Pool Analyzer coupling facility report is a two-step process. First, you gather appropriate trace data from subsystems in the IMSplex. Then, use the coupling facility report job to display configuration and performance information about the environment.

How trace data is obtained for coupling facility reports

The IMS Buffer Pool Analyzer trace data is obtained for an IMS OSAM cache the same way that it is obtained for a buffer pool analysis.

You can use the same trace files for both reporting jobs. For more information about creating trace data, see Chapter 3, "Gathering trace data," on page 15.

To improve the accuracy of the coupling facility reports, include trace data from each IMS subsystem that participates in the data sharing environment. Unlike the buffer pool reporting process, the coupling facility reporting process accepts input from different IMS subsystems and DL/I batch jobs.

When you create trace files for the coupling facility reporting process, specify similar trace durations for each subsystem to improve the accuracy of those reports. For example, if one subsystem has 20 minutes of trace data and a second subsystem has only 5 minutes of trace data, the reporting will be skewed to favor the subsystem with the longer duration. The trace data does not need to be created at the same

time for each subsystem, although if you use IMS shared queues, using traces from each subsystem at the same time will provide a more accurate analysis.

Coupling facility reports

IMS Buffer Pool Analyzer's coupling facility report is a batch reporting process that reads data from all of the trace files that are provided.

IMS Buffer Pool Analyzer's coupling facility reports on your current configuration and provides projections that you can use to determine what changes in your IMS OSAM caching configuration will benefit performance.

BPLCF coupling facility report

The IMS Buffer Pool Analyzer coupling facility (BPLCF) report provides information about your current IMS OSAM data caching environment and projections on the effect of making changes to your configuration.

You can determine what effects the following types of changes will have on your environment:

- Changing the size of the IMS OSAM cache structure
- Changing the directory ratio and element ratio
- Changing the subpools that are cached in the coupling facility
- Changing which database data sets are included in which IMS subpools

These options allow you to create models that reflect possible IMS OSAM caching configurations. The reports that are created for each model show the performance information for that configuration, which allows you to determine whether the modeled configuration achieves your objectives.

BPLCF JCL

You can generate a coupling facility report by running the batch coupling facility report job.

The following sample JCL is included in the SBPLSAMP library in member BPLCF.

```
//BPLCF    JOB
//*
//* This job runs the Buffer Pool Analyzer coupling facility report.
//*
//  SET  SBPLLINK=BPL140.SBPLLINK          SBPLLINK DSN
//  SET  ITKBSRVR=                         OPTIONAL ITKB SERVER NAME
//*
//* Any number of GTF trace data sets may be included in each
//* execution of the reporting utility.  Each GTF trace data set must
//* be allocated to a DD statement that begins with DDNAME TRACE.  If
//* there were three GTF trace data sets, a report execution could
//* include DDNAMES TRACE1, TRACE2, and TRACEXYZ.
//*
//BPL      EXEC PGM=BPLCF,REGION=64M,PARM='ITKBSRVR=*!ENT!*'
//STEPLIB  DD   DSN=*!ENT!*,DISP=SHR
//SYSPRINT DD   SYSOUT=*
//GTFRPT   DD   SYSOUT=*
//SPLRPT   DD   SYSOUT=*
//MDLRPT   DD   SYSOUT=*
//DBDRPT   DD   SYSOUT=*
//DSNRPT   DD   SYSOUT=*
//SYSUDUMP DD   SYSOUT=*
//TRACE1   DD   DSN=???,DISP=SHR
//TRACEXYZ DD   DSN=???,DISP=SHR
//DFSVSAMP DD   DUMMY
//SYSIN    DD   DUMMY
//
```

Figure 35. Sample JCL for running the coupling facility report job

You can define the following symbolic parameters in the sample job.

SBPLLINK

The SBPLLINK symbolic parameter sets the data set name of the IMS Buffer Pool Analyzer load library that is used in the STEPLIB DD statement.

ITKBSRVR

You can use the optional ITKBSRVR parameter to set the XCF group name for the IMS Tools Knowledge Base servers on the LPAR where the report job will run. If IMS Tools Knowledge Base is not used in your installation, or if you do not want to write the coupling facility report to the Knowledge Base repository, leave this symbolic parameter blank.

The following DD names are included in the JCL for the Buffer Pool report:

STEPLIB

Required DD statement that must refer to a data set that contains the IMS Buffer Pool Analyzer load modules. The library does not require APF authorization for any batch report jobs.

SYSPRINT

Print output file for error messages and SYSIN control statements. This DD statement is required. The DCB information for this DD statement, which does not need to be specified, is RECFM=FBA,LRECL=133.

GTFRPT

Print output file for the GTF trace report. If you do not specify this DD statement, the GTF trace report will not be produced. DCB information is not required, but is RECFM=FBA,LRECL=133,DSORG=PS.

SPLRPT

Print output file for the current configuration (actual) coupling facility reports. If you do not specify this DD statement, the current configuration reports will not be produced. DCB information is not required, but is RECFM=FBA,LRECL=133,DSORG=PS.

MDLRPT

Print output file for the model configuration coupling facility reports, if an alternate configuration was requested in a DFSVSAMP DD statement or a CHANGE DBD statement in the SYSIN DD. If you do not specify this DD statement, the model subpool report will not be produced. DCB information is not required, but is RECFM=FBA,LRECL=133,DSORG=PS.

DSNRPT

Print output file for the Data Set Information report. You can disable this report by using the SYSIN statement DSNRPT=NO. If you do not specify this DD statement, the Data Set Information report will not be produced. DCB information is not required, but is RECFM=FBA,LRECL=133,DSORG=PS.

SYSUDUMP

Diagnostic information. If an abend occurs, the information that is produced by specifying this DD statement can be useful for troubleshooting purposes.

TRACE*xxx*

You can include an unlimited number of TRACE*xxx* DD names in the report job. The last three characters of the DD name can be any characters that are valid in JCL. All of the specified trace data sets are read and combined for reporting purposes. At least one TRACE*xxx* DD statement is required.

Restriction: Each TRACE*xxx* DD statement can have only one trace data set. Concatenated data sets are not supported. To include multiple trace data sets, code each data set with a different DD name.

For other considerations, see "Combining multiple traces into a single report" on page 19.

DBDRPT

Optional DD statement that can be used to specify an alternate configuration to model. When provided, this data set must contain IMS buffer pool definition statements, as described in the *IMS System Definition*, and IMS Buffer Pool Analyzer control statements described in "DFSVSAMP statements" on page 62.

SYSIN

Optional DD that can be used to specify IMS Buffer Pool Analyzer coupling facility report options. When provided, this data set must contain IMS Buffer Pool Analyzer statements, as described in "SYSIN statements" on page 63.

DFSVSAMP statements

You can use IMS Buffer Pool Analyzer to specify an alternate configuration to model. The DFSVSAMP DD statement allows you to change the buffer pool configuration of each IMS subsystem that is traced.

You can update the configuration of each subsystem in the IMSplex because the buffer pool configuration specifies which subpools to cache and which database data sets to include in each subpool.

Use the same syntax in the DFSVSAMP DD statement that is used to define your buffer pool configuration to IMS. For more information about buffer pool configuration statements, see the *IMS System Definition*.

Identify the subsystem that each set of buffer pool statements applies to because you can include buffer pool specifications from multiple subsystems in the DFSVSAMP file. Use the IMSID and JOBNAME statements before each set of buffer pool specifications to identify the subsystem to which the buffer pool statements apply.

The following example shows a DFSVSAMP DD statement for the coupling facility report.

```
//DFSVSAMP DD  *
OSAM_STR_SIZE=8192

IMSID=IMB1
VSRBF=2048,800
VSRBF=4096,1600
VSRBF=8192,500
VSRBF=12288,200
IOBF=(4096,1000,N,N,O04K)
IOBF=(8192,200,N,N,O08K,A)
IOBF=(12288,200,N,N,O12K,A)
CFNAMES,CFIRLM=IRLM_CF_NAME_16B
CFNAMES,CFVSAM=VSAM,CFOSAM=(OSAM,1,4)

IMSID=IMB2
VSRBF=2048,800
VSRBF=4096,1600
VSRBF=8192,500
VSRBF=12288,200
IOBF=(4096,1000,N,N,O04K)
IOBF=(8192,200,N,N,O08K,A)
IOBF=(12288,200,N,N,O12K,A)
CFNAMES,CFIRLM=IRLM_CF_NAME_16B
CFNAMES,CFVSAM=VSAM,CFOSAM=(OSAM,1,4)

JOBNAME=BATCH1
VSRBF=8192,50
IOBF=(4096,100,N,N,O04K)
IOBF=(8192,20,N,N,O08K,A)
IOBF=(12288,20,N,N,O12K,A)
CFNAMES,CFIRLM=IRLM_CF_NAME_16B
CFNAMES,CFVSAM=VSAM,CFOSAM=(OSAM,1,4)
/*
```

Figure 36. Sample DFSVSAMP specification for the coupling facility report

The DFSVSAMP DD statement can contain the following statements:

OSAM_STR_SIZE
> The **OSAM_STR_SIZE** statement is an optional statement. If **OSAM_STR_SIZE** is not specified, IMS Buffer Pool Analyzer uses the same structure size that was present during the trace. The statement must begin in column one and should not have any blanks embedded in the statement.
>
> The syntax of the OSAM_STR_SIZE statement is:
>
> ```
>     OSAM_STR_SIZE (number)
> ```
>
> The *number* variable is the number of kilo-bytes to allocate to the structure. For example, if your desired structure size is 32,768K, specify **OSAM_STR_SIZE**=32768.
>
> The value of *number* is used in the model reports that are generated to the MDLRPT output report file.

IMSID

The IMSID statement (or JOBNAME statement) must be included before any IMS DFSVSAMP buffer pool statements are specified.

The IMSID statement is used to identify which subsystem will have the buffer pool specifications that are described in this section. All buffer pool specifications that follow the IMSID statement apply to the named IMSID. The buffer pool statements that apply to the specified subsystem end when the next IMSID statement or JOBNAME statement is encountered.

The **IMSID** statement must begin in column one. No embedded blanks are permitted in the statement. The specified value must be four characters or less.

Restriction: For an IMS control region (either DB/TM or DBCTL), you must utilize the **IMSID** statement, not the **JOBNAME** statement.

The syntax of the IMSID statement is:

```
IMSID=name
```

JOBNAME

The JOBNAME statement (or IMSID statement) must be included before any IMS DFSVSAMP buffer pool statements are specified.

Use the JOBNAME statement to identify the DL/I batch job name that will have the buffer pool specifications that follow. All buffer pool specifications that follow the **JOBNAME** statement apply to that job. The buffer pool statements that apply to the specified job name end when the next **IMSID** statement or **JOBNAME** statement is encountered. The following rules apply to the **JOBNAME** statement:

- It must begin in column one.
- No embedded blanks are permitted.
- The value that is specified must be 8 characters or less.

Requirement: For either a DB/TM or DBCTL IMS control region you must use the **IMSID** statement, not the **JOBNAME** statement.

The **JOBNAME** statement can be used only to specify an IMS DL/I batch job name.

The syntax of the **JOBNAME** statement is:

```
JOBNAME=jobname
```

SYSIN statements

You can include multiple BUFNO statements in the SYSIN DD of a coupling facility report job to customize the output that is produced.

The SYSIN DD statement is optional. Most SYSIN statements cannot be continued to multiple lines. The exception is the BUFNO statement, which can be continued on as many lines as are necessary.

The following is the syntax of the BUFNO statement:

```
BUFNO type= (number,number,...)
```

The BUFNO statement allows you to customize the subpool reports that display performance metrics for different numbers of buffers. The BUFNO statement allows you to replace the default number of buffers with numbers that you choose. You can specify different numbers of buffers for OSAM, VSAM, and VSAM with hiperspace, DEDB subpools, and in the OSAM cache. The *type=* parameter in the BUFNO statement specifies the type of subpools for which the numbers of buffers will be used. You can specify the following *type=* parameters in each BUFNO statement:

- OSAM=
- VSAM=

- HIPERSPACE= (or HS=)
- DEDB=
- CF=

In addition to the *type=* parameter, you must specify up to 45 different numbers of buffers. Each number can range from 3 to 32768 for OSAM and VSAM or from 3 to 999999 for Hiperspace(TM), DEDB, or CF types. The following figure shows the default numbers of buffers for OSAM, VSAM, HIPERSPACE, DEDB, and CF:

```
BUFNO OSAM=(4,8,16,32,64,128,192,256,384,512,768,1024,1536,
            2048,3072,4096,6144,8192,12288,16384,20480,24576,
            28672,32768)

BUFNO VSAM=(4,8,16,32,64,128,192,256,384,512,768,1024,1536,
            2048,3072,4096,6144,8192,12288,16384,20480,24576,
            28672,32768)

BUFNO HIPERSPACE=(4,8,16,32,64,128,192,256,384,512,768,1024,
                  1536,2048,3072,4096,6144,8192,12288,16384,
                  20480,24576,28672,32768,40960,49152,57344,
                  65535,81920,98302,131072,163840,196608,262144)

BUFNO DEDB=(4,8,16,32,64,128,192,256,384,512,768,1024,1536,
            2048,3072,4096,6144,8192,12288,16384,20480,24576,
            28672,32768)

BUFNO   CF=(64,128,192,256,384,512,768,1024,1536,2048,3072,4096,6144,
            8192,12288,16384,20480,24576,28672,32768,40960,49152,
            57344,65536,81920,98302)
```

Figure 37. Default numbers of buffers for various database types

You can specify as many of the five types of BUFNO statements as you need. For example, you can include only OSAM= and VSAM= statements and allow hiperspace, DEDB, and CF values to default to their default values. BUFNO statements are the only SYSIN statements that can be continued. Only columns 1 through 72 are treated as data. To continue a BUFNO statement, simply ensure that the last non-blank characters in columns 1 through 72 is a comma. A complete BUFNO statement is identified by a closing parenthesis.

Tip: The largest number for each type determines the number of buffer slots that are tracked during processing of the Buffer Pool report job. Specifying a very large value can significantly increase the job's processing time.

The following examples show valid BUFNO statements:

```
//SYSIN  DD *
BUFNO OSAM=(3,32,100,200,300,400,500,600,700,800,900,1000,2000,
            3000,4000,5000,6000,7000,8000,9000,10000,20000,
            30000)
BUFNO CF = (100, 200, 300,400,500,600,700,800,900,1000,
            2000,3000,4000,5000,6000,7000,8000,9000,
            10000,20000,30000,40000,50000,
            60000)
```

Figure 38. Examples of valid BUFNO statements

Coupling facility report print DD names

IMS Buffer Pool Analyzer's coupling facility analysis job produces a variety of reports that are written to different print DD names.

The following DD names are used for IMS Buffer Pool Analyzer coupling facility report output:

SYSPRINT
SYSPRINT includes diagnostic messages that describe any error conditions that are encountered during processing, and a list of the statements that are supplied in the SYSIN DD statement that are specified to control the options that are used to generate the report.

GTFRPT

GTFRPT includes the generalized trace facility (GTF) report that describes the GTF trace data sets that are supplied as input to this coupling facility report.

SPLRPT

SPLRPT contains the coupling facility reports. These reports document the current IMS OSAM cache coupling facility environment. If a DFSVSAMP DD statement is used to specify an alternate environment, the MDLRPT DD name contains reports that document the alternate environment. The coupling facility reports include OSAM coupling facility cache detail, OSAM summary subpool information, and detailed information about each OSAM subpool in the configuration.

MDLRPT

MDLRPT contains coupling facility reports for the alternate configuration that is requested in the DFSVSAMP DD statement. The coupling facility reports include OSAM coupling facility cache detail, OSAM summary subpool information, and detailed information about each OSAM subpool in the configuration.

DSNRPT

DSNRPT contains the data set report for OSAM data sets. These reports provide a list of the databases that are used during the trace interval, their DD names, and data set names. This optional report can be suppressed by including DSNRPT=NO in the SYSIN.

GTF trace report

The Generalized Trace Facility (GTF) report shows timestamps and the number of records for each GTF trace data set that is included in the report job.

The GTF trace report also provides information about lost records and lost blocks. Lost records occur when GTF buffers fill, and no buffer space is available for additional trace records. Lost blocks typically occur during write errors to the trace data set or during periods of very high trace activity. After trace buffers become available, GTF writes an error record that provides information about the number of trace records that were discarded for lack of buffer space. The report shows both the number of GTF error records and the total number of trace records that were discarded.

The following figure shows an example GTF trace report that includes two input GTF trace data sets. For each trace data set, the following types of information are shown:

- The data set names
- The subsystem names
- The start and end times of the trace data
- The elapsed time of each trace
- The number of records of each type

```
PAGE     1                              IMS BUFFER POOL ANALYZER 1.4.0
GTF TRACE REPORTS                                 COUPLING FACILITY

GTF TRACE DATA FOR SUBSYSTEM IMS IMB2 DDNAME=TRACE0    DSN=BPL.BPL140.TRACE

                      NUMBER                                      ELAPSED
   GTF RECORD TYPE    RECORDS       START TIME         END TIME   SECONDS
   _____    _____    _____    _____    _____
   CONTROL                54     2021.224  8:50:45  2021.224  8:55:47    302.3
   BFR POOL ANAL       5,868     2021.224  8:50:47  2021.224  8:55:47    300.0
   OSAM                1,384     2021.224  8:50:59  2021.224  8:55:44    285.5
   VSAM                  416     2021.224  8:50:47  2021.224  8:55:44    297.5
   UNKNOWN TYPE            0
   LOST RECORDS           0
   LOST BLOCKS            0
   TOTAL               7,722

GTF TRACE DATA FOR SUBSYSTEM IMS IMB1 DDNAME=TRACE3    DSN=BPL.TRACE.SVL5

                      NUMBER                                      ELAPSED
   GTF RECORD TYPE    RECORDS       START TIME         END TIME   SECONDS
   _____    _____    _____    _____    _____
   CONTROL               221     2021.074 14:26:12  2021.074 14:30:08    236.1
   BFR POOL ANAL       4,157     2021.074 14:26:14  2021.074 14:30:08    233.9
   OSAM                    0
   VSAM               52,891     2021.074 14:27:37  2021.074 14:29:12     95.8
   UNKNOWN TYPE            0
   LOST RECORDS           0
   LOST BLOCKS            0
   TOTAL              57,269
```

Figure 39. Sample report: GTF trace

The GTF trace report includes the following information fields:

CONTROL
GTF control records that are produced by GTF but are not used by IMS Buffer Pool Analyzer.

BFR POOL ANAL
Records that are generated by IMS Buffer Pool Analyzer. Records can include information about subpool definitions, database data set information, control records, or DEDB lookaside buffer activity.

OSAM/VSAM
OSAM trace records that are produced by the IMS Buffer Pool Analyzer or VSAM trace records that are produced by VSAM processing.

UNKNOWN TYPE
GTF trace records with an event ID that does not match any record types that IMS Buffer Pool Analyzer processes (440 - Buffer Pool Analyzer records, F4F - IMS OSAM records, and F61 - VSAM records). There should be no records in this category. If the report indicates the presence of unknown records, investigate the event ID that is associated with the records, determine the source of the record type, and disable recording the records. For assistance, contact IBM Software Support.

LOST RECORDS
Records that were discarded because of insufficient buffer space. The number of records shows how many occurrences of lost records were recorded. When records are lost, a value other than zero is present and the actual number of records that were discarded is also listed.

LOST BLOCKS
The number of blocks that were discarded because of insufficient buffer space. The number of records shows how many occurrences of lost blocks were recorded. When blocks are lost, a value other than zero is present and the actual number of blocks that were discarded is also listed.

OSAM data set information reports

The OSAM data set information reports display the data sets that are associated with the performance reports.

The coupling facility version of the OSAM data set report shows only the OSAM data sets that are defined in cached subpools.

The OSAM data set information reports are optional. You can suppress these reports by supplying a DSNRPT=NO statement in the SYSIN DD statement.

OSAM data set information reports provide information about data set names that can be referenced by the batch jobs and subsystems. Only those data sets that can be used are displayed:

- For batch jobs, the list includes every DBD in the PSB that has a database data set that is either present in the job's JCL, or available for dynamic allocation by DFSMDA members or from information in the RECON data sets for Fast Path and HALDB databases.
- For online systems, the list includes data set names for every database data set that has a valid ACBLIB member, including data set names that are available from the JCL, DFSMDA members, or data set names that are available from RECON definitions.

The following figure is an example OSAM data set information report. Information about block size is presented only for data sets that were active during the trace interval.

```
PAGE     1                                IMS BUFFER POOL ANALYZER 1.4.0
OSAM DATA SETS IN CF CACHED SUBPOOLS                       COUPLING FACILITY

         DCB   DB    MASTER    BLOCK
DBDNAME  NUM   TYPE  HALDBD    SIZE   DDNAME    DSNAME
_____  ___   ____  _____  _____  _____  _____
CUSTD4   1     DL/I  --------  8,192  CUSTD4    BPA1.CUSTD4.HIDAM
PDB1001  A     PART  CUSTOMRA  8,192  PDB1001A  BPA1.DBD1.A00001
PDB1002  A     PART  CUSTOMRA  8,192  PDB1002A  BPA1.DBD1.A00002
PDB1003  A     PART  CUSTOMRA  8,192  PDB1003A  BPA1.DBD1.A00003
PDB1004  A     PART  CUSTOMRA  8,192  PDB1004A  BPA1.DBD1.A00004
PDB1005  A     PART  CUSTOMRA  8,192  PDB1005A  BPA1.DBD1.A00005
PDB1006  A     PART  CUSTOMRA  8,192  PDB1006A  BPA1.DBD1.A00006
PDB1007  A     PART  CUSTOMRA  8,192  PDB1007A  BPA1.DBD1.A00007
PDB1008  A     PART  CUSTOMRA  8,192  PDB1008A  BPA1.DBD1.A00008
```

Figure 40. Sample report: OSAM data set information

The OSAM data set information report includes the following fields:

DBDNAME
The database name.

DCB NUM
The DCB number within the database for this data set, or the HALDB partition ID.

DB TYPE
The type of database. DL/I indicates a full function database. PART indicates a partition of a HALDB.

MASTER HALDB
This field is blank unless the database is a HALDB partition. For partitioned databases, the master HALDB database name of which this partition is a member. This field is blank unless the database is a HALDB partition.

BLOCK SIZE
The data set block size. This block size does not reflect updates to the database that are specified in a CHANGE statement included in the SYSIN statements.

DDNAME
The DD name that is used by this database data set.

DSNAME
The data set name of the database data set.

Coupling facility cache reports

IMS Buffer Pool Analyzer can produce either one or two sets of coupling facility cache reports.

Several types of reports are created and the first of these reports are written to the DD name SPLRPT document in the IMS coupling facility cache configuration that was in place at the time trace data was gathered. The second set of coupling facility cache reports are written if you requested a change to

the configuration through the DFSVSAMP DD statement, or a if you requested a change to the database access method or block size through a CHANGE statement in the SYSIN DD statement. The same type of reports are produced for both the actual configuration that is written to the SPLRPT DD output and to the model configuration that is written to the MDLRPT DD output.

OSAM coupling facility cache detail report

Several types of reports are created. The first reports in the SPLRPT and MDLRPT DD statements show detailed coupling facility cache configuration and usage information.

The following figure is an example of these OSAM coupling facility cache detail reports.

```
PAGE     1                                IMS BUFFER POOL ANALYZER 1.4.0
OSAM COUPLING FACILITY CACHE DETAIL                       COUPLING FACILITY

    CURRENT OSAM DATA CACHING CONFIGURATION:
        OSAM STRUCTURE NAME:          OSAMT1D
        COUPLING FACILITY NAME:         ICF02
        STRUCTURE SIZE:              16,384K
        DIRECTORY RATIO:                  50
        DIRECTORY ENTRIES:             5,282
        ELEMENT RATIO:                    50
        ELEMENT ENTRIES:               5,278

    CURRENT OSAM DATA CACHING PERFORMANCE:
        BUFFER REQUESTS:              121,959         323.90 PER SECOND
        BUFFER HITS:                   17,450          46.34 PER SECOND
        BUFFER READS:                 104,509         277.56 PER SECOND
        HIT RATIO:                       14.3%
        BUFFER LIFE:                      1.77  SECONDS

    CURRENT OSAM DATA CACHE SIZING INFORMATION:
        CFSIZER STRUCTURE SIZING INFORMATION:
            NUMBER OF OSAM BUFFERS:     1,400
            AVG BUF SIZE OF CACHED DATA:  8,192
        ADDITIONAL STRUCTURE SIZING INFORMATION:
            STRUCTURE OVERHEAD:     4,192,056                (APPROXIMATE)
            DIRECTORY SIZE:              309                (APPROXIMATE)
            ELEMENT SIZE:               2048
            NUMBER OF OSAM BUFFERS:     1,400
                NON-CACHED BUFFERS:     1,200
                CACHED BUFFERS:          200
```

Figure 41. Sample report: OSAM coupling facility cache detail

The OSAM coupling facility cache detail reports include the following fields:

OSAM STRUCTURE NAME
    The name of the OSAM cache structure, as it is specified in the DFSVSM*xx* member of PROCLIB.

COUPLING FACILITY NAME
    The name of the coupling facility in which the OSAM cache structure was allocated.

STRUCTURE SIZE
    The size of the OSAM cache structure.

DIRECTORY RATIO
    The directory ratio. For DD name SPLRPT, the directory ratio is obtained from IMS control blocks. For DD name MDLRPT, the directory ratio is taken from DFSVSAMP if it is supplied in the report job; otherwise, the value that is obtained from IMS control blocks is used.

DIRECTORY ENTRIES
    The approximate number of directory entries in the structure. This value is based on the level of software that is running in the coupling facility when the trace is taken.

ELEMENT RATIO
    The element ratio. For DD name SPLRPT, the element ratio is obtained from IMS control blocks. For DD name MDLRPT, the element ratio is taken from DFSVSAMP if it is supplied in the report job; otherwise, the value that is obtained from the IMS control blocks is used.

ELEMENT ENTRIES
The approximate number of elements in the structure. This value is based on the level of software that is running in the coupling facility when the trace is taken.

BUFFER REQUESTS
The number of times, and the number of times per second, that one of the subsystems that is sharing the OSAM cache structure needed a database block that was not present in that subsystem's local buffer pools.

BUFFER HITS
The number of times, and the number of times per second, that a database block that is required by a subsystem was available in the OSAM cache structure.

BUFFER READS
The number of times, and the number of times per second, that a database block that is required by a subsystem was not available in the OSAM cache structure.

HIT RATIO
The percentage of times that a required database block was found in the OSAM cache structure.

BUFFER LIFE
The length of time that a database block remains in the cache structure without being referenced.

NUMBER OF OSAM BUFFERS
The number of both cached and non-cached buffers that are present in the trace data. For the Model report, in the MDLRPT DD statement, this number includes all of the OSAM buffers that are defined in all the subsystems in trace data or in the DFSVSAMP DD statement. This number is presented as an input to the CFSizer tool.

This number of OSAM buffers represents the maximum number of buffers that are used by all of the subsystems that are participating in the IMSplex.

AVG BUF SIZE OF CACHED DATA
The average buffer size of blocks that are written to the coupling facility structure. The CFSizer tool uses this number to determine the number of elements that are required to fit an average data base block that is to be cached.

STRUCTURE OVERHEAD
The amount of storage that is used in the coupling facility structure for internal control blocks and data. The value that is listed is an approximate value that is determined during the trace process. The value varies depending on the coupling facility software maintenance level.

DIRECTORY SIZE
The amount of storage that is required for each directory entry in the coupling facility structure. A directory entry is used to track the status of a single database buffer. One directory entry is required for each OSAM buffer that is used by an IMS subsystem in the IMSplex, and for each buffer that is resident in the coupling facility cache.

ELEMENT SIZE
The size of each element in the coupling facility structure. Elements are used to store the actual database buffer. Each element in the IMS OSAM cache is always 2048 bytes. From 1 to 16 elements are used to store each buffer, depending on the size of each buffer.

NUMBER OF OSAM BUFFERS
The number of OSAM buffers that are used by IMS subsystems and included in the trace data or specified in the DFSVSAMP input. This value is broken down into cached and non-cached buffers.

Requirement: A directory entry is required for each OSAM buffer in each IMS subsystem in the IMSplex, either cached or non-cached. Elements are used only for cached buffers.

Model of varying cache sizes report

The model of varying cache sizes displays projections for performance metrics for differing amounts of storage that can be allocated to the cache structure.

The number of buffers and structure size indicate the amount of coupling facility storage that is used by the structure. For each structure size, the recommended ratios for directory entries (DIR) and element entries (ELE) are shown. The hit ratio, reads per second, and buffer life show the effect of changing the structure size to the value on each line of the report.

The following figure is an example of the Model of varying cache sizes for actual cache configuration report.

```
MODEL OF VARYING CACHE SIZES FOR ACTUAL CACHE CONFIGURATION:

                                ------  PROJECTIONS  ------
          NUMBER   STRUCTURE   RATIOS    HIT   READS PER    BUFFER
          BUFFERS  SIZE (KB)  DIR ELE   RATIO   SECOND       LIFE

          --------  ---------  --- ---  ------  ---------  ---------
               0      4,520      1   0    0.0%     12.30       0.00
              64      4,988     79  16   25.0%      9.22       6.93
             128      5,520     83  32   29.7%      8.64      14.81
             192      6,052     29  16   33.4%      8.18      23.46
             256      6,584     91  64   37.4%      7.69      33.30
             384      7,644     33  32   44.8%      6.79      56.58
             512      8,708    107 128   53.1%      5.77      88.78
             768     10,832     41  64   65.1%      4.28     179.41
           1,024     12,960    139 256   71.5%      3.50     292.63
           1,536     17,208     57 128   71.6%      3.49     440.62
           2,048     21,460    203 512   71.6%      3.49     587.50
           3,072     29,960     89 256   71.6%      3.49     881.25
           4,096     38,460    165 512   71.6%      3.49    1175.00
```

Figure 42. Sample report: Model of varying cache sizes for actual cache configuration

Databases with activity in the cache report

The Databases with activity in the cache report displays the names of the databases that have had activity in the OSAM cache subpool structure.

The following sample report also shows activity rates (requests and reads per second) and hit ratios for each of the database data sets with database blocks in the cache.

```
DATABASES WITH ACTIVITY IN THIS SUBPOOL:

           DCB  HALDB     BLOCK    HIT    REQUESTS      READS
  DBDNAME  NUM  MASTER    SIZE    RATIO  PER SECOND  PER SECOND

  -------  ---  -------  -------  ------  ----------  ----------
  PDB2001   A   CUSTOMRB   8,192   14.3%     153.24      131.19
  PDB1034   A   CUSTOMRA   8,192   14.5%       6.79        5.81
  PDB1009   A   CUSTOMRA   8,192   15.6%       4.20        3.54
  PDB1040   A   CUSTOMRA   8,192   14.8%       4.00        3.41
  PDB1035   A   CUSTOMRA   8,192   15.9%       3.90        3.28
  PDB1026   A   CUSTOMRA   8,192   13.0%       3.84        3.34
  PDB1036   A   CUSTOMRA   8,192   13.1%       3.77        3.28
  PDB1031   A   CUSTOMRA   8,192   14.6%       3.68        3.14
```

Figure 43. Sample report: Databases with activity in the cache

Database subpool summary report

The database subpool summary report displays all of the OSAM database buffer subpools that are present in each subsystem.

The Database subpool summary shows, both the cached subpools and the uncached subpools because uncached subpools still require directory entries in the OSAM cache structure for cross invalidation processing.

The following figure shows an example of a database subpool summary report.

```
DATABASE SUBPOOL SUMMARY
                              REQUESTS  READS
              BUFFER SUB-  C  NUMBER    PER      PER     HIT  BUFFER  MARGINAL
      SUBSYS   SIZE  POOL  F  BUFFERS  SECOND   SECOND  RATIO  LIFE   REDUCTION
     ───────  ─────  ────  ─  ───────  ───────  ──────  ─────  ────── ─────────
     IMS IMB1  4,096 004K     1,000    ---------- SUBPOOL NOT CACHED  --------
     IMS IMB1  8,192 008K  A    200     5,071      324  93.6%     0        68
     IMS IMB1 12,288 012K  A    200    ----------  SUBPOOL NOT USED   -------
     IMS IMB2  4,096 004K  A  1,000    ----------  SUBPOOL NOT USED   -------
     IMS IMB2  8,192 008K  C    200         4        0  94.1%   929         0
     IMS IMB2 12,288 012K       200    ----------  SUBPOOL NOT CACHED  -------
              *TOTAL         2,800     5,075      324  93.6%
```

Figure 44. Sample report: Database subpool summary

Performance metrics are gathered only for subpools that are cached. Uncached subpools are shown with the message SUBPOOL NOT CACHED in place of the performance information. The subpool summary report sample shows the following information:

SUBSYS
The name of the subsystem that was obtained from trace files or from IMS buffer pool statements in the DFSVSAMP DD statement in the report job. For an IMS control region, the name is IMS followed by the IMSID. For a DL/I batch job, SUBSYS is the job name.

BUFFER SIZE
The buffer size of the IMS subpool.

SUB-POOL
The subpool name that is defined in that IMS subsystem.

CF
The coupling facility cache option.

- Blank indicates that the CF is not cached.
- C- indicates that the changed database blocks are cached.
- A- indicates that all of the database blocks are cached.

NUMBER BUFFERS
The number of buffers that are allocated to the subpool.

REQUESTS PER SECOND
The number of database block requests per second for the subsystem's subpool.

READS PER SECOND
The number of buffer misses when the buffer is not found in the subpool. The number of reads per second equates to the number of OSAM cache requests per second.

HIT RATIO
The percentage of buffer requests that are met from the subsystem's buffer pool.

BUFFER LIFE
The average duration that an unreferenced buffer remains in the subpool.

MARGINAL REDUCTION
The marginal reduction value indicates how effectively each subpool could use additional storage. A large number indicates a subpool that could have a lower I/O rate if it had more storage. A subpool with a marginal reduction of 0 indicates that adding storage would not improve the performance significantly.

Subpool detail report

IMS Buffer Pool Analyzer provides detailed information about performance for each subpool that is defined.

The subpool detail report includes the following four sections:

- Actual subpool configuration
- Actual subpool performance

- Model of varying buffer pool sizes for actual subpool configuration
- Databases with activity in this subpool

The following figure shows an example subpool detail report:

```
PAGE   160                        IMS BUFFER POOL ANALYZER 1.4.0
 VSAM 4,096 SUBPOOL DFLT DETAIL          IMS ZICA BUFFER POOL REPORT

  ACTUAL SUBPOOL CONFIGURATION:
   IMSID/JOBNAME:          IMS ZICA
   BUFFER SIZE:            4,096 VSAM INDEX
   BUFFER POOL ID:           DFLT
   NUMBER BUFFERS:         4,000

  ACTUAL SUBPOOL PERFORMANCE:
   BUFFER REQUESTS:        2,661,697   8,837.27 PER SECOND
   BUFFER POOL HITS:       2,365,445   7,853.66 PER SECOND  88.8%
   BUFFER READS:            296,252     983.61 PER SECOND  11.1%
   HIT RATIO:                 88.8%
   BUFFER LIFE:               4.06 SECONDS
   MARGINAL REDUCTION:         124

  MODEL OF VARYING BUFFER POOL SIZES FOR ACTUAL SUBPOOL CONFIGURATION:

                    ----------------- PROJECTIONS ------------------
      NUMBER       POOL      HIT      READS PER     BUFFER        MARGINAL
      BUFFERS      SIZE      RATIO     SECOND        LIFE         REDUCTION

      -------    --------   ------    --------     --------      ----------
            4       16K     18.2%     7,221.91       0.00            41,512
          100      400K     68.3%     2,793.87       0.03             4,281
          200      800K     73.7%     2,318.18       0.08             1,867
          300    1,200K     76.1%     2,110.66       0.14             1,230
          400    1,600K     77.6%     1,973.95       0.20               922
          500    2,000K     78.8%     1,871.40       0.26               720
          600    2,400K     79.7%     1,791.32       0.33               590
          700    2,800K     80.4%     1,725.76       0.40               502
          800    3,200K     81.1%     1,669.96       0.47               429
          900    3,600K     81.6%     1,622.20       0.55               378
        1,000    4,000K     82.1%     1,580.19       0.63               254
        2,000    8,000K     85.3%     1,297.94       1.54               158
        3,000   12,000K     87.3%     1,122.22       2.67               124
        4,000   16,000K     88.8%       983.57       4.06                 0
        5,000   20,000K     88.8%       983.48       5.08                 0
        6,000   24,000K     88.8%       983.48       6.10                 0
        7,000   28,000K     88.8%       983.48       7.11                 0
        8,000   32,000K     88.8%       983.48       8.13                 0
        9,000   36,000K     88.8%       983.48       9.15                 0
       10,000   40,000K     88.8%       983.48      10.16                 0
       20,000   80,000K     88.8%       983.48      20.33                 0
       30,000  120,000K     88.8%       983.48      30.50                 0
       40,000  160,000K     88.8%       983.48      40.67                 0

  DATABASES WITH ACTIVITY IN THIS SUBPOOL:
               DCB  VSAM   HALDB     CI    HIT     REQUESTS      READS
      DBDNAME   NUM  COMP   MASTER    SIZE  RATIO  PER SECOND  PER SECOND

      -------   ---  -----  ------    -----  -----  ----------  ----------
      PG15IDF    X   INDEX  PGZIDF    4,096 79.2%      102.28       21.22
      NIZXCI1X   3   INDEX            4,096 69.2%       58.92       18.13
      NIZXCI2X   3   INDEX            4,096 84.6%       84.59       12.94
      NIZXFDRX   1   INDEX            4,096 57.5%       29.99       12.73
      PGZPMTX    1   INDEX            4,096 97.2%      463.02       12.52
```

Figure 45. Sample report: Subpool detail

```
VSAM 4,096 SUBPOOL VDEF DETAIL                       IMS IMB1 BUFFER POOL REPORT

   ACTUAL SUBPOOL CONFIGURATION:
      IMSID/JOBNAME:               IMS IMB1
      BUFFER SIZE:                    4,096  VSAM    BOTH
      BUFFER POOL ID:                 VDEF
      NUMBER BUFFERS:                 1,600
      NUMBER HIPERSPACE BUFFERS:        400

   ACTUAL SUBPOOL PERFORMANCE (BASE):
      BUFFER REQUESTS:              141,446          156.74 PER SECOND
      BUFFER POOL HITS:            140,949          156.19 PER SECOND      99.6%
      HIT RATIO:                    99.6%
      BUFFER LIFE:                 2905.17  SECONDS
      MARGINAL REDUCTION:               8

   ACTUAL SUBPOOL PERFORMANCE (HIPERSPACE):
      BUFFER REQUESTS:                 497            0.55 PER SECOND
      HIPERSPACE HITS:                   0            0.00 PER SECOND       0.0%
      HIT RATIO:                     0.0%
      BUFFER LIFE:                 3631.46  SECONDS
      MARGINAL REDUCTION:               0

   ACTUAL SUBPOOL PERFORMANCE (BASE+HIPERSPACE):
      BUFFER REQUESTS:              141,446          156.74 PER SECOND
      BUFFER POOL HITS:            140,949          156.19 PER SECOND      99.6%
      HIPERSPACE HITS:                   0            0.00 PER SECOND       0.0%
      DASD READS:                      497            0.55 PER SECOND       0.3%
      HIT RATIO:                    99.6%
      BUFFER LIFE:                 3631.46  SECONDS
      MARGINAL REDUCTION:               0
```

Figure 46. Sample report: Subpool detail with hiperspace buffers

The first section of the subpool detail report provides information about the actual subpool configuration. In the preceding example, the actual subpool configuration is displayed, which indicates that this is the IMS configuration that is used during the trace and that is written to the SPLRPT DD data set. The subpool configuration for a model environment that is written to the MDLRPT DD data set would display the following message SUBPOOL CONFIGURATION FROM DFSVSAMP. The subpool configuration information identifies the IMSID or DL/I batch job name that was traced and the buffer size, buffer pool ID, and number of buffers in the subpool.

The second section of the subpool detail report provides information about the actual subpool performance. The preceding example displays the actual subpool performance that is based on the trace data. If a model configuration is in place, the report message in the MDLRPT DD data set would display DFSVSAMP SUBPOOL CONFIGURATION PERFORMANCE, and the performance metrics would reflect the subpool configuration as it was requested in DFSVSAMP, along with any CHANGE DBD statements that were specified on the SYSIN DD statement.

The performance metrics that are listed in the actual subpool performance section show the number of buffer requests, buffer pool hits, DASD buffer reads, the per-second rates of these metrics, and the percentages of requests that resulted in buffer pool hits, DASD reads, then the hit ratio, buffer life, and marginal reduction value.

The third section of the subpool detail report shows projections for subpool performance metrics for different amounts of storage that are allocated to the subpool, and it displays the resulting performance for different numbers of buffers that might be allocated to the subpool. The number of buffers on each line of the report is determined based on the default numbers of buffers for a subpool. Or, you can specify the number of buffers on each line using **BUFNO** statements in the SYSIN DD statement for the report job.

IMS Buffer Pool Analyzer uses the information in performance projections to review subpool performance and determine whether performance can be improved by increasing or decreasing the number of buffers that are allocated to each subpool.

The fourth section of the subpool detail report lists all of the database data sets that had activity in the subpool, and performance metrics are shown for each data set. The data sets are presented in order of the highest through the lowest number of reads-per-second, which means that the highest I/O data sets are displayed first in each subpool's list of active databases.

Modeling an alternate configuration

IMS Buffer Pool Analyzer allows you to model buffer pool configuration changes and produce the subpool reports that are based on this alternate configuration.

You can add, remove, or change subpool definitions, change the association of database data sets with subpools, and change the block or CI size of a data set.

Changing the access method or block size of a database

To model the impact of a change to the access method (OSAM to VSAM, or VSAM to OSAM), or to the block size of a database data set use the coupling facility report.

About this task

If you plan to make a database change, the coupling facility report lets you review the impact of the change on your OSAM cache configuration and to plan changes to your configuration before you actually implement the database change.

Procedure

To model the impact of a database change, perform the following steps:

1. Supply one or more CHANGE DBD statements in the SYSIN DD statement of the coupling facility report job.
2. Run the coupling facility report job.

Results

The coupling facility report produces reports for the supplied trace data that document the performance of the old database configuration (in the SPLRPT DD statement) and the performance of the new database configuration (in the MDLRPT DD statement).

Changes to the IMS buffer pool configuration

You can specify updated buffer pool configurations for any or all of the IMS subsystems in your trace files.

Changes to an IMS subsystem's buffer pool configuration can directly impact the performance of the OSAM cache. The MDLRPT DD output will contain the alternate configuration reports that reflect the updated buffer pools.

Changing the OSAM cache configuration

You can update the directory and element ratios by making changes to the buffer pool definition CFNAMES statement.

Procedure

To update the directory and element ratios, perform the following steps:

1. Modify and include the updated CFNAMES statement in the report job DFSVSAMP DD statements so that these statements include the alternate configuration model, the updated directory, and element ratios.
2. Alternatively, you can change the OSAM cache structure size and model the impact of the changed structure size.

Alternate configuration reports

Each time that the coupling facility alternate configuration report job is run, it can produce either one or two sets of reports.

The first set of alternate configuration reports reflects the actual subpool and OSAM cache configurations, and it is written to the SPLRPT DD output.

The second set of alternate configuration reports, which are referred to as *model reports*, are produced only if you request a configuration change through the use of a CHANGE DBD statement in the SYSIN DD statement or by specifying any statements in the DFSVSAMP DD statement. The model alternate configuration reports are always written to the MDLRPT DD output.

Tuning the OSAM coupling facility cache

Because the it is not the same as an OSAM subpool, there are several methods for tuning the OSAM coupling facility cache.

Increasing the size of a subpool results in longer search time for the subpool (because there are more buffers) and increased storage utilization. No additional I/O is associated with the additional buffers. When a buffer hit occurs in the added buffers, a read I/O is completely eliminated.

Caching requires that blocks be written to the cache when they are read (or updated, depending on the cache option that you select). Writing data to the OSAM coupling facility takes CPU time and I/O. In addition, when a buffer hit occurs an I/O is still required to retrieve the buffer from the coupling facility.

The essential point is that the OSAM cache is not simply an extension to an OSAM subpool. But, because the OSAM coupling facility cache is a shared resource among all IMS subsystems in the data sharing environment, it is not subject to buffer invalidation. So, in some cases, the OSAM coupling facility cache can be more efficient than the OSAM subpools.

There are only three practical tuning options for the OSAM coupling facility cache. These tuning options are as follow:

1. Select the database data sets that are to be cached in the coupling facility structure.
2. Select the directory and element ratios.
3. Select the OSAM cache structure size.

Selecting database data sets to be cached

Choose the database data sets DBDS that you want to isolate and cache.

Before you begin
Approach enabling caching very seriously, especially for an initial implementation.

Procedure

1. Choose one or a few heavy-update-activity shared OSAM database data sets (DBDS) as candidates for caching.
2. After you select the DBDSs that you want to cache, isolate those data sets to a separate subpool to be cached.
3. Create new DFSVSM*xx* members for each IMS subsystem in the data sharing environment, without implementing the new member, that assigns the selected DBDSs to the isolated subpool.

Selecting the directory and element ratios

After you have selected or updated the list of DBDSs to be cached, created, and updated DFSVSM*xx* members, you can use the coupling facility report to determine appropriate values for the directory ratio and element ratio for the size of your coupling facility structure.

Procedure

To create DFSVSAMP input statements for the coupling facility report job, perform the following steps:

1. Copy the updated DFSVSM*xx* members that you have created to reflect the updated subpool configuration.
2. Add the IMSID= and JOBNAME= statements between the different DFSVSM*xx* members.
3. Place updated DFSVSM*xx* members in the DFSVSAMP file even if you have no trace data for some of the subsystems.
4. If you ever run IMS DL/I batch jobs, also create DFSVSAMP input for batch jobs, so that the number of directory entries in the structure is sufficient.
   a) Create input for JOBNAME=BATCH1 through BATCH*xxx*, depending on the maximum number of DL/I batch jobs that can execute concurrently in the sysplex.
   b) The subpools for every batch job can be the same; however, ensure that each batch definition has the maximum number of OSAM buffers that any DL/I batch job in your environment has.
5. If you want to leave additional directory entries for expansion of online buffer pools, create another batch job in the DFSVSAMP input stream with the number of OSAM buffers that you want to reserve for expansion.
6. When you have the DFSVSAMP input customized for your new environment, run the coupling facility report with the DFSVSAMP input.

   The coupling facility report job should have trace files from all of the IMS subsystems in the data sharing environment.
7. To produce the correct mix of workloads so that the average DBDS block size is correct, follow the recommended directory ratio and element ratio values that are displayed in the report output that is defined in the MDLRPT DD statement.

Selecting the OSAM cache structure size

Review the output of a coupling facility report job and locate the OSAM cache structure detail report in the MDLRPT DD statement specification.

About this task

The cache structure configuration and performance information of the OSAM cache structure detail report provide information about the models of hit ratios with varying buffer sizes. You can use this information as a guide to determine the amount of storage (structure size) that you want to allocate to the structure.

Procedure

After you determine the structure size, update the CFRM policy with the new value and install the updated policy.

Chapter 8. Troubleshooting

IMS Buffer Pool Analyzer issues messages and codes that can help you diagnose and correct problems that you experience with the product.

Topics:

-
-
-

Runtime messages (BPL)

This reference section provides detailed information about the error messages that are issued by IMS Buffer Pool Analyzer. The technical information in this section can help you troubleshoot and diagnose IMS Buffer Pool Analyzer problems.

Message format

IMS Buffer Pool Analyzer messages adhere to the following format:

```
BPLnnnx
```

In the message text,

BPL
: Indicates that the message was issued by IMS Buffer Pool Analyzer

nnn
: Indicates the message identification number

x
: Indicates the severity of the message:

> **A**
> : Indicates that operator intervention is required before processing can continue.
>
> **E**
> : Indicates that an error occurred, which might or might not require operator intervention.
>
> **I**
> : Indicates that the message is informational only.
>
> **W**
> : Indicates that the message is a warning to alert you to a possible error condition.

Each message also includes the following information:

Explanation
: Explains what the message text means, why it occurred, and what its variables represent.

System action
: Explains what the system will do in response to the event that triggered this message.

User response
: Describes whether a response is necessary, what the appropriate response is, and how the response will affect the system or program.

BPL000I   *modid* **debug information**

Explanation

This is an informational debug message that can be written when **OPT=DEBUG** is specified in the BPLTRACE start command. *modid* is a four-character module ID

indicating the module that issued the debug message. There are 26 variations of trace messages, which are of value when investigating a software problem with the data gathering process. For an IMS control region with many databases, the number of trace messages can be very large. Because these messages are written to MVS SYSLOG, it is recommended that you use the DEBUG option only as directed by the IBM Software Support.

System action

None.

User response

To disable debug messages, do not include the **OPT=DEBUG** parameter in the BPLTRACE start command or in the BPLTRACE PROC.

BPL101E   **OPEN FAILED FOR IMSACB DD** *ddname*

Explanation

An attempt was made to open the IMS control region ACBLIB data set(s) that was allocated to *ddname*. The open was not successful.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Review the MVS SYSLOG for error messages associated with the open of the library. If no messages are found to indicate the cause of the problem, contact IBM Software Support for assistance.

BPL102E   **ACBLIB BLDL FOR DBD MEMBERS FAILED RC=***nn* **REASON=***nn*

Explanation

A BLDL for members of the ACBLIB data set(s) was unsuccessful. The return code and reason codes from the BLDL macro are shown in the message text.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Review the MVS SYSLOG for error messages associated with the error. If no messages are found

to indicate the cause of the problem, contact IBM Software Support for assistance.

BPL103E   **BLDL FAILED ACBLIB MEMBER** *dbdname*

Explanation

The indicated DBD member of ACBLIB was not found.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Verify that an online change for the ACBLIB data sets was not in progress at the time of the error. Contact IBM Software Support for assistance.

BPL104E   **INVALID ACBLIB DIRECTORY DATA FOR DBD** *dbdname*

Explanation

ACBLIB member, *dbdname*, did not contain valid PDS directory information for an IMS DBD member of ACBLIB.

System action

IMS Buffer Pool Analyzer ignores the database with the error and continues processing.

User response

Contact IBM Software Support for assistance.

BPL105E   **ACBLIB MEMBER** *dbdname* **NOT FOUND**

Explanation

A FIND macro for ACBLIB member *dbdname* failed.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Review the MVS SYSLOG for error messages associated with the error. If no messages are found to indicate the cause of the problem, contact IBM Software Support for assistance.

BPL106E   **DBD VALIDATION FAILED FOR ACBLIB MEMBER** *dbdname (r)*

Explanation

Validation of ACBLIB member *dbdname* failed. The reason for the failure is indicated by the reason code *r* at the end of the message.

System action

IMS Buffer Pool Analyzer ignores the database with the error and continues processing.

User response

Verify that the ACBLIB member named in the message is valid. Contact IBM Software Support for assistance.

BPL107E	ACBLIB MEMBER *dbdname* HAS NO DATA

Explanation

No data was found while reading ACBLIB member *dbdname*.

System action

IMS Buffer Pool Analyzer ignores the database with the error and continues processing.

User response

Make sure that an online change was not in progress at the time of the error. Review the MVS SYSLOG for error messages associated with the error. If no messages are found to indicate the cause of the problem, contact IBM Software Support for assistance.

BPL108E	ACBLIB MEMBER *dbdname* HAS INCOMPLETE DATA

Explanation

The ACBLIB member *dbdname* did not contain the entire DBD.

System action

IMS Buffer Pool Analyzer ignores the database with the error and continues processing.

User response

Ensure that an online change was not in progress at the time of the error. Review the MVS SYSLOG for error messages associated with the error. If no messages are found to indicate the cause of the problem, contact IBM Software Support for assistance.

BPL109E	INVALID PARAMETER INFORMATION PASSED TO *module_name*

Explanation

The list of database names that should be provided to *module_name* BPLACBR or BPLMACB was missing.

System action

IMS Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL123E	CONCATENATED DATASET FOR DDN=*ddname* DBD=*dbdname*

Explanation

The TIOT in the target address space indicates that there are multiple data sets associated with the indicated *ddname*. Database *dbdname* should have only one data set allocated to the *ddname*, not a concatenation of data sets.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Review allocation for the indicated *ddname*. Contact IBM Software Support for assistance.

BPL124I	JOBNAME *name* BEING TRACED HAS PSB *psbname*

Explanation

Trace data is being gathered for the indicated DL/I or DBB batch job. The name of the PSB being used is also shown.

System action

None.

User response

This is an informational message that requires no response.

BPL125E	NO DATABASE DATA SETS FOUND IN TARGET JOB

Explanation

A review of the DLI or DBB batch job environment showed that there were no database data sets allocated to the job and that there would, therefore, be no database I/O activity to trace.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Determine whether the batch job accesses the IMS database. If the batch job performs any IMS database activity, contact IBM Software Support for assistance.

BPL126E	**ERROR LOCATING ALL DATABASE DIRECTORIES**

Explanation

An error occurred processing the database directory information for a batch DLI or DBB job.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL127E	**ERROR RETRIEVING PSB INFORMATION FOR JOB -** *xxxx*

Explanation

Buffer Pool Analyzer encountered an unexpected condition while locating the PSB in use for a DL/I batch job.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Verify that the DL/I batch job did not complete when the error occurred. The reason for the unexpected condition is indicated in the message text. Contact IBM Software Support for assistance.

BPL131E	**LOAD FAILED FOR MODULE IGGCSI00 RC=***nn* **ABEND=***nnn*

Explanation

The BPLTRACE task tried to load module IGGCSI00, the catalog search interface module. The load failed with the indicated abend code and return code.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Verify that the catalog search interface module (IGGCSI00) is installed on the system and is available to the BPLTRACE started task.

BPL133E	**CATALOG SEARCH INTERFACE RETURN CODE** *nn* **REASON CODE** *nnnnnn*

Explanation

A catalog request made to the Catalog Search Interface failed with the indicated return code and reason code.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

The CSI return code and reason codes are documented in *DFSMS/MVS Managing Catalogs*. Contact IBM Software Support for assistance.

BPL134E	**CATALOG SEARCH INTERFACE DATA ERROR REASON CODE** *nnnnnnnn*

Explanation

The CSICRETN field indicated a failure in the Catalog Search Interface module.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

The CSI return code and reason codes are documented in *DFSMS/MVS Managing Catalogs*. Contact IBM Software Support for assistance.

BPL135E   CATALOG SEARCH INTERFACE FIELD ERROR REASON CODE *nnnnnnnn*

Explanation

The CSIEFLAG field indicated a failure in the catalog search interface module for a field.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

The CSI return code and reason codes are documented in *DFSMS/MVS Managing Catalogs*. Contact IBM Software Support for assistance.

BPL136E   UNKNOWN CATALOG ENTRY TYPE RETURNED

Explanation

The CSIETYPE field returned by the Catalog Search Interface module contained an unknown entry type.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL137E   CATALOG SEARCH (IGGCSI00) ERROR - *reason*

Explanation

The CSICFLAG field indicated a failure in the Catalog Search Interface module. The CSI return code and reason codes are documented in *DFSMS/MVS Managing Catalogs*.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL138E   UNEXPECTED CATALOG ENTRY TYPE FOUND – *type*

Explanation

The Catalog Search Interface module returned an unexpected data set type for a database data set.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL139E   INCORRECT CATALOG ENTRY - DBD=*dbdname* **DDN=***ddname*

Explanation

These messages follow the preceding BPL139E message:

```
BPL139E DSN=dataset-name
BPL139E RETURNED DSN=dataset-name
```

The catalog search interface module returned a data set name that was not the same as the data set name requested. The data set names requested and returned are shown in the messages that are issued.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

List the catalog entry for the requested DSN shown in the message to verify that it contains catalog information for the expected data set. Contact IBM Software Support for assistance.

BPL140E   VSAM CATALOG INFO MISSING - DBD=*dbdname* **DDN=***ddname*

Explanation

This message follows the preceding BPL140E message:

```
BPL140E DSN=dataset-name
```

The Catalog Search Interface did not return a value for the VSAMTYPE parameter.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

List the catalog entry for the requested DSN shown in the message to verify that it contains complete catalog information. Contact IBM Software Support for assistance.

BPL151E   **BPLDSNI PARM MISSING ADDRESS SPACE INDICATOR**

Explanation

Required parameter information was not provided to module BPLDSNI.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL152E   **BPLDSNI PARM HAS TOO MANY ADDRESS SPACES**

Explanation

Invalid parameter information was provided to module BPLDSNI.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL153E   **SETLOCK {OBTAIN | RELEASE} FAILED RC=*nn***

Explanation

An MVS SETLOCK request failed with the indicated return code. The type of SETLOCK request (OBTAIN or RELEASE) is shown in the message.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL154E   **{GETMAIN | FREEMAIN} FOR ECSA FAILED RC=*nn***

Explanation

A GETMAIN or FREEMAIN macro, as shown in the message, failed. The return code from the GETMAIN/FREEMAIN macro is shown in the message.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL155E   **SRB ERROR IEAMSCHD RC=*nn* COMP=*nn* CODE=*nn* REASON=*nn* (STATUS *nnnn*)**

Explanation

An IEAMSCHD macro failed with the indicated return code (register 15), SRBCOMP, SRBCODE, and SRB reason code fields. The SRBSTAT field contains program status information.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL156E   **IEAMSCHD FAILED RC=*nn***

Explanation

An IEAMSCHD macro failed with the indicated return code (register 15).

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL157E   **BPLSRB00 FAILED - SEE SYSLOG MESSAGE(S) (STATUS=*nnnn*)**

Explanation

A failure occurred during DSNAME processing.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Check the MVS SYSLOG for other BPL messages preceding this message. Contact IBM Software Support for assistance.

BPL158E　　**BPLSRB00 WAS SUCCESSFUL BUT NO DSNAMES WERE RETURNED (STATUS=*nnnn*)**

Explanation

A failure occurred during DSNAME processing.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL159E　　**DYNAMIC {ALLOCATION | CONCATENATION} FAILED RC=*nn* ERROR CODE=*nnnn* INFO CODE=*nnnn***

Explanation

A failure occurred during dynamic allocation or concatenation processing. The return code, error code, and information code are indicated in the message. This message follows the preceding BPL159I message: BPL159I DSN=*dataset-name*.

System action

IMS Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Review the reason for the dynamic allocation error. Contact IBM Software Support for assistance.

BPL159I　　**DYNAMIC {ALLOCATION | CONCATENATION} COMPLETE DDNAME=*ddname***
　　　　　　　or

DSN=*dataset-name*

Explanation

The indicated data set was dynamically allocated or concatenated to the BPLTRACE task.

User response

None. This message is informational.

BPL171E　　**DDNAME SEARCH FAILED - *control-block* ADDRESS WAS 0**

Explanation

A failure occurred during DSNAME processing.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL172E　　**SWAREQ FAILED FOR *control-block* RC=*nn***

Explanation

A SWAREQ macro failed for the named control block type. The return code is indicated in the message.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL173E　　**ALESERV {ADD | DEL} FAILED RC=*nn***

Explanation

An ALESERV macro failed. The return code and type of request (add or delete) are indicated in the message.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL181E   **GTF NOT ACTIVE OR NOT RECORDING E440 USER RECORDS**

Explanation

A GTRACE macro indicated that GTF was not recording the appropriate user trace records required for Buffer Pool Analyzer.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL182E   **GTRACE ERROR - RETURN CODE *nn***

Explanation

A GTRACE macro failed with the indicated return code.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Verify that the GTF trace started by the BPLTRACE task was running at the time of the error. Contact IBM Software Support for assistance.

BPL183E   **UNABLE TO DETERMINE THE ACTIVE ACBLIB DD**

Explanation

An error occurred determining the ddname of the active ACBLIB.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL184E   **UNABLE TO DETERMINE THE ACTIVE ACBLIB DATASET NAMES**

Explanation

The IMSACBA or IMSACBB ddname was not found in the TIOT of the target address space.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL185E   **ERROR LOCATING VALID DBD NAMES**

Explanation

The scan to find what DBDs were present in the IMS subsystem failed to find all the DBD entries.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL186E   **FATH PATH ERROR - *control-block* ADDRESS WAS 0**

Explanation

There was a valid DEDB database found in the database directory, but either the ESCD or the first DMCB address was not present.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL187E   **FAILED TO FIND ALL {DEDB | HALDB} DATABASES**

Explanation

The number of either Fast Path DEDB or HALDB partition databases was not the same as the number expected.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL188E **FIND FAILED FOR ACTIVE DEPENDENT REGIONS**

Explanation

During data gathering, the BPLTRACE task was unable to determine the z/OS address space IDs (ASIDs) of the IMS dependent regions.

System action

Buffer Pool Analyzer ends the BPLTRACE task.

User response

Contact IBM Software Support for assistance.

BPL189I **IMS CONFIGURATION STATUS:** *resource* **INFO COMPLETE**

Explanation

IMS Buffer Pool Analyzer has completed gathering information for the resource named in the message from the IMS configuration. The resource field indicates the status of the information gathering phase and will be one of the following: DBD NAME, ACBLIB, IMS DIRECTORY, ALLOC DSN, DFSMDA, RECON, or CATALOG.

ACBLIB is displayed if the information was obtained from the ACB library. IMS DIRECTORY is displayed if the information was obtained from the IMS directory when the IMS management of ACBs is enabled.

System action

None.

User response

None. This message provides the status of IMS Buffer Pool Analyzer's information gathering phase. The Buffer Pool Analyzer trace can be stopped once this message appears with COLLECTION COMPLETE, and batch reports will still be able to produce valid buffer pool reports.

BPL201E **OPEN FAILED FOR** *ddname*

Explanation

An OPEN request for the indicated *ddname* failed.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Review the MVS SYSLOG for error messages associated with the open of the library. If no messages are found to indicate the cause of the problem, contact IBM Software Support for assistance.

BPL202E **BLDL FAILED FAILED RC=***nn* **REASON=***nn*

Explanation

A BLDL request failed. The return code and reason code are shown in the message.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Review the MVS SYSLOG for error messages associated with the library. If no messages are found to indicate the cause of the problem, contact IBM Software Support for assistance.

BPL203E **DFSMDA MEMBER FOR DBD** *dbdname* **DOES NOT INCLUDE DDNAME** *ddname*

Explanation

The IMS dynamic allocation module for DBD *dbdname* did not include data set name information for the *ddname* shown. This *ddname* is one of the required *ddname*s for this database, and the *ddname* should have been found in the dynamic allocation module.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Review the named DFSMDA member, and contact IBM Software Support for assistance.

BPL204E **CLOSE ERROR - DDNAME** *ddname*

Explanation

A data set CLOSE request failed.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Review the MVS SYSLOG for error messages associated with the close of the DD name. If no messages are found to indicate the cause of the problem, contact IBM Software Support for assistance.

BPL205E	**LOAD FAILED FOR DFSMDA MEMBER** *name* **ABEND=***nnn* **REASON=***nn*

Explanation

A LOAD request failed for the module name indicated in the message. The abend code and reason code are shown in the message.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Review the MVS SYSLOG for error messages associated with the error. If no messages are found to indicate the cause of the problem, contact IBM Software Support for assistance.

BPL206E	**INVALID DFSMDA MEMBER** *name* *rsn*

Explanation

An error occurred processing the named dynamic allocation module. The reason code (*rsn*) assists the IBM Software Support in determining the nature of the error.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL231E	**ERROR PROCESSING FASTPATH RECON RECORDS (***rsn***)**

Explanation

An error occurred processing RECON records associated with a Fast Path database. The reason code (*rsn*) provides the IBM Software Support with additional information on the cause of the error.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance. The problem probably occurred because a database defined in IMS as a Fast Path DEDB was defined in the RECON data set as a non-Fast Path database.

BPL232E	**ERROR PROCESSING HALDB RECON RECORDS (***rsn***)**

Explanation

An error occurred processing RECON records associated with a HALDB database. The problem probably occurred because a database defined in IMS as a HALDB database was defined in the RECON data set as a non-HALDB database. The reason code (*rsn*) provides the IBM Software Support with additional information on the cause of the error.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL241E	**UNABLE TO FIND 3 RECON DSNS FROM THE SAME SOURCE**

Explanation

The DBRC API requires that all three RECON data sets be available. Buffer Pool Analyzer was unable to verify that it had correctly found all three RECON data set names because some were present in the job's JCL and others were found in the dynamic allocation library.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Ensure that the batch job or IMS control region JCL have either all three RECON data set names included or that they have no data set names in the JCL.

BPL242E   **UNABLE TO FIND AT LEAST 2 RECON DATASETS**

Explanation

Two or more RECON data sets must be allocated to the BPLTRACE task. However, two or more were not found. It was not possible to proceed with accessing the RECON data sets.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Review other messages that might have occurred before this, such as those for dynamic allocation failures for RECON data sets. Contact IBM Software Support for assistance.

BPL243E   **RESERVE FAILED RC=*nn***

Explanation

A RESERVE macro failed with the indicated return code.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL244E   **RECON CLOSE ERROR RC=*nn* ACBERFLG=*nn***

Explanation

A CLOSE macro failed for a RECON data set with the indicated return code and ACB error flag values.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Review any RECON error messages that might have contributed to a close failure. The ACB error flag might also provide information related to the failure.

BPL245E   **UNABLE TO DETERMINE RECON VERSION (*rsn*)**

Explanation

A failure occurred processing the RECON header records. Buffer Pool Analyzer was unable to determine the version of the RECON data sets.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Review the output from a LIST.RECON STATUS command to determine the RECON release. Contact IBM Software Support for assistance.

BPL246E   **VSAM MODCB (*type*) RC=*nn* R0=*nn***

Explanation

A VSAM MODCB macro failed with the indicated return code and register 0 values. The (*type*) value is an indication of the specific MODCB that was in progress.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL247E   **ERROR IN SWAREQ RC=*nn***

Explanation

A SWAREQ macro failed. The return code is indicated in the message.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL248E	**ERROR DURING GETDSAB RC=***nn* ***R0=***nnnn*

Explanation

A GETDSAB macro failed. The return code and register 0 values are indicated in the message.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL249E	**ERROR DETERMINING ACTIVE RECON COPIES**

Explanation

Buffer Pool Analyzer experienced an error while reviewing the RECON header records.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Review the output from a LIST.RECON STATUS command to determine the RECON status. Contact IBM Software Support for assistance.

BPL250E	**UNABLE TO FIND BOTH COPY1 AND COPY2 RECONS**

Explanation

Buffer Pool Analyzer encountered an error while reviewing the RECON header records.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Review the output from a LIST.RECON STATUS command to determine the RECON status. Contact IBM Software Support for assistance.

BPL251E	**ERROR DETERMINING SPARE RECON**

Explanation

Buffer Pool Analyzer encountered an error while reviewing the RECON header records.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Review the output from a LIST.RECON STATUS command to determine the RECON status. Contact IBM Software Support for assistance.

BPL252E	**OPEN FAILED FOR** *ddname* **RC=***nn* **ACBERFLG=***nn*

Explanation

An OPEN macro failed for a RECON data set (allocated to the *ddname* shown) with the indicated return code and ACB error flag values.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Review any RECON error messages that might have contributed to an open failure. The ACB error flag might also provide information related to the failure.

BPL253E	**VSAM** *verb* **FAILED RC=***nn* **RPLERRCD=***nn*

Explanation

A VSAM GET or POINT macro failed for a RECON data set with the indicated return code and RPL error code values.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Review any RECON error messages that might have contributed to the error. The RPL error code might also provide information related to the failure.

BPL254E	**RECON DSN MISMATCH (JCL/ DFSMDA) FOR ddname**

Explanation

The RECON data set name present in the IMS DBRC task or batch job's JCL for the indicated RECON data set does not agree with the data set name in the dynamic allocation member in the job's STEPLIB or IMSDALIB.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Review the RECON data set names that are present in the batch job or IMS DBRC task JCL, along with the dynamic allocation members in job's STEPLIB or IMSDALIB DD names for inconsistencies.

BPL257E   **ATTACHX FAILED FOR DBRC API ROUTINE RC=***rc*

Explanation

The MVS ATTACHX macro returned an unexpected return code when it was called to attach load module BPLRCN2.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Review the MVS SYSLOG for other errors associated with the ATTACH failure, as well as the return code from the MVS ATTACH macro. For further assistance, contact IBM Software Support.

BPL258E   **UNEXPECTED IMS RELEASE ID FOUND**

Explanation

The IMS release in the trace directory contained an unexpected value.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Review the MVS SYSLOG for other BPL messages that might be related to this error. Also, ensure that

the target job or IMS control region is running on a supported release of IMS.

BPL259E   *modname* **FAILED ABEND=***abcode*

Explanation

Buffer Pool Analyzer's DBRC API processing routine failed with the indicated abend code.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Review the MVS SYSLOG for additional messages that describe the reason for the subtask failure.

BPL261E   **PARM FIELD CONTAINS AN UNKNOWN KEYWORD**

Explanation

The **PARM=** field in the BPLTRACE PROC contained an unknown keyword.

System action

Buffer Pool Analyzer ends the BPLTRACE task.

User response

Review the PROC and the values specified for the symbolic parameters to determine the cause of the error.

BPL262E   **IMSID** *imsid* **NOT FOUND ON THIS MVS SYSTEM**

Explanation

The IMSID specified in the BPLTRACE startup was not found on this MVS image.

System action

Buffer Pool Analyzer ends the BPLTRACE task.

User response

Ensure that BPLTRACE is running on the correct MVS LPAR, and verify that the proper IMSID was specified. If the IMS subsystem is expected to be started at a later time, use the delay time (**DT=**) or wait time (**WT=**) parameters to allow BPLTRACE to wait until the IMS subsystem is available.

BPL263E   **IMS RELEASE NOT SUPPORTED FOR** *subsystem-name*

Explanation

The job or IMSID specified in the BPLTRACE start command was found, but the release of IMS being run is not supported.

System action

Buffer Pool Analyzer ends the BPLTRACE task.

User response

Ensure that the subsystem that experienced the error is of one of the supported IMS releases.

BPL264E   **OPEN FAILED FOR REQUIRED DDNAME** *ddname*

Explanation

An MVS OPEN failed for the specified *ddname*.

System action

Buffer Pool Analyzer ends the BPLTRACE task.

User response

Ensure that the *ddname* is specified properly in the JCL. For example, the SYSLIB DD must refer to a PDS. Review the MVS SYSLOG to determine if there are other error messages associated with the open failure.

BPL265E   **GTF IS ALREADY ACTIVE - TRACE NOT INITIATED**

Explanation

A GTF trace is already active on this MVS image. Only one GTF trace can be active on an MVS LPAR.

System action

Buffer Pool Analyzer ends the BPLTRACE task.

User response

Ensure that there are no GTF traces executing when the BPLTRACE task is started. Also, Ensure that if multiple BPLTRACE tasks are active, only one task at a time attempts to activate the GTF.

BPL266E   **START COMMAND FAILED RC=*nnnn***

Explanation

An MCGRE macro issued to start the GTF trace address space failed with the indicated return code.

System action

Buffer Pool Analyzer ends the BPLTRACE task.

User response

Ensure that BPLTRACE is running APF authorized. Contact IBM Software Support for assistance.

BPL267E   **GTF TRACE INITIALIZATION FAILED-CHECK GTFPROC NAME FOR POSSIBLE ERRORS**

Explanation

BPLTRACE started the BPLGTF PROC, but the GTF trace was not successfully initialized.

System action

Buffer Pool Analyzer ends the BPLTRACE task.

User response

Review the MVS SYSLOG to determine why the BPLGTF started task failed.

BPL268E   **PARAMETER VALUE SPECIFIED FOR KEYWORD *keyword* IS TOO LONG**

Explanation

The value for the identified keyword exceeded the allowable length.

System action

Buffer Pool Analyzer ends the BPLTRACE task.

User response

Review the BPLTRACE value for the identified keyword where specified in the PROC or on the start command. Ensure the value conforms to the required lengths.

BPL269E   **PARAMETER VALUE SPECIFIED FOR *keyword* IS NOT A VALID NUMBER**

Explanation

The value for the identified keyword is not a numeric value.

System action

Buffer Pool Analyzer ends the BPLTRACE task.

User response

Review the BPLTRACE value for the identified keyword where specified in the PROC or on the start command. Ensure that the value specified for the identified keyword is numeric.

BPL270E   **GTFTIME PARAMETER IS INVALID OR MISSING**

Explanation

The value for the GTFTIME (or **GT=** PROC parameter) was either not specified or not valid.

System action

Buffer Pool Analyzer ends the BPLTRACE task.

User response

Review the BPLTRACE value for the GT parameter where specified in the PROC or on the start command. Ensure that the value specified is a valid value.

BPL271E   **PARAMETER KEYWORDS IMSID AND JOBANAME WERE BOTH SPECIFIED**

Explanation

Values were specified for both **JOB=** and **IMSID=** in the BPLTRACE started task.

System action

Buffer Pool Analyzer ends the BPLTRACE task.

User response

Only one of these parameters can be specified in a given BPLTRACE started task.

BPL272E   **PARAMETER DID NOT INCLUDE A VALUE FOR EITHER IMSID OR JOB**

Explanation

A value was not specified for either JOB= or IMSID= in the BPLTRACE started task.

System action

Buffer Pool Analyzer ends the BPLTRACE task.

User response

One of these parameters is required for BPLTRACE to identify the target of the trace process.

BPL273I   **WAITING FOR IMS BATCH INITIALIZATION TO COMPLETE - PHASE *x***

Explanation

A batch job that was requested to be traced has not completed IMS initialization. BPLTRACE will wait and recheck the status of the batch job every 2.5 seconds. The phase indicated in the message might be 1 - indicating that the batch SCD has not been built, or 2 - indicating that the batch initialization has not completed.

System action

None.

User response

None. This message is informational.

BPL274E   **ERROR FINDING ASCB FOR GTF STARTED TASK**

Explanation

Buffer Pool Analyzer failed to find the ASCB for the BPLGTF started task.

System action

Buffer Pool Analyzer ends the BPLTRACE task.

User response

Contact IBM Software Support for assistance.

BPL275E   **IMS BATCH JOB *jobname* NOT FOUND**

Explanation

The IMS DLI or DBB batch job named *jobname* was not found on this MVS system.

System action

Buffer Pool Analyzer ends the BPLTRACE task.

User response

Ensure that the job is running on the correct MVS LPAR (the same LPAR as BPLTRACE) and that the job is a DLI or DBB job. BMP jobs use the IMS control region

database buffer pools and, therefore, are not traced by their job name.

BPL276E   **FIND FAILED FOR BATCH SCD FOR JOB** *jobname*

Explanation

Buffer Pool Analyzer failed to find the batch SCD for the identified *jobname*.

System action

Buffer Pool Analyzer ends the BPLTRACE task.

User response

Contact IBM Software Support for assistance.

BPL277E   **ATTACHX FAILED FOR CONFIGURATION ROUTINE RC=***nn*

Explanation

An ATTACHX macro failed with the indicated return code.

System action

Buffer Pool Analyzer ends the BPLTRACE task.

User response

Check the MVS SYSLOG for other messages that might be related to the failure. The return code might also help you find the reason for the error. Contact IBM Software Support for assistance.

BPL278E   **ALESERV** *type* **FAILED RC=***nn*

Explanation

An MVS ALESERV macro failed. The type of the ALESERV (ADD or DELETE) is identified in the message text, as is the return code from the ALESERV request.

System action

Buffer Pool Analyzer ends the BPLTRACE task.

User response

Contact IBM Software Support for assistance.

BPL279E   *module* **FAILED ABEND=***abcode*

Explanation

The identified module name ended abnormally. The abend code is included in the message text.

System action

Buffer Pool Analyzer ends the BPLTRACE task and any BPLGTF task.

User response

If the abend is a user abend, other messages indicate the reason for the failure. For a system abend, review the *z/OS MVS System Codes* manual, and contact IBM Software Support.

BPL280E   **QEDIT {FREE | LIMIT} FAILED RC=***nn*

Explanation

An MVS QEDIT macro failed with the indicated return code. The request might have been a `free CIB` request or a `limit` request.

System action

Buffer Pool Analyzer ends the BPLTRACE task.

User response

Contact IBM Software Support for assistance.

BPL281E   **UNEXPECTED POST**

Explanation

The BPLTRACE task was posted, but Buffer Pool Analyzer was unable to determine the reason for the post.

System action

Buffer Pool Analyzer ends the BPLTRACE task, and any BPLGTF trace task.

User response

Keep the abend dump, and contact IBM Software Support for assistance.

BPL282E   **GTF TIME INTERVAL EXPIRED BEFORE CONFIGURATION DATA COLLECTION COMPLETED**
or
INCREASE GTF TRACE TIME AND RETRY

Explanation

The GTF trace timer interval expired before Buffer Pool Analyzer completed data gathering and tracing.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the BPLGTF task.

User response

Ensure that a reasonable GTF time (**GT=**) value was specified. While running, the F BPLTRACE,STATUS command can be used to obtain additional information about the status of the trace. Contact IBM Software Support if the configuration does not complete within a reasonable time.

BPL283E   **GTF NOT ACTIVE OR NOT RECORDING E440 USER RECORDS**

Explanation

A GTRACE macro indicated that GTF was not recording the appropriate user trace records required for Buffer Pool Analyzer.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL284E   **GTRACE ERROR - RETURN CODE** *nn*

Explanation

A GTRACE macro failed with the indicated return code.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Verify that the GTF trace that was started by the BPLTRACE task was running at the time of the error. Contact IBM Software Support for assistance.

BPL285E   **JOB** *xxxxxxxx* **ENDED BEFORE TRACE INIT COMPLETE - TRACE STOPPED**

Explanation

A batch job that was requested to be traced ended before the Buffer Pool Analyzer initialization process completed.

System action

Buffer Pool Analyzer stops the GTF trace and the Buffer Pool Analyzer trace address spaces.

User response

When selecting an IMS batch job to monitor, choose a batch job that will run for a longer time.

BPL286E   **SCD ADDRESS INVALID CODE** *reason*

Explanation

Validation of the target IMS subsystem's SCD address failed.

System action

Buffer Pool Analyzer ends the GTF trace and the Buffer Pool Analyzer trace address spaces.

User response

Ensure that the target job or IMSID did not end while Buffer Pool Analyzer was initializing. Contact IBM Software Support for further assistance.

BPL288I   *nnnnnn* **CONFIG ENTRIES TRACED**

Explanation

The BPLTRACE task traced the indicated number of IMS configuration records.

System action

None

Programmer response

None

BPL289I   *nnnnnn* **CONFIG ENTRIES WERE RETRIED**

Explanation

The BPLTRACE task tried to trace the indicated number of configuration records again because GTF discarded the initial trace attempt.

System action

The indicated number of records were traced once more.

Programmer response

If records had to be traced again, consider increasing the buffer space with the BPLGTF task.

BPL290I   **REQUESTED DELAY COMPLETED (*nnnn*) MINUTES**

Explanation

The delay time (**DT=**) value requested at BPLTRACE startup has completed.

System action

Buffer Pool Analyzer continues processing.

User response

None. This message is informational.

BPL291I   **IMS OSAMGTF TRACE OUTPUT ENABLED FOR *jobname***

Explanation

The IMS OSAM GTF trace was started for the identified batch job or IMS control region.

System action

Buffer Pool Analyzer continues processing.

User response

None. This message is informational.

BPL292I   **GTF SUCCESSFUL INITIALIZATION DETECTED**

Explanation

Buffer Pool Analyzer confirmed that GTF successfully initialized.

System action

Buffer Pool Analyzer continues processing.

User response

None. This message is informational.

BPL293I   **GTF TRACE STOP COMMAND ISSUED**

Explanation

Buffer Pool Analyzer issued the MVS STOP command to close the GTF address space.

System action

Buffer Pool Analyzer continues processing.

User response

None. This message is informational.

BPL294I   **IMS OSAMGTF TRACE OUTPUT DISABLED FOR *jobname***

Explanation

The IMS OSAM GTF trace was stopped for the identified batch job or IMS control region.

System action

Buffer Pool Analyzer continues processing.

User response

None. This message is informational.

BPL295I   **SUBSYSTEM *jobname* ENDED BEFORE TRACE INTERVAL EXPIRED**

Explanation

The batch job or IMS control region being traced ended before the GTF timer interval (**GT=**) completed.

System action

Buffer Pool Analyzer continues as though the GTF time (**GT=**) had completed.

User response

None. This message is informational.

BPL296I   **UNEXPECTED OPERATOR COMMUNICATION - *type***

Explanation

The BPLTRACE task received an unexpected post from the operator communication ECB. This could have been a mount command, STC communication, or some other communication request.

System action

The request is ignored.

User response

None. This message is informational.

BPL297I **STOP COMMAND ACKNOWLEDGED**

Explanation

The operator issued a MODIFY BPLTRACE,STOP or a STOP BPLTRACE command.

System action

Buffer Pool Analyzer initiates shutdown processing.

User response

None. This message is informational.

BPL298I **MODIFY PARAMETER MISSING OR INVALID**
or
VALID MODIFY PARAMETERS ARE STATUS, START, OR STOP

Explanation

An invalid MVS **MODIFY** command was issued. The message text shows the three valid parameters on a **MODIFY** command.

System action

Buffer Pool Analyzer ignores the command.

User response

Correct the command, and enter it again.

BPL299I **JOB *jobname* HAS ENDED OR CHANGED STEPS. TRACE STOPPED**

Explanation

The batch job or IMS control region being traced ended before the GTF timer interval (**GT=**) completed.

System action

Buffer Pool Analyzer continues as though the GTF time (**GT=**) had completed.

User response

None. This message is informational.

BPL300W **DYNAMIC UNALLOC FAILED RC=*xx* ERROR CODE=*xxxx* INFO CODE=*xxxx***

Explanation

Dynamic unallocation of the SYSLIB DD failed with the indicated return code and dynamic allocation error and information codes.

System action

The BPLTRACE task continues.

User response

If an enqueue conflict should occur for the PARMLIB data set between the BPLTACE and BPLGTF tasks, cancel both tasks. Contact IBM Software Support for assistance.

BPL301I **IMS BUFFER POOL ANALYSIS OPTIONS IN EFFECT**

Explanation

This indicates the beginning of messages that describe the option requested for this trace session. Messages BPL203I to BPL306I follow this message and describe the options selected. These messages are always issued during Buffer Pool Analyzer trace initialization.

When the BPLTRACE task is started, Buffer Pool Analyzer writes these messages to the MVS SYSLOG to identify the values in use for this process.

System action

Buffer Pool Analyzer continues processing.

User response

None. This message is informational.

BPL302I *job-type =name*

Explanation

The type of IMS task to be monitored, along with the name, are shown in the message. The message will show either a type of JOBNAME or IMSID and the name of the job or IMS control region.

System action

Buffer Pool Analyzer continues processing.

User response

None. This message is informational.

BPL303I **DELAY TIME (MINUTES) =*nnnn***

Explanation

The amount of delay time selected (or defaulted to) on the START command is shown in the message.

System action

Buffer Pool Analyzer continues processing.

User response

None. This message is informational.

BPL304I	MAXIMUM WAIT TIME (MINUTES) =*nnnn*

Explanation

The amount of wait time selected (or defaulted to) on the START command is shown in the message.

System action

Buffer Pool Analyzer continues processing.

User response

None. This message is informational.

BPL305I	GTF TRACE TIME (MINUTES) =*nnnn*

Explanation

The amount of GTF trace time selected (or defaulted to) on the START command is shown in the message.

System action

Buffer Pool Analyzer continues processing.

User response

None. This message is informational.

BPL306I	GTF PROC NAME =*taskname*

Explanation

The name of the GTF trace PROC is shown. The value, typically BPLGTF, is shown as specified in the START command or defaulted to in the BPLTRACE PROC.

System action

Buffer Pool Analyzer continues processing.

User response

None. This message is informational.

BPL320E	BLKSIZE OF SYSLIB DD IS LESS THEN 240

Explanation

The block size of the PARMLIB data set (SYSLIB DD) in the BPLTRACE PROCLIB is less than 240. The block size must be at least 240.

System action

The trace task abends.

Programmer response

Ensure that the PARMLIB data set has a RECFM of FB and that the block size is at least 240.

BPL321I	PARM FIELD PROCESSING IN PROGRESS

Explanation

In response to a STATUS request, this message indicates that **PARM=** field processing has not yet completed.

System action

None.

User response

None. This message is informational.

BPL322I	DELAY *status*

Explanation

In response to a STATUS or START MODIFY command, this message provides the status of the delay. Possible *status* field information is one of the following:

- COMPLETED (which shows the requested delay time)
- IN PROGRESS (which shows the amount of delay time remaining)
- BYPASSED BY OPERATOR (which shows the original delay time)

System action

None. The BYPASSED BY OPERATOR format of the message indicates that a MODIFY START command was entered and that the delay time was bypassed.

User response

None. This message is informational.

BPL323I   TRACE REQUESTED FOR *task*

Explanation

In response to a status command, this message shows either the IMSID (where *task* is IMS *imsid*) or the job name (where *task* is JOB *jobname*).

System action

None.

User response

None. This message is informational.

BPL324I   WAITING FOR JOB START -
**   *mmmm:ss* OF *nnnn* MINUTES**
**   REMAINING**

Explanation

In response to a STATUS request, this message indicates the amount of time that Buffer Pool Analyzer will continue looking for the requested IMS subsystem.

System action

None.

User response

None. This message is informational.

BPL325I   BUILDING GTF CONTROL CARDS

Explanation

In response to a STATUS request, this message indicates that GTF trace control cards are being built and written to the PARMLIB data set.

System action

None.

User response

None. This message is informational.

BPL326I   IMS OSAMGTF TRACE STATUS:
**   *status***

Explanation

In response to a STATUS request, this message indicates the status of the IMS OSAMGTF trace. It can be NOT STARTED, STARTED, or STOPPED.

System action

None.

User response

None. This message is informational.

BPL327I   GTF TRACE *status*

Explanation

In response to a STATUS command, this message shows the status of the GTF trace. Status can be any of the following:

- COMPLETE
- BEING STOPPED
- START SUCCESSFUL
- PENDING START VALIDATION
- IN PROGRESS (includes total time and remaining time of the trace)

System action

None.

User response

None. This message is informational.

BPL328I   IMS CONFIGURATION STATUS:
**   *status***

Explanation

In response to a STATUS request, this message indicates the status of the configuration process. It can be NOT STARTED, COMPLETE, IN PROGRESS, or STARTED.

System action

None.

User response

None. This message is informational.

BPL329I   STATUS DISPLAY COMPLETE

Explanation

In response to a STATUS request, this message appears at the end of all status messages to indicate that the status request is complete.

System action

None.

User response

None. This message is informational.

BPL330I   **NOT IN A DELAY WAIT - START REQUEST IGNORED**

Explanation

In response to a START request, this message indicates that Buffer Pool Analyzer is not waiting for the delay time (**DT=**) value.

System action

The command is ignored.

User response

If the trace is not executing, verify that the target job name or IMSID is running on the same MVS LPAR on which BPLTRACE is running.

BPL331I   **START REQUEST COMPLETE - REMAINING DELAY TIME BYPASSED (*mmmm:ss* OF *nnnn* MINUTES)**

Explanation

In response to a START request, this message indicates:

- The amount of time that was bypassed
- That Buffer Pool Analyzer will begin searching for the requested job or IMSID

System action

Delay time is bypassed as requested.

User response

None. This message is informational.

BPL351E   **ATTACHX FAILED FOR CATALOG SERVER RC=*xx***

Explanation

An MVS ATTACH for module BPLCATS failed with the indicated return code.

System action

The BPLTRACE task ends abnormally.

User response

Review the return code from the ATTACH macro identified in the message to determine the reason for the failure. Contact IBM Software Support for additional assistance.

BPL352E   **SUBTASK BPLCATS FAILED ABEND=*xxxxx***

Explanation

The BPLTRACE task determined that one of the catalog server tasks has abended with the indicated abend code.

System action

The BPLTRACE task ends abnormally.

User response

Review the MVS SYSLOG for additional messages that might indicate the reason for the failure. Contact IBM Software Support for additional assistance.

BPL353E   **SUBTASK BPLCATS TERMINATE REQUEST FAILED**

Explanation

The BPLTRACE task requested that the catalog server tasks end, but one or more tasks failed to end.

System action

The BPLTRACE task ends abnormally.

User response

Review the MVS SYSLOG for additional messages that might indicate the reason for the failure. Contact IBM Software Support for additional assistance.

BPL361W   **DFSVSAMP INPUT HAS INCONSISTENT CFOSAM *parameter***

Explanation

TA statement in the DFSVSAMP input data contained a CFNAMES statement that included an inconsistent value for one of the CFOSAM keyword parameters. The DFSVSAMP input should contain the same values for the structure name (STRNAME), directory ratio (DIRRATIO), and element ratio (ELERATIO). The value that was inconsistent is specified in the message text.

System action

The report job ends continues. The inconsistent value is ignored.

User response

Review the DFSVSAMP input statements to ensure that all subsystems contain consistent specifications of the CFOSAM parameters.

BPL362E   **CFOSAM KEYWORD MISSING VALUE FOR** *parameter*

Explanation

A statement in the DFSVSAMP input data contained a CFNAMES statement that was missing the value for one of the CFOSAM keyword parameters. The CFOSAM keyword must specify a value for the structure name (STRNAME) field.

System action

The report job ends continues. The CFNAMES statement is ignored.

User response

Review the DFSVSAMP input statements to ensure that all subsystems contain a value for the CFOSAM structure name parameter.

BPL382E   **IXCQUERY FAILED FOR OSAM STRNAME RC=***rc* **REASON=***reason*

Explanation

The IXCQUERY macro returned with an unexpected return code while querying the OSAM structure name. The return code and reason code are indicated in the message text.

System action

Data gathering continues, but OSAM coupling facility information will not be available for reporting.

User response

Contact IBM Software Support for assistance.

BPL383E   **ERROR PROCESSING IXCQUERY RESULTS -** *code*

Explanation

IMS Buffer Pool Analyzer encountered a problem interpreting the output of the IXCQUERY macro. The code indicates the reason for the error.

System action

Data gathering continues, but OSAM coupling facility information will not be available for reporting.

User response

Contact IBM Software Support for assistance.

BPL384E   **IXLCSP FAILED RC=***rc* **REASON=***reason*

Explanation

The IXLCSP macro returned with an unexpected return code. The return code and reason code are indicated in the message text.

System action

Data gathering continues, but OSAM coupling facility information will not be available for reporting.

User response

Contact IBM Software Support for assistance.

BPL401I   **NO OSAM CACHED SUBPOOLS FOUND**

Explanation

The OSAM cache analysis report process did not find any cached OSAM subpools either the actual traced environment or the DFSVSAMP modeled environment.

System action

Reporting continues with either the DFSVSAMP model environment or the actual environment.

User response

Ensure that DFSVSAMP or trace files contain the data desired for this execution of the OSAM cache analysis utility.

BPL421W   **MISSING CF CONFIGURATION INFO FOR** *type*

Explanation

Trace data provided to the OSAM cache analysis utility did not include configuration information for either

ELEMENTS or BUFFERS as specified in the message text.

System action

A default value is provided for the missing value and reporting continues.

User response

Review the messages in the MVS SYSLOG during the execution of the BPLTRACE process to see if any of the BPL381E, BPL382E, BPL383E, or BPL384E messages indicates that CF information might not have been included in the trace files.

BPL423W   INCONSISTENT IMSPLEX DEFINITIONS IN TRACES-*type*

Explanation

Buffer Pool Analyzer gathers information related to IMSPLEX definitions during the data gathering process that is written to the trace files. Inconsistent definitions were found while comparing the values among all the trace files that were input to the OSAM cache analysis report job. The type of information that was inconsistent is indicated in the message text, which can be one of the following:

Message text	Description
LOCK STR	Lock structure name
OSAM STR	OSAM structure name
VSAM STR	VSAM structure name
OSAM MSZ	OSAM structure maximum size
OSAM ASZ	OSAM structure actual size
VSAM MSZ	VSAM structure maximum size
VSAM ASZ	VSAM structure actual size
DIRRATIO	OSAM cache structure directory ratio
ELERATIO	OSAM cache structure element ratio
CACHEACT	OSAM caching active flag
CFNAME	Coupling facility name where the OSAM structure resides
# DIR	Number of directory entries contained in the OSAM structure
# ELEM	Number of element entries contained in the OSAM structure

Message text	Description
STR OVHD	The size of the OSAM structure overhead included in the structure size
DIR SIZE	The size of a directory entry contained in the OSAM structure

System action

The inconsistent value is ignored.

User response

Ensure that the inconsistency is expected. This message could indicate that IMS systems in different data sharing groups were included in the same OSAM cache analysis report request.

Unless you plan to merge data sharing groups, including different groups in the same cache analysis is not desirable because different data sharing groups use separate OSAM coupling facility cache structures.

BPL521E   OPEN FAILED FOR DDNAME *ddname*

Explanation

Open failed for the indicated *ddname*.

System action

The report job ends abnormally.

Programmer response

Verify that the indicated *ddname* is included in the JCL for the report job, and review the job log for any messages related to the OPEN failure.

BPL522E   GETDSAB FAILED RC=*nn* FOR DDNAME SYSIN

Explanation

An MVS GETDSAB macro failed for DDNAME SYSIN.

System action

The report job ends abnormally.

Programmer response

Verify that the SYSIN DDNAME is included in the JCL for the report job. Contact IBM Software Support for assistance.

BPL523E   TIME REQUEST FAILED

Explanation

An MVS TIME macro failed.

System action

The report job ends with a condition code 12.

Programmer response

Contact IBM Software Support for assistance.

BPL524E	INVALID VARY PERCENT - MUST BE 0 - 100

Explanation

A VARYPCT statement contains an invalid value. Permissible values for the VARYPCT statement are 0% to 100%.

System action

The report job ends with a condition code of 12.

User response

The VARYPCT statement preceding the message contains the error. Review the statement and specify a valid value for VARYPCT=.

BPL525E	VALUE OF *xxxxxxxx* MISSING OR INVALID - ABOVE STMT IGNORED

Explanation

A statement in the DFSVSAMP input data contained syntax that was not acceptable.

System action

The report job ends with a condition code of 12.

Programmer response

The statement preceding the message contains the error. Review the statement, and correct any errors. The value that caused the error is identified in the message text.

BPL526E	HIPERSPACE BUFFERS INVALID FOR BUFSIZE < 4096

Explanation

A statement in the DFSVSAMP input data contained a hiperspace buffer (or HSOptional/HSRequired) specification for a subpool with a buffer size of less than 4096 bytes.

System action

The report job ends with a condition code of 12.

Programmer response

The statement preceding the message contains the error. Hiperspace buffers are allowed only for subpools with buffer sizes of 4096 bytes or more.

BPL527E	SUBPOOL IS *xxxxxxxx* - ABOVE STMT IGNORED

Explanation

A statement in the DFSVSAMP input data contained a subpool specification (either **IOBF=** or **VSRBF=** statement) that was either invalid or a duplicate of a prior subpool definition.

System action

The report job ends with a condition code of 12.

Programmer response

The statement preceding the message contains the error. Review the statement, and correct the error.

BPL528E	FIRST POOLID FOLLOWS FIRST VSRBF - ABOVE STMT IGNORED

Explanation

A statement in the DFSVSAMP input data contained a VSAM pool definition (**POOLID=** statement) that follows the first VSAM subpool definition (**VSRBF=** statement).

System action

The report job ends with a condition code of 12.

Programmer response

If multiple VSAM pools are to be defined in an environment, all subpools must be associated with a VSAM pool. A **POOLID=** statement must precede each set of subpool definitions (**VSRBF=** statements).

BPL529E	POOLID ABOVE IS *xxxxxxxx* - ABOVE STMT IGNORED

Explanation

A statement in the DFSVSAMP input data contained a VSAM pool definition (**POOLID=** statement) that either duplicates a prior pool name or exceeds the maximum number of VSAM pool definitions.

System action

The report job ends with a condition code of 12.

Programmer response

For "DUPLICATE" error conditions, name every VSAM POOLID differently. For "TOO MANY" error conditions, reduce the number of VSAM pools to a maximum of 15 (minus any RESVPOOL pool identifiers specified).

BPL530E	**DATABASE** *xxxxxxxx* **HAS AN INVALID SUBPOOL SPECIFIED -** *xxxx*

Explanation

A statement in the DFSVSAMP input data specified a **DBD=** requirement for the named database to use a buffer pool with the specified name. The pool name requested in the **DBD=** statement was not specified on any **POOLID=** statement.

System action

The report job ends with a condition code of 12.

Programmer response

Correct the **DBD=** statement so that it specifies a valid pool name.

BPL531W	**SUBPOOL** *xxxx* **HAS NO DBD'S ASSIGNED VIA**® **DBD= STATEMENTS**

Explanation

A statement in the DFSVSAMP input data specified a VSAM pool or OSAM subpool but there were no databases defined to use the pool or subpool.

System action

Reports are produced, but the reporting job produces a return code of 4.

Programmer response

Consider removing the VSAM pool or OSAM subpool if there are no databases to be assigned to use the pool or subpool. If a database should be assigned to use the pool or subpool, ensure that a valid DBD= statement is present to assign the database to the appropriate buffer pool.

BPL532E	**INVALID KEYWORD - STATEMENT IGNORED**

Explanation

A statement in the SYSIN DD contained invalid syntax.

System action

The report job ends with a condition code of 12..

Programmer response

Verify that the SYSIN DD contains only valid statements. DSNRPT= and DBDRPT= are valid keywords that must be the first data in a statement.

BPL533E	**MISSING OR INVALID VALUE SPECIFIED FOR** *xxxxxxxx* **STATEMENT ABOVE**

Explanation

A statement in the SYSIN DD contained an invalid value.

System action

The report job ends with a condition code of 12.

Programmer response

Verify that the SYSIN DD contains only valid statements. DSNRPT= and DBDRPT= statements must specify a value of either YES or NO.

BPL534E	**MISSING OR INVALID VALUE SPECIFIED FOR KEYWORD** *xxxxxxxx* **IN CHANGE STATEMENT ABOVE**

Explanation

A CHANGE DBD statement either did not include the indicated keyword, or specified an invalid value for the keyword.

System action

The report job ends with a condition code of 12.

User response

Check the syntax of the CHANGE DBD statement preceding the message. Validate that the indicted keyword was specified and that the value for the keyword is a valid value.

BPL535E	**KEYWORD** *name* **IN CHANGE STATEMENT ABOVE** *error*

Explanation

A CHANGE DBD statement either did not include the indicated keyword, or specified an invalid value for the keyword.

System action

The report job ends with a condition code of 12.

User response

Check the syntax of the CHANGE DBD statement preceding the message. Validate that the indicted keyword was specified and that the value for the keyword is a valid value.

BPL536E	ABOVE BUFNO STATEMENT HAS INVALID BUFFER SIZE *size*

Explanation

A BUFNO statement included an invalid number of buffers within the specifications. The invalid value, or at least the first eight characters of the value, are included in the message text.

System action

The report job ends with a condition code of 12.

User response

Check the value indicated in the message. It is not a valid value for the type of BUFNO statement preceding this message.

BPL537E	NUMBER OF BUFFER SPECIFICATIONS EXCEEDS 45

Explanation

A BUFNO statement included more than 45 values. Buffer Pool Analyzer has a limit of 45 values that can be specified in a BUFNO statement.

System action

The report job ends with a condition code of 12.

User response

Reduce the number of specifications in the BUFNO statement to 45 or less.

BPL538E	DUPLICATE BUFFER SIZE SPECIFIED - *size*

Explanation

A BUFNO statement included the same value more than once. The number is indicated in the message.

System action

The report job ends with a condition code of 12.

User response

Eliminate the duplicate value(s) specified in the BUFNO statement that precedes this message.

BPL539E	ERROR PROCESSING ABOVE BUFNO STMT - *reason*

Explanation

An internal error occurred processing the preceding BUFNO statement.

System action

The report job ends with a condition code of 12.

User response

Ensure that no more than one BUFNO for the same pool type was specified in the SYSIN DD input. Contact IBM Software Support for assistance.

BPL540E	INVALID NUMBER OF BUFFERS IGNORED - *reason*

Explanation

A BUFNO statement included an entry in the list of values that was not acceptable. The "reason" in the message text indicates the cause of the error - either that there was an entry with a value of less than 3 or an entry with a value of more than 32768.

System action

The report job ends with a condition code of 12.

User response

Review the values of the entries in the flagged BUFNO statement, and ensure that there are no entries with a value of less than 3; and, for OSAM or VSAM, that there are no entries with a value more than 32768.

BPL541E	OPEN FAILED FOR DDNAME *ddname*

Explanation

Open failed for the indicated *ddname*.

System action

The report job ends abnormally.

Programmer response

Verify that the indicated *ddname* is included in the JCL for the report job, and review the job log for any messages related to the open failure.

BPL542E   **BUFFER POOL ANALYZER TRACE RECORDS MISSING - *xxxxxxxx***

Explanation

The Buffer Pool Analyzer reporting job scanned the first trace data set and was unable to find one or more Buffer Pool Analyzer trace records. The record type included in the message (TRACEDIR, ASID, SUBPOOL, DDT, or TGT NAME) indicates the type of 440 trace record that was not found.

System action

The Report job ends abnormally because the trace data is not usable.

Programmer response

Review the GTF trace configuration information defined in the BPLGTF PROC used for data gathering. Consider increasing the following values:

- **BLOK=**

  Note that the amount of storage specified is obtained from ECSA.

- **SIZE=***value*

  value can be specified only for the supported releases of z/OS.

- **NCP=***value*

After the PROC values have been reviewed and adjusted, rerun a new trace and report. If the problem continues, contact IBM Software Support for assistance.

BPL543E   **MULTIPLE JOBNAMES PRESENT IN GTF DATA – *jobname1* AND *jobname1***

Explanation

OSAM GTF records include the job name of the IMS subsystem that generated the record. There were multiple job names present in the GTF trace.

System action

The report job ends abnormally.

Programmer response

Make sure that the report data for only one job name is included in the report.

BPL544E   **BPL GTF RECORD HAS INVALID DATASET TYPE**

Explanation

A user 440 GTF record, generated by the BPLTRACE task, has an invalid data set type.

System action

The report job ends abnormally.

Programmer response

Contact IBM Software Support for assistance.

BPL545E   **SWAREQ FAILED RC=*nn***

Explanation

An MVS SWAREQ macro failed.

System action

The report job ends abnormally.

Programmer response

Contact IBM Software Support for assistance.

BPL546E   **GTF CONTROL RECORD VALIDATION FAILED - *reason***

Explanation

The Buffer Pool Analyzer verification on GTF control record information indicated that there was a problem.

System action

The report job ends abnormally.

Programmer response

The *reason* can be:

- An INVALID AID/FID, which indicates that the first record in the trace data set was not a control record.
- NOT A USER TRACE, which indicates that user type GTF trace records were not included in the trace.

- SYSIDS PRESENT, which indicates that GTF trace data sets were merged with IPCS trace facilities. GTF trace data sets cannot be merged before input to the reporting facility.

Contact IBM Software Support for assistance.

BPL547E	UNKNOWN SUBTYPE GTF 440 RECORD IGNORED

Explanation

A user 440 GTF record, generated by the BPLTRACE task, has an unknown record type indication.

System action

The report job ends abnormally.

Programmer response

Contact IBM Software Support for assistance.

BPL548E	INCONSISTENT TIMEZONES FOUND IN GTF RECORD

Explanation

GTF trace records contained more than one time zone offset.

System action

The report job ends abnormally.

Programmer response

Buffer Pool Analyzer uses timestamps to validate the consistency of GTF data. If the time zone offset was changed during a trace interval, the trace data cannot be used. Rerun the trace when the time zone on the LPAR will not be changed.

BPL549E	ERROR PROCESSING SPLIT RECORD - CODE *code*

Explanation

Buffer Pool Analyzer encountered an error merging records that GTF had split across blocks in the GTFC trace file.

System action

The report job ends abnormally.

Programmer response

Retain the dump, and contact IBM Software Support for assistance.

BPL550E	NO TRACE DDNAMES FOUND

Explanation

Buffer Pool Analyzer was unable to locate any TRACE*xxx* DD statements in the report job JCL.

System action

The report job ends abnormally.

Programmer response

Review the report job JCL and ensure that at least one DD name beginning with TRACE is present.

BPL561E	UNEXPECTED GTF RECORD TYPE PASSED TO OSAM PROCESSOR (*type*)

Explanation

Buffer Pool Analyzer passed an incorrect record to the OSAM subroutine.

System action

The report job ends abnormally.

Programmer response

Contact IBM Software Support for assistance.

BPL562E	NUMBER OF OSAM DDNAMES EXCEEDED 65,536

Explanation

More than 65,536 ddnames were found in the GTF trace data sets. IMS Buffer Pool Analyzer has a limit of 65,536.

System action

The report job ends abnormally.

Programmer response

If more than one trace data set was included in execution of the report, make sure that the traces are from a single IMS subsystem. Contact IBM Software Support for assistance.

BPL563E	INCONSISTENT *type* VALUE FOR DDN *ddname*

Explanation

There was inconsistent information in the trace data set(s) for the *type* (BLKSIZE, BFRSIZE, or SUBPOOL) for the indicated *ddname*.

System action

The report job ends abnormally.

Programmer response

If more than one trace data set was included in execution of the report, ensure that the trace information is consistent in all traces. Contact IBM Software Support for assistance.

BPL565E	OSAM GTF *type* DOES NOT MATCH BPL GTF DATA OSAM=*value1* BPL=*value2*

Explanation

There was inconsistent information in the trace data set(s) for the *type* (BLKSIZE, SUBPOOL, or DMB/DCB number) between the OSAM trace records and the Buffer Pool Analyzer trace records.

System action

The report job ends abnormally.

Programmer response

If more than one trace data set was included in execution of the report, make sure that the trace information is consistent in all traces. Contact IBM Software Support for assistance.

BPL568E	NO VALID BUFFER SUBPOOLS WERE FOUND FOR *xxxxxxxx* IN POOL *xxxx*

Explanation

The named database could not be assigned a subpool in the named OSAM buffer pool. This occurs only for OSAM pools specified in the DFSVSAMP input data set.

System action

Reporting continues. The database identified in the message is not reported in the modeled results.

Programmer response

Review the OSAM definitions specified for the named subpool. Make sure there is a subpool with buffers large enough to hold blocks from this DBD.

BPL569E	DDNAME *xxxxxxxx* DMB/DCB number FOUND IN OSAM GTF DATA HAD NO CORRESPONDING BPL DATA

Explanation

OSAM GTF trace records were found in the trace data set for the named DDNAME, but the Buffer Pool Analyzer GTF trace records did not identify that DDNAME as a valid IMS database DDNAME.

System action

The trace records for this DDNAME are reported, but the DBDNAME and DSNAME associated with this DDNAME will be blank.

Programmer response

Contact IBM Software Support for assistance.

BPL570E	OSAM GTF DATA CONTAINS SUBPOOL *xxxx* SIZE *nnnnn* NOT PRESENT IN BPL DATA

Explanation

OSAM GTF trace records were found in the trace data set that indicated an OSAM subpool with the indicated subpool name and size. Buffer Pool Analyzer GTF trace records did not identify this subpool.

System action

The trace records for this buffer pool are reported, but some information might not be valid.

Programmer response

Contact IBM Software Support for assistance.

BPL581E	INVALID PRINT FUNCTION DD IDENTIFIER

Explanation

An error occurred writing a message to the print data set. The print function identifier was invalid.

System action

The report job ends abnormally.

User response

Contact IBM Software Support for assistance.

BPL582E   PRINT OPEN REQUESTED WHEN ALREADY OPEN

Explanation

An error occurred opening the print data set. The print file was already open.

System action

The report job ends abnormally.

User response

Contact IBM Software Support for assistance.

BPL583E   PRINT CLOSE REQUESTED WHEN NOT OPEN

Explanation

An error occurred closing the print data set. The print file was not open.

System action

The report job ends abnormally.

User response

Contact IBM Software Support for assistance.

BPL584E   MESSAGE EXCEEDS 132 BYTES

Explanation

An error occurred writing a message to the print data set. The message length exceeded the maximum allowable value.

System action

The report job ends abnormally.

Programmer response

Contact IBM Software Support for assistance.

BPL585E   PRINT REQUESTED BEFORE OPEN

Explanation

Buffer Pool Analyzer requested that a message be printed before the print file was opened.

System action

The report job ends abnormally.

User response

Contact IBM Software Support for assistance.

BPL586E   INVALID PRINT FUNCTION REQUEST BYTE

Explanation

An error occurred writing a message to the print data set. The message request type was invalid.

System action

The report job ends abnormally.

Programmer response

Contact IBM Software Support for assistance.

BPL587E   OPEN FAILED FOR PRINT DDNAME *ddname*

Explanation

Open failed for the indicated *ddname*.

System action

None.

Programmer response

Verify that the indicated *ddname* is included in the JCL for the report job, and review the job log for any messages related to the open failure.

BPL588E   OUTPUT FOR DDNAME *ddname* DISCARDED-DD COULD NOT BE OPENED

Explanation

Output could not be written to the named DD because the DD failed to open.

System action

None.

User response

Verify that the indicated ddname is included in the JCL for the report job.

BPL589E   CLOSE FAILED FOR DDNAME *ddname*

Explanation

An attempt to close the named DD failed with a return code 4.

System action

The report job ends abnormally.

User response

Review the batch job's job log for messages that might be related to the close failure.

BPL590E	TITLE TEXT EXCEEDS POSSIBLE LENGTH

Explanation

A request to specify title text exceeded the maximum length for title text.

System action

The report job ends abnormally.

User response

Contact IBM Software Support for assistance.

BPL591E	ITKB REQUESTED BUT SOFTWARE NOT AVAILABLE

Explanation

The IMS Tools Knowledge Base load modules were not present in the report job STEPLIB data sets.

System action

The report job completes successfully, but output is only written to SYSOUT, not to the IMS Tools Knowledge Base repository.

User response

Ensure that the IMS Tools Knowledge Base SHKTLOAD library is included in the report job's STEPLIB, or that the ITKBSRVR= specification in the PARM field is omitted or specified as blanks if you do not wish to use the IMS Tools Knowledge Base.

BPL592E	BLDL FAILED FOR ITKB MODULE HKTXXLI RC=*rc* R0=*reason*

Explanation

An MVS BLDL for IMS Tools Knowledge Base module HKTXXLI returned an unexpected return code and reason code. The return code and reason code are shown in the message text.

System action

The report job ends abnormally.

User response

Determine the reason for the BLDL failure, and correct the problem. Contact IBM Software Support for assistance.

BPL593E	ITKB function FAILED DD=*ddname* RC=*rc* R0=*reason*

Explanation

An IMS Tools Knowledge Base request failed. The function code will be INITIAL, OPEN, WRITE, CLOSE, or TERMINATE. The DD name associated with the request, and the function return code and reason code are listed in the message.

System action

The report job ends abnormally.

User response

Contact IBM Software Support for assistance.

BPL594I	ITKB SERVER NOT ACTIVE-ITKB REPORTING BYPASSED

Explanation

The ITKBSRVR= specification in the PARM field of the report job was not found on the LPAR where the report job is executing.

System action

The report job completes successfully, but output is only written to SYSOUT, not to the IMS Tools Knowledge Base repository.

User response

Correct the IMS Tools Knowledge Base server name specified in the PARM field of the report job so that it specifies a server name that is active on the MVS LPAR where the report job executes.

BPL595I	BUFFER POOL ANALYZER PRODUCT NOT DEFINED TO ITKB - ITKB BYPASSED

Explanation

The IMS Tools Knowledge Base repository was not properly initialized for the Buffer Pool Analyzer product.

System action

The report job completes successfully, but output is only written to SYSOUT, not to the IMS Tools Knowledge Base repository.

User response

Ensure that the installation procedures documented in the Customization chapter of this manual were followed properly. The BPLITKB job in the SBPLSAMP library performs the required IMS Tools Knowledge Base initialization.

BPL596I	ITKB REPORT ID id NOT DEFINED - ITKB REPORT BYPASSED

Explanation

The IMS Tools Knowledge Base repository definitions for the report ID specified were missing or set to discard output.

System action

The report job completes successfully, but output is only written to SYSOUT, not to the IMS Tools Knowledge Base repository.

User response

Ensure that the installation procedures documented in the Customization chapter of this manual were followed properly. The BPLITKB job in the SBPLSAMP library performs the required IMS Tools Knowledge Base initialization.

BPL612I	TRACE DATASET WRAPPED TRACE START TIME REFLECTS THE OLDEST GTF DATA IN THE DATASET

Explanation

The Buffer Pool Analyzer determined that the trace data set wrapped.

System action

None.

Programmer response

Try to make sure that the trace data set is allocated with enough space to hold the data generated.

BPL613I	SOME TRACE RECORDS ARE MISSING. PER SECOND INFORMATION WILL BE IMPACTED DEPENDING ON THE NUMBER OF MISSING TRACE RECORDS

Explanation

Buffer Pool Analyzer determined that GTF could not write records to the trace data set fast enough to match the number of trace records generated.

System action

None.

Programmer response

Review "Customizing the sample started tasks" on page 10, especially the performance options for the BPLGTF PROC. Specifically, consider increasing the **NCP=**, **BLOK=** and, if applicable, the **SIZE=** parameters. Note that the **BLOK=** parameter specifies the amount of ECSA that GTF will use; therefore, ensure that you consider the amount of available ECSA before changing this parameter.

BPL621E	UNEXPECTED GTF RECORD TYPE PASSED TO VSAM PROCESSOR (*type*)

Explanation

Buffer Pool Analyzer passed an incorrect record to the VSAM subroutine.

System action

The report job ends abnormally.

Programmer response

Contact IBM Software Support for assistance.

BPL622E	NUMBER OF VSAM DDNAMES EXCEEDED 65,536

Explanation

More than 65,536 VSAM component names were found in the GTF trace data sets. IMS Buffer Pool Analyzer has a limit of 65,536.

System action

The report job ends abnormally.

Programmer response

If more than one trace data set was included in the report job, make sure that the traces are from a single IMS subsystem. Contact IBM Software Support for assistance.

BPL623E **INCONSISTENT** *type* **VALUE FOR DSN** *dsname*

Explanation

There was inconsistent *type* information in the trace data set(s) (BLKSIZE or CISIZE, BFRSIZE, or SUBPOOL) for the indicated DSNAME.

System action

The report job ends abnormally.

Programmer response

If more than one trace data set was included in the report job, make sure that the trace information is consistent in all traces. Contact IBM Software Support for assistance.

BPL625E **VSAM GTF** *type* **DOES NOT MATCH BPL GTF DATA VSAM=***value1* **BPL=***value2*

Explanation

There was inconsistent *type* information (VSAMTYPE, VSHRPOOL, VSAMCLST, or DMB/DCB number) in the trace data set(s) between the OSAM trace records and the Buffer Pool Analyzer trace records.

System action

The report job ends abnormally.

Programmer response

If more than one trace data set was included in the report job, make sure that the trace information is consistent in all traces. Contact IBM Software Support for assistance.

BPL626E **UNKNOWN DSN FOUND IN VSAM GTF RECORD –** *dsn*

Explanation

A VSAM GTF trace record contained a data set name that was not included in a Buffer Pool Analyzer GTF trace record.

System action

Reporting continues, although this data set will be reported with a blank DBDNAME and DDNAME.

Programmer response

If more than one trace data set was included in the report job, make sure that the trace information is consistent in all traces. Contact IBM Software Support for assistance.

BPL627E **UNKNOWN SUBPOOL FOUND IN VSAM GTF RECORD –** *bufsize*

Explanation

A VSAM GTF trace record contained a buffer size that was not included in a Buffer Pool Analyzer GTF trace record.

System action

Reporting continues. A buffer pool is created with the indicated buffer size.

Programmer response

If more than one trace data set was included in execution of the report, make sure that the trace information is consistent in all traces. Contact IBM Software Support for assistance.

BPL628E **NO VALID BUFFER SUBPOOLS WERE FOUND FOR** *xxxxxxxx* **IN POOL** *xxxx*

Explanation

The named database could not be assigned a subpool in the named VSAM pool. This occurs only for VSAM pools specified in the DFSVSAMP input data set.

System action

Reporting continues. The database identified in the message is not reported in the modeled results.

Programmer response

Review the VSAM definitions specified for the named subpool. Make sure that there is a subpool with buffers large enough to hold blocks from this DBD. Also if

one of the database data sets is a KSDS, make sure that both data and index subpools exist that are large enough to hold control intervals from data sets associated with this DBD.

BPL641E   **ERROR PROCESSING FASTPATH RECON RECORDS (*code*)**

Explanation

An error occurred processing DEDB RECON records.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL642E   **ERROR PROCESSING HALDB RECON RECORDS (*code*)**

Explanation

An error occurred processing HALDB RECON records.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL643E   **DBRC API FUNCTION *function* RC=*rc* REASON=*reason***

Explanation

A call to the IMS DBRC API failed. The function, return code, and reason code are shown in the message text.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Review the DBRC API return code and reason code in the *IMS System Programming APIs*. Contact IBM Software Support for assistance.

BPL644W   **ERROR PROCESSING HALDB PARTITION ID FOR DBD *dbdname***

Explanation

An unknown partition ID was encountered in one of the data sets for the database name indicated in the message text.

System action

Buffer Pool Analyzer continues processing.

User response

Review the DBRC for the named database to ensure that all database data sets have a valid partition ID. Reports produced using this trace file might experience unexpected results.

BPL651I   **ABEND S*aaa* RC*bb* DDNAME SYSLIB**

Explanation

An abend occurred processing the SYSLIB DD in the BPLTRACE process. The abend code (*aaa*) and return code (*bb*) document the reason for the abend.

System action

If the abend is a system abend 013 with return code 18, resulting from the BPLTRACE member not already being in the PARMLIB data set, IMS Buffer Pool Analyzer writes a new member. Otherwise, the BPLTRACE task abends.

Programmer response

Determine the reason for the abend and address the cause of the problem. If a space abend such as B37/D37/E37 abend occurs, then increase the size of the PARMLIB data set.

BPL652E   **COMPRESS FAILED - *reason***

Explanation

IMS Buffer Pool Analyzer encountered an error attempting to compress the SYSLIB DD. The reason for the error is shown in the message.

System action

The BPLTRACE task abends.

Programmer response

The error might be caused by one of the following reasons:

- ATTACH RETURN CODE *nn*

Determine the reason for the ATTACH macro to have failed, and contact IBM Software Support for assistance.

- IEBCOPY RETURN CODE *nn*

Determine the reason for IEBCOPY to have failed, and contact IBM Software Support for assistance.

BPL653E   **UNABLE TO RECOVER - *reason***

Explanation

IMS Buffer Pool Analyzer encountered an error attempting to recover from an abend during SYSLIB DD processing. The reason for the error is shown in the message.

System action

The BPLTRACE task abends.

User response

The error might be caused by one of the following reasons:

- IGNORE OPTION NOT AVAILABLE

  The DCB Abend Exit cannot ignore the error. This might be caused by the error being encountered during close processing.

- MULTIPLE ABENDS OCCURRED

  More than one abend occurred processing the SYSLIB DD. IMS Buffer Pool Analyzer only attempts recovery from a single abend.

- UNEXPECTED ABEND

  An abend other than a system 013 return code 18 occurred. IMS Buffer Pool Analyzer does not recover from other abends.

Review the abend code and any other messages associated with the error that might appear in the MVS SYSLOG. Contact IBM Software Support for assistance.

BPL661E   **GTF NOT ACTIVE OR NOT RECORDING E440 USER RECORDS**

Explanation

A GTRACE macro indicated that GTF was not recording the appropriate user trace records required for Buffer Pool Analyzer.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL662E   **GTRACE ERROR - RETURN CODE *nn***

Explanation

A GTRACE macro failed with the indicated return code.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Verify that the GTF trace started by the BPLTRACE task was running at the time of the error. Contact IBM Software Support for assistance.

BPL663E   **UNEXPECTED CONDITION PROCESSING DEDB ACTIVITY *a***

Explanation

An error occurred processing a DEDB access. An unexpected condition occurred interpreting the DEDB buffer request and the activity required to meet the buffer request.

System action

Buffer Pool Analyzer writes this message and message BPL664E and ignores the DEDB buffer request.

User response

Retain the messages issued during the problem occurrence, and contact IBM Software Support for assistance.

BPL664E   ***name* TRACE FIELD ENDS WITH *value***

Explanation

An error occurred processing a DEDB access. An unexpected condition occurred interpreting the DEDB buffer request and the activity required to meet the buffer request. This message shows the EPSTTRAC or DMHRTRAC fields and the last 16 bytes of trace IDs present in those control blocks.

System action

Buffer Pool Analyzer writes this message and message BPL663E and ignores the DEDB buffer request.

User response

Retain the messages issued during the problem occurrence, and contact IBM Software Support for assistance.

BPL681E   **UNEXPECTED GTF RECORD TYPE PASSED TO DEDB PROCESSOR** *type*

Explanation

Buffer Pool Analyzer passed an incorrect record to the DEDB subroutine.

System action

The job report ends abnormally.

User response

Contact IBM Software Support for assistance.

BPL682E   **NUMBER OF DEDB DD NAMES EXCEEDED 16,384**

Explanation

More than 16,384 DEDB component names were found in the GTF trace data sets. Buffer Pool Analyzer has a limit of 16,384.

System action

The job report ends abnormally.

User response

If more than one trace data set was included in the report job, ensure that the traces are from a single IMS subsystem. Contact IBM Software Support for assistance.

BPL683E   **INCONSISTENT** *type* **VALUE FOR DSN** *dsname*

Explanation

There was inconsistent *type* information (BLKSIZE or CISIZE, BFRSIZE, or SUBPOOL) in the trace data set(s) for the indicated DSNAME.

System action

The job report ends abnormally.

User response

If more than one trace data set was included in the report job, ensure that the trace information is

consistent in all traces. Contact IBM Software Support for assistance.

BPL685E   **VSAM GTF** *type* **DOES NOT MATCH BPL GTF DATA VSAM=***value1* **BPL=***value2*

Explanation

There was inconsistent *type* information (VSAMTYPE, VSHRPOOL, SVSAMCLST, or DMB/DCB number) in the trace data set(s) between the OSAM trace records and the Buffer Pool Analyzer trace records.

System action

The job report ends abnormally.

User response

If more than one trace data set was included in the report job, ensure that the trace information is consistent in all traces. Contact IBM Software Support for assistance.

BPL686E   **UNKNOWN DSN FOUND IN VSAM GTF RECORD -** *dsn*

Explanation

A VSAM GTF trace record contained a data set name that was not included in a Buffer Pool Analyzer GTF trace record.

System action

Reporting continues, although this data set will be reported with a blank DBDNAME and DDNAME.

User response

If more than one trace data set was included in the report job, ensure that the trace information is consistent in all traces. Contact IBM Software Support for assistance.

BPL687E   **UNKNOWN SUBPOOL FOUND IN** *db_type* **GTF RECORD -** *text*

Explanation

In the message text, *db_type* can be either VSAM or DEDB. Depending on the *db_type* value, this message can have different meanings, as follows:

If *db_type* **is VSAM**
 A VSAM GTF trace record contained a buffer size (*text*) that was not included in a Buffer Pool Analyzer GTF trace record.

If *db_type* **is DEDB**

A pool named *text* (8 characters) was included in a DEDB GTF trace record for a buffer pool request. However, this pool was not found in Fast Path DEDB buffer pools defined by the DEDB= control statement in the DFSVSM*xx* member of the IMS PROCLIB data set. This error can occur in one or more of the following cases:

- **Case A:** If this message is reported in SPLRPT, it means that the buffer pool named *text* was not included in the Buffer Pool Analyzer trace records (GTF 440 records).

- **Case B:** If this message is reported in MDLRPT, it means that the buffer pool named *text* is not defined in the DFSVSAMP DD statement of BPLMAIN JCL.

- **Case C:** In addition, if the pool name (*text*) consists of the CI size and blanks (such as "bbbb512b" or "bbb2048b", where b represents a blank), this pool is a buffer pool created by IMS due to an error in the definition of Fast Path DEDB buffer pools.

System action

See the system action for the applicable database type:

If *db_type* **is VSAM**

Reporting continues. A buffer pool is created with the indicated buffer size.

If *db_type* **is DEDB**

IMS Buffer Pool Analyzer creates a buffer pool named *text* with 1000 buffers and continues processing. The size of this buffer pool is the buffer size specified in the DEDB GTF trace record.

User response

See the user response for the applicable database type:

If *db_type* **is VSAM**

If more than one trace data set was included in the report job, ensure that the trace information is consistent across traces. Contact IBM Software Support for assistance.

If *db_type* **is DEDB**

Take one or more of the following actions for the cases that apply to you:

- **Case A:** If more than one trace data set was included in the report job, ensure that the trace information is consistent across traces.

- **Case B:** Ensure that DEDB buffer pools are defined correctly in the DFSVSAMP DD statement

in the BPLMAIN JCL. Also check if Case A is applicable.

- **Case C:** Correct the definition of Fast Path DEDB buffer pools in DFSVSM*xx*, and reacquire the trace by running BPLTRACE before running BPLMAIN.

Note that you cannot specify a buffer pool named *text* in DFSVSAMP in the BPLMAIN JCL because it is a buffer pool that can be created only if there is an error in the definition of Fast Path DEDB buffer pools.

If no error is found in your IMS system, contact IBM Software Support.

BPL688E   **DBD/AREA ENTRY WAS NOT FOUND FOR AREA** *name*

Explanation

A DEDB area present in a DEDB GTF record was not found in the list of DEDB areas found in the IMS configuration.

System action

The job report ends abnormally.

User response

Contact IBM Software Support for assistance.

BPL701E   **OPEN FAILED FOR DDNAME** *ddname*

Explanation

The BPLTRACE task attempted to open the stated DD name but the OPEN failed.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Review the MVS SYSLOG for error messages associated with the OPEN failure. Contact IBM Software Support for assistance.

BPL702E   **BLDL FAILED FOR MODULE** *module*

Explanation

A BLDL for the named module failed.

System action

IMS Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

The named module was not found in either the IMS control region STEPLIB/JOBLIB or LINKLIST. Contact IBM Software Support for assistance.

BPL703E   **LOAD FAILED FOR MODULE** *name* **RC=***rc* **ABCODE=***abcode*

Explanation

An MVS LOAD for the named module failed. The return code and abend code are listed in the message text.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Review the MVS SYSLOG for additional messages associated with the LOAD failure. Contact IBM Software Support for assistance.

BPL704E   **FIND FAILED FOR MEMBER** *name* **RC=***rc*

Explanation

An MVS FIND macro for the named module failed with the indicated return code.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Review the MVS SYSLOG for additional messages associated with the FIND failure. Contact IBM Software Support for assistance.

BPL705E   **UNEXPECTED EOF IN RESLIB MEMBER DBFDEDB0 PROCESSING CESD RECORDS**

Explanation

While reading member DBFDEDB0 from RESLIB, an MVS READ macro indicated that end-of-file was encountered when end-of-file was not expected.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL706E   **FAILED TO LOCATE CSECT** *name* **IN RESLIB MODULE DBFDEDB0**

Explanation

Buffer Pool Analyzer failed to locate the named CSECT name within RESLIB load module DBFDEDB0.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL707E   **POINT FAILED IN MEMBER** *name* **RC=***rc*

Explanation

An MVS POINT macro issued within the named member returned an unexpected return code.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL708E   **UNEXPECTED EOF IN RESLIB MEMBER DBFDEDB0 PROCESSING RELOCATION RECORDS**

Explanation

Buffer Pool Analyzer encountered an unexpected end-of-file condition while reading RESLIB member DBFDEDB0.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL709E   **ERROR PROCESSING RESLIB MEMBER DBFDEDB0**

Explanation

Buffer Pool Analyzer encountered an unexpected condition while processing RESLIB member DBFDEDB0.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL710E   **ERROR PLANTING DEDB INTERCEPT CODE** *code*

Explanation

Buffer Pool Analyzer encountered an error while planting the BPA intercept to create DEDB GTF records. An internal code is indicated in the message text.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL712E   **DEDB INTERCEPT FAILURE FOR MODULE DBFSGAB*n* CODE** *code*

Explanation

Buffer Pool Analyzer encountered an error while planting the BPA intercept in module DBFSGAB0 to create DEDB GTF records. An internal code is indicated in the message text.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL713I   **BUFFER POOL ANALYZER DEDB I/O INTERCEPTS** *status*

Explanation

Buffer Poor Analyzer successfully planted the intercepts required to generate DEDB activity GTF records (for a status of ACTIVATED) or has removed the intercepts required to track DEDB activity (for a status of DEACTIVATED).

System action

None.

User response

None. This message is informational.

BPL714E   **UNABLE TO LOCATE MODULE DBFDEDB0 IN IMS CTL**

Explanation

Buffer Pool Analyzer encountered an error trying to locate module DBFDEDB0 in the IMS control region.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL715E   **CLOSE FAILED FOR DDNAME** *ddname*

Explanation

An MVS CLOSE failed for the indicated DD name.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Review the MVS SYSLOG for any other errors that might be associated with the CLOSE failure. Contact IBM Software Support for assistance.

BPL721E   **INTERCEPT DEACTIVATE FAILED FOR DBFSGAB0 CODE** *code*

Explanation

Buffer Pool Analyzer failed to deactivate the intercept for IMS module DBFSGAB0.

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL722E	INTERCEPT DEACTIVATE FAILED FOR DBFSGAB2 CODE *code*

System action

Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact IBM Software Support for assistance.

BPL723W	BUFFER POOL ANALYZER TRACE RECORDS MISSING - TYPE=*xxxxxxxx* DDNAME=*ddname*

Explanation

While scanning the first configuration records of the trace data set, the IMS Buffer Pool Analyzer reporting job was unable to find one or more Buffer Pool Analyzer trace records of record type *xxxxxxxx*.

In the message,

ddname
    Indicates the DD name of the trace data set to be scanned.

xxxxxxxx
    Indicates the type of 440 trace record that was not found. The value can be either SUBPOOL or DDT.

System action

IMS Buffer Pool Analyzer continues processing and reads the second configuration trace records to gather information for reporting.

System programmer response

Check the job log to see if BPL542E message follows this message. If so, follow the instructions for BPL542E message. Otherwise, ignore this message.

BPL724I	IMS MANAGEMENT OF ACBS IS ACTIVE

BSDSNAME=*bsds_name*

Explanation

IMS catalog and IMS directory are active. *bsds_name* shows the name of the IMS bootstrap data set. IMS Buffer Pool Analyzer accesses the IMS directory using the high-level qualifier (*dir_hlq*) specified in the IMS control region JCL.

System action

IMS Buffer Pool Analyzer continues processing.

User response

None. This message is informational.

BPL725E	DFS3CATQ MODULE NOT FOUND

Explanation

The DFS3CATQ module, which is a required module to obtain information about IMS-managed ACBs, is not found in the STEPLIB, JOBLIB, or LINKLIST concatenation of the IMS control region JCL.

System action

The BPLTRACE task ends abnormally.

User response

Contact the IMS administrator at your site for assistance. If no error is found in your IMS system, contact IBM Software Support.

BPL726E	DFS3CATQ API ERROR, FUNCTION=*function* RC=*return_code* RSN=*reason_code*

Explanation

IMS Buffer Pool Analyzer trace task (BPLTRACE) detected an IMS catalog API (DFS3CATQ) error. The error occurred while obtaining information about IMS-managed ACBs from the IMS directory. *function*, *return_code*, and *reason_code* show information returned from the IMS catalog API.

System action

The BPLTRACE task ends abnormally.

User response

Look up the return code and reason code from the IMS catalog API (DFS3CATQ) in *IMS System Programming*

APIs. Correct the problem and restart the BPLTRACE task.

BPL727E	DIRECTORY INFORMATION NOT FOUND IN IMS MANAGED ACBS FOR DBD=*dbdname*

Explanation

IMS Buffer Pool Analyzer issued an IMS catalog API list request (DFS3CATQ FUNCTION=LIST) to obtain directory information. The directory information for the indicated DBD-type ACB was not found in the IMS-managed ACBs environment.

System action

IMS Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact the IMS administrator at your site for assistance. If no error is found in your IMS system, contact IBM Software Support.

BPL728E	INVALID DIRECTORY INFORMATION IN IMS MANAGED ACBS FOR DBD=*dbdname*

Explanation

The indicated DBD-type ACB, which is managed by IMS, does not contain valid directory information.

System action

IMS Buffer Pool Analyzer ignores the database with the error and continues processing.

User response

Contact the IMS administrator at your site for assistance. If no error is found in your IMS system, contact IBM Software Support.

BPL729E	ACB NOT FOUND IN IMS MANAGED ACBS FOR DBD=*dbdname*

Explanation

IMS Buffer Pool Analyzer issued an IMS catalog API get request (DFS3CATQ FUNCTION=GET) to obtain object definition from the IMS catalog. Although the directory information for the indicated DBD-type ACB exists in the IMS-managed ACBs environment, definition of the DBD-type ACB was not found.

System action

IMS Buffer Pool Analyzer ends the BPLTRACE task and the trace process.

User response

Contact the IMS administrator at your site for assistance. If no error is found in your IMS system, contact IBM Software Support.

BPL730E	DBD VALIDATION FAILED IN IMS MANAGED ACBS FOR DBD=*dbdname* (*rsn*)

Explanation

The validation of the indicated ACB (*dbdname*), which is managed by IMS, failed. The reason for the failure is indicated by the reason code (*rsn*).

System action

IMS Buffer Pool Analyzer ignores the database with the error and continues processing.

User response

Contact the IMS administrator at your site for assistance. If no error is found in your IMS system, contact IBM Software Support.

Abend codes

Abend codes are issued by the data-gathering or trace components of IMS Buffer Pool Analyzer. The technical information in this section can help you troubleshoot and diagnose IMS Buffer Pool Analyzer problems.

The following abend codes can occur during execution of IMS Buffer Pool Analyzer:

U4081
> User 4081 abends are issued by the data-gathering or trace components of IMS Buffer Pool Analyzer.
>
> These abends occur only in the BPLTRACE started task and do not affect the job or IMS subsystem being monitored.

Messages written to the MVS SYSLOG before the abend indicate the reason for the abend.

U4082

User 4082 abends are issued by the reporting components of IMS Buffer Pool Analyzer. Messages indicating the reason for the abend are written to the print output.

Gathering diagnostic information

Before you report a problem with IMS Buffer Pool Analyzer to IBM Software Support, gather the appropriate diagnostic information.

Procedure

Provide the following information for all IMS Buffer Pool Analyzer problems:

- A clear description of the problem and the steps that are required to re-create the problem
- All messages that were issued as a result of the problem
- A complete log of the job
- If an abend has occurred, the abend dump
- Product release number and the number of the last program temporary fix (PTF) that was installed
- The version of IMS that you are using and the type and version of the operating system that you are using

Provide the following information if the report content is incorrect:

- All reports generated by IMS Buffer Pool Analyzer
- GTF trace data sets that are used as the input for the report job

Notices

This information was developed for products and services offered in the U.S.A.

This material may be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing
IBM Corporation
North Castle Drive

Armonk, NY 10504-1785
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at http://www.ibm.com/legal/copytrade.shtml.

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

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions:

Applicability: These terms and conditions are in addition to any terms of use for the IBM website.

Personal use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights: Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED,

INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at http://www.ibm.com/privacy and the section titled "Cookies, Web Beacons, and Other Technologies" in IBM's Online Privacy Statement at http://www.ibm.com/privacy/details. Also, see the "IBM Software Products and Software-as-a-Service Privacy Statement" at http://www.ibm.com/software/info/product-privacy.

Index

Numerics

440 records 4

A

abend codes 118
ACBLIB member, valid 66
access method
    change 74
accessibility
    overview 8
add definitions 74
allocating data sets 9
alternate buffer pool configuration
    IMSplex 62
    multiple subsystems 62
alternate configuration
    model 74
    modeling report 33
    reports 75
APF authorization 9
authorization, APF 9

B

batch jobs 66
batch reporting process 60
batch reporting utility
    reports generated by 27
BATCH*nnn* parameter 76
benefits and features 3
BFR pool analysis 65
block analysis
    control statements 22
    report 21
    reports DD statements 22
block analysis report
    GTF trace report 23
    SYSPRINT report 23
block size
    changing 74
    report column 66
block size of a database
    change 74
blocks
    lost 34, 65
BPLCF coupling facility report 60
BPLCF JCL 60
BPLGTF
    PROC 12, 19
    procedure, customizing 10
BPLMAIN
    reports 6
    sample JCL 28
BPLTRACE
    bypassing delay time 18

BPLTRACE *(continued)*
    controlling trace activity 17
    customizing 10
    gathering information 52
    getting status 17
    GTF (generalized trace facility) functions 17
    overhead 18
    parameters 15
    started task 15, 17
    stopping 18
buffer allocations, reviewing 54
buffer pool
    analysis 59
    configuration, changing 74
    DD names 32
    reports output 32
    reports, customizing 29
    reports, IMS 27
buffer pool sizes, reviewing 53
buffer pools
    defining 50
    organization 50
    OSAM organization 50
    tuning database 49
buffers, number of 59
BUFNO statement 63

C

cache
    details 68
    enabling 75
    OSAM 75
    reports 68
    size varying 70
cached
    subpools 70
calculations
    buffer life 6
    hit ratio 6
    marginal reduction 6
CF parameter 63
CFNAMES statement 58
CFRM (coupling facility resource management) policy 76
CFSizer tool 58
change block size 74
change CI size data sets 74
CHANGE DBD statements 75
change definitions 74
CHANGE statement 67
changes, summary 1
CI size data sets, change 74
combining multiple 15
configuration
    modeling alternate 33
    OSAM cache 75
    subpool cache 75

organization *(continued)*
      buffer pools, subpools 50
OSAM / VSAM 65
OSAM buffer pools
      defining 50
      organization 50
OSAM cache
      configuration 75
      configuration, changing 60
      IMS, using 59
      overview 57
      performance 74
      structure 58
      structure report 70
      trace data 59
OSAM cache (IMS)
      implementing 57
OSAM cache configuration
      change 74
OSAM cache reports, coupling facility 57
OSAM cache structure detail report 76
OSAM cache structure size, select 76
OSAM caching configuration 60
OSAM coupling facility
      cache, tuning 75
OSAM coupling facility configuration 68
OSAM data sets 66
OSAM database buffer subpools subsystem 70
OSAM parameter 63
OSAM structure storage sizing 58
OSAM_STR_SIZE statement 62
overhead, trace 18
overview
      coupling facility IMS OSAM cache 57
      database buffer pool reports 27
      prerequisites 4
      process flow 5
      tuning process 52

P

paging activity, reviewing 52
parameters
      DT 15
      for started task 15
      GT 15
      GTFPROC 15
      IMSID 15
      JOB 15
      RGN 15
      WT 15
prerequisites 4
procedures
      BPLGTF 10
      BPLTRACE 10
process flow 5
product documentation terms and conditions 121

R

ratios
      directory 58, 76
      element 58, 76

ratios *(continued)*
      models of hit 76
reader comment form 7
records
      440 4
      F4F 4
      F61 4
      lost 34, 65
      lost records
            reducing size 37
      reducing number lost 37
      trace 4
remove definitions 74
report output, coupling facility report print DD 64
reports
      alternate configuration 75
      block analysis 21
      block analysis reports 22
      BPLCF (IMS Buffer Pool Analyzer coupling facility) 60
      buffer pool 29
      buffer pool DD names 32
      buffer pool, IMS 27
      combining multiple traces into one 19
      configuration modeling statistics, alternate
            statistics, configuration modeling 33
      coupling facility 59, 60
      coupling facility cache 67
      coupling facility IMS OSAM cache 57
      coupling facility report 76
      data set information 37
      database block analysis 23
      database buffer pool 28
      database buffer pool summaries 27
      database buffer subpool performance information and
      modeling 6
      database data set level 6
      database statistics 39
      database subpool statistics 42, 44
      Database subpool statistics 41
      database subpool summary 70
      databases with activity in the cache 70
      DFSVSAMP
            configuration 33
            statement 33
      generated by batch reporting utility 27
      GTF trace 65
      GTF trace data set 6
      GTF trace report 23
      GTF trace statistics 34
      hit ratio calculation 6
      hits per second 6
      IMS Buffer Pool Analyzer coupling facility (BPLCF) 60
      marginal reduction 6
      model of varying cache sizes 70
      OSAM cache structure detail 76
      OSAM coupling facility cache detail 68
      OSAM data set information 66
      output, buffer pool 32
      reads per second 6
      requests per second 6
      Subpool block analysis 24
      subpool detail 45, 71
      SYSOUT 34
      SYSPRINT 23

tuning *(continued)*
    reviewing subpool buffer allocations 54
    validating results 56
tuning database buffer pools 49
type, unknown GTF record 65

U

uncached subpools 70
unknown GTF record type 65

V

varied hit ratio report 43
virtual storage utilization information 52
VSAM / OSAM 65
VSAM buffer pools
    defining 50
VSAM buffer poolsbuffer pools
    organization 50
VSAM cross invalidation (XI) structure 58
VSAM organization 50
VSAM parameter 63

W

WT parameter 15

IBM®

Product Number: 5697-H77