International Technical Support Organization

SG24-4611-00

DB2 for MVS/ESA Version 4
Data Sharing Performance Topics

December 1995

IBM

International Technical Support Organization
San Jose Center

International Technical Support Organization

SG24-4611-00

DB2 for MVS/ESA Version 4
Data Sharing Performance Topics

December 1995

Take Note!

Before using this information and the product it supports, be sure to read the general information under "Special Notices" on page xiii.

First Edition (December 1995)

This edition applies to DB2 Version 4, Program Number 5695-DB2.

Order publications through your IBM representative or the IBM branch office serving your locality. Publications are not stocked at the address given below.

An ITSO Technical Bulletin Evaluation Form for reader′s feedback appears facing Chapter 1. If the form has been removed, comments may be addressed to:

IBM Corporation, International Technical Support Organization
Dept. 471 Building 80-E2
650 Harry Road
San Jose, California 95120-6099

Abstract

This document describes availability and performance enhancements in Version 4 of IBM DATABASE 2 for MVS/ESA (also called DB2 for MVS/ESA) introduced by data sharing in a parallel sysplex environment. This document discusses at length many tools that can be used to measure performance in this environment. Design and tuning recommendations are based on performance measurements conducted at the IBM Santa Teresa Laboratory.

This document is intended for customers and technical professionals who want to take advantage of the availability and performance enhancements in Version 4 introduced by data sharing. A background knowledge of DB2 data sharing is assumed.

(185 pages)

Contents

Figures

Tables

Special Notices

This publication is intended to help customers and technical professionals understand and evaluate DB2 for MVS/ESA Version 4 data sharing performance. The information in this publication is not intended as the specification of any programming interfaces that are provided by Version 4. See the PUBLICATIONS section of the IBM Programming Announcement for the DB2 product for more information about what publications are considered to be product documentation. Also see "Related Publications" on page xvi for a list of some publications.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM′s product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM′s intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter 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 the IBM Director of Licensing, IBM Corporation, 500 Columbus Avenue, Thornwood, NY 10594 USA.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer′s ability to evaluate and integrate them into the customer′s operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

Throughout this book, the licensed program DB2 for MVS/ESA and a particular DB2 subsystem are each referred to as ″DB2″. In each case, the context makes the meaning clear. The term MVS is used to represent MVS/Enterprise Systems Architecture (MVS/ESA). CICS is used to represent CICS/MVS and CICS/ESA; IMS is used to represent IMS/ESA.

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

CICS	DATABASE 2
DB2	DFSMS
DFSMS/MVS	DFSMSdss
DFSORT	DRDA
IBM	MVS/ESA
OS/2	RACF

RMF	S/390
Sysplex Timer	System/390
VTAM	3090

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc.

PC Direct is a trademark of Ziff Communications Company and is used by IBM Corporation under license.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Windows is a trademark of Microsoft Corporation.

Other trademarks are trademarks of their respective companies.

Preface

This document describes availability and performance enhancements in Version 4 of IBM DATABASE 2 for MVS/ESA (also called DB2 for MVS/ESA) introduced by data sharing in a parallel systems complex (sysplex) environment. This document discusses at length many tools that can be used to measure performance in this environment. Design and tuning recommendations are based on performance measurements conducted at the IBM Santa Teresa Laboratory.

This document is intended for customers and technical professionals who want to take advantage of the availability and performance enhancements in Version 4 introduced by data sharing. A background knowledge of DB2 data sharing is assumed.

How This Document is Organized

The document is organized as follows:

- Chapter 1, “Introduction”

  This chapter discusses alternative architectures and why DB2 uses the Shared Data architecture. The data sharing performance objectives and the measurement criteria are outlined.

- Chapter 2, “Sysplex Overview”

  This chapter provides an overview of the sysplex environment and shows how DB2 uses this environment.

- Chapter 3, “Data Sharing Overview”

  This chapter describes DB2 data sharing concepts and introduces data sharing terminology.

- Chapter 4, “Locking in a Data Sharing Group”

  This chapter describes DB2 data sharing locking concepts in more detail.

- Chapter 5, “Measurements Analysis”

  This chapter describes the results of the measurements of the IBM Relational Warehouse Workload (IRWW) by the DB2 Performance Group at the IBM Santa Teresa Laboratory.

- Chapter 6, “The IBM Relational Warehouse Workload in Parallel Sysplex”

  This chapter analyzes the performance profile of DB2 data sharing.

- Chapter 7, “Planning and Design Guidelines”

  This chapter describes the factors that affect performance in a data sharing environment. The information provided in this chapter should enable you to make informed decisions in order to derive maximum benefit in this environment.

- Chapter 8, “Instrumentation and Tuning”

  This chapter provides information about where to start and what to look for when tuning the performance in a data sharing environment. Information is also provided about tools that are widely used for this purpose.

Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this document.

- *IBM DATABASE 2 for MVS/ESA Version 4 Release Guide*, SC26-3394-01
- *IBM DATABASE 2 for MVS/ESA Version 4 Data Sharing: Planning and Administration*, SC26-3269-01
- *IBM DATABASE 2 for MVS/ESA Version 4 Administration Guide*, SC26-3265
- *IBM DATABASE 2 for MVS/ESA Version 4 SQL Reference*, SC26-3270

IBM DATABASE 2 Performance Monitor (DB2 PM) Publications

- *IBM DATABASE 2 Performance Monitor for MVS (DB2 PM) Version 4 Online Monitor User's Guide*, SH12-6165
- *IBM DATABASE 2 Performance Monitor for MVS (DB2 PM) Version 4 Batch User's Guide*, SH12-6164
- *IBM DATABASE 2 Performance Monitor for MVS (DB2 PM) Version 4 Report Reference, Volume 1*, SH12-6163
- *IBM DATABASE 2 Performance Monitor for MVS (DB2 PM) Version 4 Report Reference, Volume 2*, SH12-6166

Sysplex Publications

- *MVS/ESA Sysplex Overview*, GC28-1208
- *MVS/ESA Sysplex Systems Management*, GC28-1209
- *MVS/ESA Sysplex Hardware and Software Migration*, GC28-1210
- *MVS/ESA Sysplex Application Migration*, GC28-1211
- *MVS/ESA Setting Up a Sysplex*, GC28-1449
- *MVS/ESA SP V5 Programming: Sysplex Services Guide* , GC28-1495-02
- *MVS/ESA SP V5 Programming: Sysplex Services Reference* , GC28-1496-02

Resource Measurement Facility (RMF) Publications

- *RMF User's Guide*, GC28-1058
- *RMF Reference Summary*, SX33-9031

International Technical Support Organization Publications

- *DB2 V4 Non-Data-Sharing Performance Topics*, SG24-4562
- *DB2 PM Usage Guide Update*, SG24-2584
- *MVS/ESA SP Version 5 Sysplex Migration Guide*, SG24-4581
- *System/390 MVS Parallel Sysplex Performance*, GG24-4356
- *DB2 V3 Performance Topics*, GG24-4284
- *DB2 V2.3 Nondistributed Performance Topics*, GG24-3823

- *Capacity Planning for DB2 Applications*, GG24-3512
- *DB2 V2.2 Design Guidelines for High Performance*, GG24-3383

A complete list of International Technical Support Organization publications, known as redbooks, with a brief description of each, can be found in:

> *International Technical Support Organization Bibliography of Redbooks,* GG24-3070.

To get a catalog of ITSO redbooks, VNET users may type:

```
TOOLS SENDTO WTSCPOK TOOLS REDBOOKS GET REDBOOKS CATALOG
```

A listing of all redbooks, sorted by category, may also be found on MKTTOOLS as ITSOPUB LISTALLX. This package is updated monthly.

How to Order ITSO Redbooks

IBM employees in the USA may order ITSO books and CD-ROMs using PUBORDER. Customers in the USA may order by calling 1-800-879-2755 or by faxing 1-800-445-9269. Almost all major Credit Cards are accepted. Outside the USA, customers should contact their local IBM office. Guidance may be obtained by sending a PROFS note to BOOKSHOP at DKIBMVM1 or E-mail to bookshop@dk.ibm.com.

Customers may order hardcopy ITSO books individually or in customized sets, called GBOFs, which relate to specific functions of interest. IBM employees and customers may also order ITSO books in online format on CD-ROM collections, which contain redbooks on a variety of products.

ITSO Redbooks on the World Wide Web (WWW)

Internet users may find information about redbooks on the ITSO World Wide Web home page. To access the ITSO Web pages, point your Web browser (such as WebExplorer from the OS/2 3.0 Warp BonusPak) to the following:

```
http://www.redbooks.ibm.com/redbooks
```

IBM employees may access LIST3820s of redbooks as well. Point your web browser to the IBM Redbooks home page:

```
http://w3.itsc.pok.ibm.com/redbooks/redbooks.html
```

Acknowledgments

This project was designed and managed by:

- Ravi Kumar, International Technical Support Organization, San Jose Center

The authors of this document are:

• Hans Duerr	IBM Spain
• Yoshio Ishii	IBM Brazil
• Rolf Loeben	IBM Germany.

This publication is the result of a residency conducted at the International Technical Support Organization, San Jose Center.

Thanks to the following people for the invaluable advice and guidance provided in the production of this document:

• Ted Blank	IBM Development, Santa Teresa
• Kim Braia	IBM Development, Santa Teresa
• John Campbell	IBM Development, Santa Teresa
• Viviane-Anavi Chaput	IBM France
• Curt Cotner	IBM Development, Santa Teresa
• Robert Gilles	IBM Development, Santa Teresa
• James Guo	IBM Development, Santa Teresa
• Dale Hoyt	IBM Development, Santa Teresa
• Jeff Josten	IBM Development, Santa Teresa
• Jeanne Mizuno Kays	IBM Development, Santa Teresa
• Gary King	IBM S/390 Performance Design, Poughkeepsie
• Gopal Krishnan	IBM Development, Santa Teresa
• Bryce Krohn	IBM Education System Center, San Jose
• Christina Lee	IBM Development, Santa Teresa
• Bob Lyle	IBM Development, Santa Teresa
• Roger Miller	IBM Development, Santa Teresa
• Chris Mills	IBM United Kingdom
• Jerry Mukai	IBM Development, Santa Teresa
• Randy Nakagawa	IBM Development, Santa Teresa
• Main Nguyen	IBM Development, Santa Teresa
• Fred Orosco	IBM Development, Santa Teresa
• Tom Russell	ITSO, Poughkeepsie Center
• Akira Shibamiya	IBM Development, Santa Teresa
• Bryan Smith	IBM Development, Santa Teresa
• Hugh Smith	IBM Development, Santa Teresa
• James Teng	IBM Development, Santa Teresa
• Horacio Terrizzano	IBM Development, Santa Teresa
• Annie Tsang	IBM Development, Santa Teresa
• Jane Wang	IBM Development, Santa Teresa
• Jay Yothers	IBM Development, Santa Teresa
• Bob Young	IBM Development, Santa Teresa
• Kathy Zeidenstein	IBM Development, Santa Teresa.

Thanks to Shirley Weinland Hentzell for editing the document.

Chapter 1. Introduction

IBM DATABASE 2 for MVS/ESA Version 4 (DB2 V4) the eighth release of DB2 for the MVS platform, introduces several enhancements in the areas of client/server processing, availability, performance, and user productivity. The overall objective of these enhancements is to provide DB2 users with improved performance of critical applications in an operational environment very close to 24 hours a day, 7 days a week.

The DB2 V4 environment is available only for the MVS/ESA platform either to new users of DB2 or to current users migrating from Version 3 subsystems.

DB2 V4 introduces support for data sharing. The DB2 data sharing function addresses price-to-performance aspects, availability, and scalability requirements with an integrated solution based on System/390 parallel sysplex and system software. The performance aspects of DB2 data sharing are covered in this document. The performance aspects of DB2 non-data-sharing are covered in a separate document.

1.1 Requirements Addressed

Data Sharing introduced by DB2 for MVS/ESA Version 4 addresses the following requirements:

- Improved price-to-performance ratios

  The S/390 CMOS microprocessor technology is leveraged to provide better price-to-performance ratios.

- Good scalability characteristics

  The coupling technology is used to grant the majority of locks synchronously and to provide efficient cross invalidation mechanisms for down level and refresh of buffers cached locally. These factors enable data sharing performance for mixed workloads with heavy inter-system read-write interest to scale well with a relatively flat overhead as additional members are added to the group.

- Flexible growth

  Parallel sysplex gives you the ability to add processing power in more granular increments and to do so in a nondisruptive fashion. As processors are added to the configuration, the data does not have to be repartitioned in any way.

  Because DB2 data can be accessed from multiple DB2 members concurrently, you now have many more options to choose from in configuring your system environment. For example, you may choose to run your query and your online work on separate systems without having to do copy management of the DB2 data. Or, you may choose to activate a new DB2 member in the group temporarily to handle the extra workload of year-end or end of quarter processing.

- Larger capacity

  With a parallel sysplex, you can configure more millions of instructions per second (MIPS) and more central processors (CPs) in front of the data. This not only increases the processing power at your disposal to access the data, it also increases the degree of inter-SQL parallelism for your online transaction processing (OLTP) applications. S/390 central processing complexes (CPCs) can provide a degree of parallelism today within a single machine for the OLTP applications through the use of tightly coupled multiprocessors (MP). The parallelism is inter-SQL and covers both concurrent OLTP transaction and concurrent query. With a tightly coupled MP, you can run multiple OLTP applications in parallel on different CPs. Currently, the highest-end ES/9000 processor is a 10-way MP. Now, with a loosely coupled parallel sysplex configuration, the degree of parallelism can be greatly increased since many more CPs can be configured. Also, with the parallel sysplex the workload can be dynamically balanced across systems in the sysplex. Since each DB2 has equal access to the shared data (there is no affinity between a processor and a set of data), the system that can best run a new piece of work can be selected. Processes can be scheduled on any system in the parallel sysplex based on the prevailing available capacity across the systems. In contrast, with other architectures processes have to be scheduled based on the location of data. Thus, in other parallel database architectures where there is affinity between the processors and the data, one of the processors can become overloaded if access requests to its piece of the database are too frequent, resulting in what is called ″data-skew″ problem.

- Availability

  Continuous availability during both planned and unplanned outages is possible. Because there are multiple paths to the data, you can lose one of the DB2 members (from either a planned or unplanned outage), and still be able to access the data through the remaining DB2 members. Because there is no affinity between the processor and the data, the workload of the DB2 system that is down can be spread evenly across the other members. Two common examples of planned outages are changing the DB2 system parameters (ZPARM) and applying program temporary fix (PTF) maintenance.

- Avoid modifying applications

  DB2 applications do not need to be rewritten to take advantage of data sharing. Except for the addition of a few miscellaneous structured query language (SQL) ″resource unavailable″ types of reason codes, data sharing does not in any way affect the SQL interface. All the controls that ensure data integrity in a multisystem environment happen covertly in DB2.

1.2 Alternative Architectures for Mixed Workloads

There are several architectures that can be considered for mixed workloads.

- Shared Everything

  This is the current architecture used by S/390 prior to the advent of data sharing. Each processor can access the shared memory and the direct access storage device (DASD). A task can be dispatched on any one of the physical processors. It has many advantages that have been demonstrated convincingly over many years—dynamic workload balancing, reduced systems management, ease of exploiting new capacity, application investment, and others. Some additional requirements need to be addressed.

- Shared Disk

Each processor has its own memory, but the processors share common disks. This architecture does not scale well for OLTP and mixed workloads where the applications and data are interlinked, because of

  - Hot data spots
  - Logical and physical contention
  - Escalating cost of message passing.

- Shared Nothing

  Each processor has its own dedicated memory and DASD. This architecture is the easiest to describe and the easiest to understand. It is particularly well suited to the query environment where the data partitioning is optimal and there is no appreciable data-access skew. For OLTP and mixed workloads, strong affinity routing is key to good performance. But strong affinity routing is difficult to achieve for large-enterprise servers, as applications and data are interlinked. The problems with Shared Nothing systems are:

  - As there is affinity between a given processor and a particular subset of the data, a skew in the data-access patterns can result in overloading a processor.
  - Partitioning requires administrators to make choices that may compromise the incremental algorithms of different workloads.
  - Does not inherently provide parallelism within and between applications.

- Shared Data

  A radically new architecture which is a major extension to the shared-disk architecture model, Shared Data builds on the traditional advantages of the S/390 Shared Everything implementation. The Shared Data architecture is well suited to all types of workload, OLTP, query, and mixed. It can also take advantage as needed of strong affinity routing. This architecture offers a higher level of availability at reasonable cost, along with dynamic workload balancing and nondisruptive growth.

1.3 Architecture Used by DB2 Version 4 Data Sharing

DB2 Version 4 data sharing is implemented using the Shared Data architecture. In this architecture, each system is a symmetrical multiprocessor with its own local memory shared between each of the central processors. The DASD can be Enterprise Systems Connection (ESCON) connected and shared by all the systems. Each system is connected to the new coupling technology which is fast and efficient. The design goal is to reduce intersystem communications and provide highly efficient global mechanisms for data locking and data coherency. Data is processed locally on each system, with the coupling technology providing a high-speed group data cache.

Data sharing offers higher availability as there is no affinity between any one processor and piece of data. If you lose any system, you can still get to the data through any other system in the group. The workload of the failed system can be spread across the survivors, instead of running one system at half capacity so that it can act as a backup for another. The solution offers N+1 availability; that is, you can configure one extra system to guard against a single point of failure. Thus, you achieve high availability at low cost.

Since no processor is linked to any particular piece of data, incoming work can be assigned to the system that can best run the work at any given point in time.

Therefore, skews in the data access patterns do not result in one of the systems becoming overloaded.

Additional processing capacity can be added nondisruptively. Since there is no affinity between processors and data, you need not repartition the data as new systems are added.

The shared data architecture was designed specifically for the parallel sysplex environment. Use of the shared coupling facility by all systems is key. There are local DB2 optimizations, such as caching data locally for processing and using the coupling facility as a store-in cache. Also, not all locks are propagated beyond local system.

The shared data architecture extends the S/390 platform and structure, building on its strengths. There is full read-write capability across subsystems and CPCs. The solution offers near linear scalability and incremental growth. The key feature is dynamic workload routing around the sysplex based on available capacity and not on location of data. Thus, the shared data architecture offers high levels of availability at low cost and good balanced performance across the broad spectrum of application types including mixed workloads.

1.4 Performance Objectives

The performance objectives for IRWW support DB2 design and development. The design solution balances good performance with random scheduling of transactions across the data sharing group in support of N+1 availability. Optimizations can be made to isolate particular applications and their associated data on separate systems to reduce the costs of data sharing, but this would compromise the availability benefits.

There are three main performance objectives:

1. Near equivalent throughput where there is no intersystem read-write interest. For example, DB2 in a parallel sysplex where there is currently no data sharing should deliver performance nearly equivalent to that in a 1-way data sharing in a parallel sysplex environment, that is, there is no intersystem read-write interest.

2. Where there is heavy intersystem read-write interest in a transaction environment (that is, 100% data sharing) across both members in a 2-way data sharing configuration, restrict the initial throughput loss to at or below 17% when evaluating 2-way data sharing.

   The second objective translates to an additional CPU capacity delta of 17% to achieve equivalent throughput. This objective assumes 100% data sharing. In many customer environments, a 50% data sharing would be more reasonable.

   The CPU capacity delta of 17% required for data sharing includes a fixed portion of approximately 3% for MVS sysplex overhead. The 3% overhead is with the first generation CMOS processors (9672/Rx1) when migrating from a nonsysplex environment to a sysplex environment. The 3% is regardless of the load, so at 50% read-write data sharing, we would have 3% + (14% / 2) or 10% added capacity delta requirement.

3. Limit the incremental throughput loss per member when scaling above 2-way to at or below 0.5%.

As each additional member is added to the group, the aim is to limit the additional internal throughput loss on each member to approximately 0.5%.

These objectives assume a homogeneous CMOS environment with matching CP speeds for CPCs running MVS and coupling facility control program (CFCP), random scheduling of transactions and keys across the data sharing group.

1.5 Documenting Performance Measurements

The workload for the performance measurements is specifically designed to evaluate data sharing in OLTP environment. Non-data-sharing overhead is kept to a minimum. The objective of these measurements is to achieve a high level of data sharing, and not a high transaction rate. The performance throughput in your environment may vary.

The details of the environment and performance measurements are documented in Chapter 5, “Measurements Analysis” on page 39.

Chapter 2. Sysplex Overview

This chapter describes briefly the sysplex environment and shows how DB2 uses this environment.

2.1 Sysplex

A sysplex is a set of MVS systems communicating and cooperating with each other through certain multisystem hardware components and software services to process customer workloads.

A sysplex implementation can use a coupling facility. A coupling facility is a combination of hardware and software, which enables the communication of subsystems within a sysplex. A sysplex without a coupling facility is referred to as a *base sysplex*. When the implementation includes a coupling facility, it is called a *parallel sysplex*.

The parallel sysplex is emerging as the computing environment of the future. Sysplex offers improved price-to-performance ratios through cost-effective processor technology and enhanced software. This technology builds on existing data processing skills and continues to run existing applications.

A sysplex can increase the capacity available for important applications, remove the need to balance work manually between systems, improve availability, and help to deliver more consistent performance. A significant difference between the sysplex and a conventional MVS/ESA system is the improved growth potential and level of availability that it offers. The sysplex increases the overall processing capacity available to applications by harnessing multiple MVS/ESA operating systems together.

Up to 32 CPCs can participate in a sysplex. This cooperation increases the total amount of work that can be processed even when one of the CPCs reaches its capacity limitations. To make this participation possible, multiple components, both software and hardware, are required. Figure 1 on page 8 shows the components of a typical parallel sysplex.

A coupling facility enables parallel processing and improved data sharing for authorized programs running in the sysplex. The cross-system extended services (XES) component of MVS enables applications and subsystems to take advantage of coupling facility services.

Cross-system coupling facility (XCF) services allow multiple instances of an application or subsystem, running on different systems in a sysplex , to share status information and communicate with each other. MVS components also use XCF signaling services for communication in a sysplex. XCF paths can be implemented with channel-to-channel connections or through the coupling facility.

A sysplex timer is required to synchronize the time-of-day (TOD) clocks for systems in a sysplex that run on different CPCs.

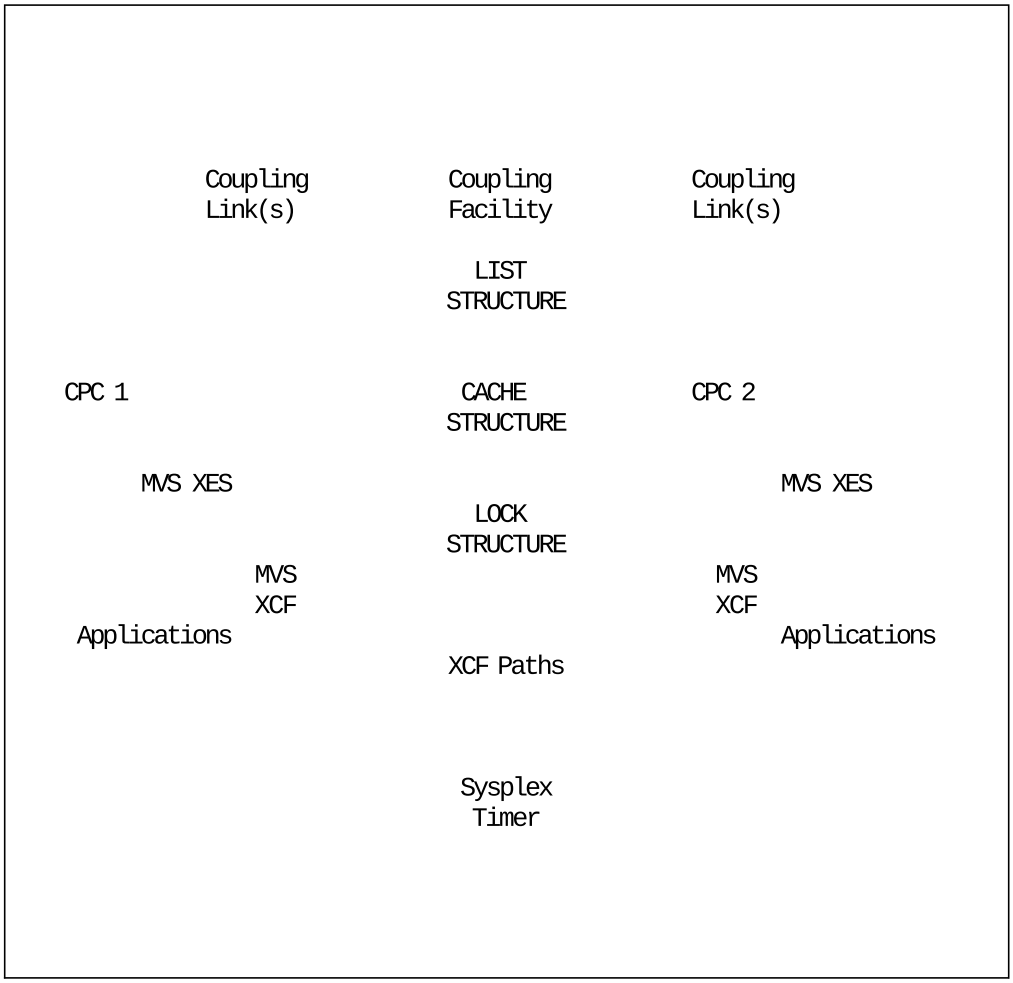

Figure 1. Parallel Sysplex Environment— Two Central Processing Complexes

2.1.1 Coupling Facility

Within the coupling facility, storage is dynamically partitioned into structures. MVS XES services manipulate data within these structures. Each of the structures has a unique function:

Lock structure Provides quick detection of lock contentions. Most locks can be granted synchronously without contention. Only a very small percentage of the total volume of lock requests are expected to have some form of lock contention. XCF intersystem message passing is required when there is lock contention.

Cache structure Used as a high-speed buffer for storing shared data with common read-write access. The cache structure also has a buffer invalidation mechanism to ensure consistency of locally cached data.

List structure Allows list format data sharing—for example for implementing functions such as shared work queues and shared status information.

Many more details on this subject are provided in *MVS/ESA SP Version 5 Sysplex Migration Guide*, Chapter 1, "Introduction to Sysplex."

2.1.2 DB2 in a Sysplex

DB2 data sharing takes significant advantage of sysplex services, specifically:

- The MVS XES services are used for communicating with the coupling facilities.
- The coupling facility structures are used as follows:
  - The lock structure is used to implement global locking for DB2 data sharing.

    All locks on shared resources must be known to all members. These locks are recorded in the lock structure.
  - The cache structure is used to store changed and optionally unchanged pages.

    The changed data can support a quick refresh of local buffers after a cross invalidation. This is particularly suited to frequently accessed data.

    The cache structure may optionally be used for clean (unchanged) pages to potentially reduce DASD input and output (I/O).
  - DB2 data sharing creates a shared communication area (SCA), in the list structure. The SCA is used to interchange information between sharing DB2 subsystems.
- XCF messages are used:
  - To communicate locking information between Internal Resource Lock Managers (IRLMs) and XESs
  - To notify other sharing members to retrieve database or system-status, or control-information changes made to SCA.
  - To communicate among sharing members for processing DB2 database commands and utilities.

Chapter 3. Data Sharing Overview

This chapter describes DB2 data sharing concepts and introduces data sharing terminology. Readers familiar with this subject may want to skip directly to Chapter 4.

For readers not familiar with this subject, this chapter may serve as an introduction to the following chapters and to the *DB2 for MVS/ESA V4 Data Sharing: Planning and Administration*.

3.1 Architecture

Before DB2 V4, only distributed data supported read and write access between multiple DB2 subsystems. With DB2 V4, data sharing allows read-write access to DB2 data from different DB2s residing in multiple CPCs in a parallel sysplex. For high transaction volumes, DB2 data sharing with a parallel sysplex is a better solution than distributed data. DB2 V4 uses the data sharing architecture to obtain high-performance and high-availability data sharing. The systems administration cost for DB2 data sharing is less than that for distributed data.

Each CPC needs an MVS instance and one or more DB2 instances to access DB2 data. The set of DB2 instances that access a specific set of shared data is called a "group." Each individual DB2 instance is called a "member" of the group. Only members of a group can share data in read-write mode. DB2 V4 does not support read-write data sharing between different groups.

DB2 data sharing is designed to work in a parallel sysplex. DB2 data sharing uses group buffer pools, a lock structure, and a shared communication area in the coupling facility to manage data integrity and provide data consistency in the data sharing environment.

3.1.1 Group

A data sharing group has one or more members. At the time of this publication the maximum is 32 members.

A group is established when the first member obtains data sharing functions either through installation or by means of migration. The DB2 catalog and directory of this first member become the catalog and directory of the group. When a new member is added to the group, it shares the existing catalog and directory.

From the point of view of data access, a group can be considered as one large DB2 subsystem that spans multiple CPCs—a multiple CPC DB2.

Figure 2 on page 12 shows a sysplex with five MVS instances (MVSa, MVSb, MVSc, MVSd, and MVSe) and eight DB2 subsystems (DB21 to DB28) making up three data sharing groups (G1, G2, G3):

- G1 has members DB21, DB25, and DB26.
- G2 has members DB22, DB23, DB27, and DB28.
- G3 has member DB24. It is either a single member-group, or a non-data-sharing DB2 subsystem, or DB2 V3 subsystem.

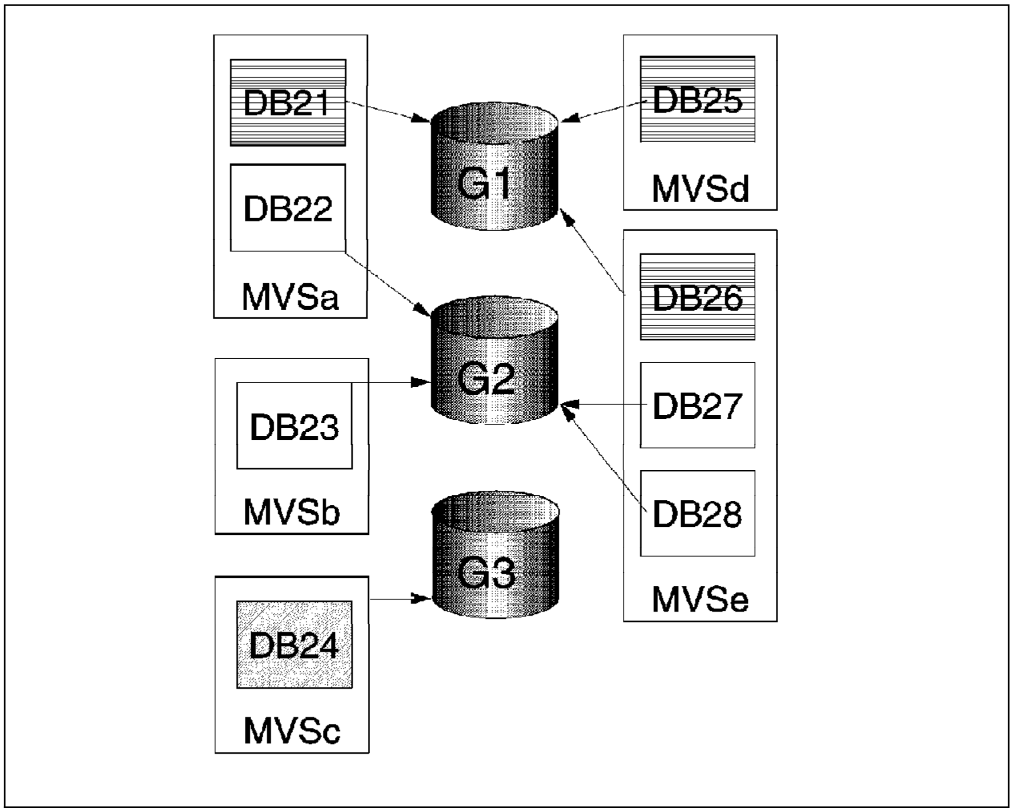

Figure 2. Three Data Sharing Groups

3.1.2 Member

A data sharing member is a DB2 subsystem.

- Each member has its specific subsystem identifier.
- Each member has its own IRLM.
- Each member has its own set of address spaces.
- Each member has its own bootstrap data set (BSDS) files.
- Each member has its own active log files.
- Each member has its own archive log files.
- Each member has its own workfile database.
- Each member has its own virtual buffer pools.
- Each member has its own optional hiperpools.
- Each member has its own environmental descriptor manager (EDM) pool.
- Each member has its own sort pool and rid pool.
- Each member has its own DSNZPARM definitions.

A data sharing member shares DB2 resources with other members of the group.

All members share the DB2 catalog and directory.
All members share the databases, table spaces, tables and indexes.
All members share the aliases, views, and synonyms.
All members share the plans and packages.
All members share the authorization.
All members share the image copies.
All members share in the result of any utility execution.
All members share the recovery information.
All members share the group buffer pools.
All members share the lock structure.
All members share the SCA.
All members share the application programming defaults contained within the DSNHDECP load module.

As a consequence of the sharing of resources:

- A data definition language (DDL) statement executed on any member creates an object, which all members can use.
- A plan bound on one member may be executed on another member.
- A grant to execute that plan can be executed on any member and authorizes execution on all members.
- Recovery of a table space can be executed on any one member and all members have the table space recovered.

Local Data for One Member Only: Each member may have local data which is defined as data never accessed from other members. Local data resides in table spaces assigned to a virtual buffer pool defined on one member only (no other members should have that virtual buffer pool defined). Other members will not be able to open that table space, because they do not have the virtual buffer pool for it. A group buffer pool (see section 3.2, "Group Buffer Pools" on page 14) should not be defined for this virtual buffer pool; it would never be used. DASD connectivity offers an alternative way of limiting data access to one DB2 member only.

Read-Only Sharing Within a Group: All members of a group may have table spaces assigned to their virtual buffer pools. If there is no group buffer pool for these virtual buffer pools, these table spaces can be shared by all members only in read mode.

Any one member may update these table spaces, provided they are closed for all other members.

See *IBM DATABASE 2 for MVS/ESA Version 4 Data Sharing: Planning and Administration*, Chapter 3, "Shared Read-Only Data."

Read-Only Sharing Between Groups: A group may use shared read-only data to share data with other groups. This is defined and works exactly like it is implemented in earlier DB2 Versions for shared read-only data between DB2 subsystems.

Read-Write Access between Groups: DB2 distributed data support provides full read-write access to data between different DB2/MVS groups, as well as to DB2 data on other platforms. DB2 distributed data features a full two-phase commit protocol with unit-of-work integrity across multiple systems.

3.2 Group Buffer Pools

Group buffer pools (GBP) are cache structures in the coupling facility, which are extensions of the local buffer pools in the individual members. A local buffer pool is a virtual buffer pool plus an optional hiperpool. There must be one GBP for each local buffer pool containing table spaces to be shared. For example, if a table space is assigned to virtual buffer pool BP4, each member must have the virtual buffer pool BP4 defined, in order to be able to read pages from that table space. To enable all members to share that table space, GBP4 (group buffer pool 4) must exist. Furthermore, any data sharing group must have GBP0, because the DB2 Catalog and Directory use BP0 and these must be shared by all members.

Group buffer pools are used in a data sharing environment:

- To store updated pages
- To register the pages stored in the local buffer pools of members in a "Page Directory."
- To optionally cache clean (unchanged) pages in the group buffer pool. This can potentially reduce DASD I/O and is especially useful in the OLTP environment for relatively small page sets where a very high hit ratio for reads can be achieved from accessing the GBP.

Updated pages are stored in the GBP to improve transaction response time at commit and to allow a quick refresh of these pages, without having to resort to DASD I/O. This is particularly well suited to frequently accessed (hot) data. Examples of hot data are:

- Space map pages
- Index leaf pages
- Small tables
- Large tables with hot spots.

The "page directory" in the GBP records (registers) the pages stored in the local buffer pools. Whenever a member reads a new page into the local virtual buffer pool from the GBP or from DASD, the interest in the page from that member must be registered in the page directory. This is necessary because another member may subsequently update the same page, causing all other locally cached versions of the page to be outdated. The page directory allows the management of the consistency of data pages within the local buffer pools of all members, (see section 3.2.4, "Buffer Pool Coherency" on page 16).

3.2.1 Inter-DB2 Read-Write Interest

Group buffer pools are used when two or more members are sharing the same page set or partition and at least one of the members has write interest. In this case, the page set or partition becomes "GBP-dependent." The condition that causes a page set or partition to become GBP-dependent is called "inter-DB2 read-write interest."

DB2 data sharing uses a dynamic mechanism to detect that sharing exists— that is, to detect inter-DB2 read-write interest. This mechanism works at the page set or partition level, with full partition independence and independence between indexes and data. This mechanism is explained in section 4.1.2, "Page set Physical Locks" on page 21.

3.2.2 Group Buffer Pool Dependence

DB2 data sharing performs the following functions for a GBP-dependent page set:

- Page registration

  Every page in the local buffer pool is registered in the page directory of the corresponding GBP.

- Force at commit processing

  At commit time, every changed page not yet written to the GBP is forced out into it. More details are given in section 3.2.3, "Force at Commit."

- Cross-system invalidation

  Whenever a changed page is written to the GBP, all copies of the page in the local buffer pools of other members are marked as invalid.

- Test for page validity

  Because of the page cross invalidation process, all pages in the local buffer pools must be tested for validity before use. See Table 1 on page 22 for an exception to the above rule. If all other members in the group are read-only, the one updating member does not check its pages for validity.

- Page refresh from GBP

  Instead of reading from DASD, pages can be obtained from the GBP. In order to optimize accesses to the coupling facility, XES obtains the page when the member registers interest in the page in the page directory.

- Castout processing

  Changed pages in the GBP are written to DASD.

- Use of a global commit log sequence number (GCLSN) for lock avoidance.

  Lock avoidance must work at group level.

3.2.3 Force at Commit

DB2 places changed pages for a page set or partition on a deferred-write queue. Typically, these pages are written to DASD using asynchronous deferred-write processing. In a DB2 data sharing environment, the changed pages for a GBP-dependent page set or partition must be written to the GBP that resides in a coupling facility structure. This write process must complete before transaction commit processing finishes. The locks protecting the changed pages cannot be released until the changed pages have been written to the coupling facility structure.

3.2.4 Buffer Pool Coherency

DB2 data sharing ensures buffers contain only valid pages. When one member updates a page, earlier copies of that page may exist in local buffer pools of other members. Those pages are no longer valid. The "cross-system invalidation" process invalidates all those pages. This process is called buffer pool coherency and is maintained across all local buffer pools and the corresponding GBPs.

The *DB2 for MVS/ESA V4 Data Sharing: Planning and Administration* has detailed description of buffer pool coherency and cross-system invalidation.

3.2.5 Group Buffer Pool Castout

The process of writing modified pages from the GBP to DASD is called group buffer pool castout. The DB2 member which first allocates the GBP structure is the GBP castout structure owner. Each GBP has a single castout structure owner and a series of castout classes. DB2 page set or partitions using that GBP are mapped to a specific castout class. The partitions of partitioned table spaces and indexes can be mapped across different castout classes.

GBP castout ownership is assigned at the data-set (page set or partition) level. GBP castout ownership at the data-set level is assigned to a particular member of the group (the first member to express write interest in the data set).

A GBP checkpoint is triggered by the GBP checkpoint interval. At that time the GBP structure owner notifies the castout class owners to do the actual castout process. Thus, at GBP checkpoint time, castout is triggered for all changed pages across all data sets.

You can control the castout process by changing the two GBP thresholds:

- GBP castout threshold
- Class castout threshold.

If the GBP castout threshold is exceeded, castout is triggered for selected castout classes, until sufficient pages are written out.

When the castout class threshold is triggered, the member who is the castout class owner performs the castout for the page sets that are in the specific castout class.

The *IBM DATABASE 2 for MVS/ESA Version 4 Data Sharing: Planning and Administration* has detailed description of the castout process.

3.3 Global Locking

For data consistency in a data sharing environment, locks must be known and respected between all members. DB2 data sharing uses global locks to ensure that each member is aware of all members′ locks.

Two locking mechanisms are used by DB2 data sharing to ensure data consistency, logical locks and physical locks.

The two types can be briefly compared as follows:

- Logical locks

  Logical locks are used to control concurrent access from application processes, such as transactions or batch programs.

- Physical locks

  Physical locks are used by DB2 members to control physical resources:

  – Page set physical locks are used to track the level of interest in a particular page set or partition and thus determine the need for GBP coherency controls.

  – Page physical locks are used to preserve the physical consistency of pages.

A more detailed definition of these two mechanisms is provided in Chapter 4, “Locking in a Data Sharing Group” on page 19.

3.3.1 Global Lock Manager

Each exclusive lock (X-lock) for a GBP-dependent page set is registered in the coupling facility lock structure. When members have interests that conflict with that exclusive lock, a global lock manager is established and given ownership of that specific lock. The global lock manager resolves contentions on that lock. More details about global lock managers are provided in section 4.3, “Global Lock Manager” on page 26.

3.3.2 Contentions

A contention can occur when more than one member needs access to one resource. The accesses may be incompatible, or the coupling facility locking structure may not be able to determine that compatibility exists. When a contention occurs, the lock request is suspended until the contention is resolved.

Three types of contention can occur in a data sharing environment:

- False contention
- XES contention
- Real contention.

These contentions are described in detail in section 4.4, “Contentions” on page 27.

Chapter 4. Locking in a Data Sharing Group

This chapter describes DB2 data sharing locking concepts in more detail.

Additional information can be found in the *IBM DATABASE 2 for MVS/ESA Version 4 Data Sharing: Planning and Administration*.

4.1 Global Locking

For data consistency in a data sharing environment, locks must be known and respected among all members. DB2 data sharing uses global locks to ensure that members are aware of each others′ locks.

In DB2 data sharing, all locks on data base resources are global locks. A global lock is one that guarantees serialization across the data sharing group. In contrast a local lock guarantees serialization for one member only. Some global locks are propagated from the local IRLM through XES to the coupling facility (CF), and some global locks do not need to be propagated.

DB2 data sharing uses two types of global locks:

- Logical Locks
- Physical Locks

It is important to mention that the IRLM keeps track of lock modes associated with multiple concurrent transactions running against a particular member and propagates only the most restrictive state beyond the local IRLM.

Logical Locks: DB2 uses transaction locks for data consistency. With DB2 data sharing, these locks are also called logical locks, or L-locks. Logical locks are used for intertransaction concurrency controls within a data sharing group. Logical locks are owned by a transaction; the lock duration can be commit, plan deallocation, or manual. Logical locks are not negotiable and can have deadlocks and timeouts.

The logical locks are of four types:

- Table space locks
- Table locks
- Page locks
- Row locks.

Physical Locks: DB2 data sharing introduces physical locks, or P-locks. Some physical locks are used to ensure inter-DB2 data coherency, others are used to ensure physical data consistency. Physical locks are owned by a DB2 member; duration is determined by interest of the member. Physical locks are negotiable and have no deadlocks, nor timeouts. From the point of view of mainline SQL performance, the most important physical locks are page set physical locks and page physical locks. The difference is in their uses:

- Page set physical locks are used to track inter-DB2 read-write interest at the page set or partition level.
- Page physical locks are used to provide write serialization between members, ensuring the physical consistency of a page. Page physical locks can be thought of as intermember page latches.

Both are further explained in section 4.1.2, "Page set Physical Locks" on page 21 and section 4.1.3, "Page Physical Locks" on page 22.

Physical locks can be subject to negotiation between members. A negotiation can lead to a less restrictive physical lock, allowing another member access to the same resource.

4.1.1 Table Space Logical Locks

DB2 uses table space locks to verify access compatibility between different processes within a DB2 subsystem. For example, if one transaction has a table space intent share (IS) lock, a batch program requiring an exclusive (X) lock on the same table space is not allowed to run. In a data sharing environment these locks need to be known by all members; they must be registered in the lock structure of the coupling facility.

The parent logical locks are used to:

- Control concurrency between application processes in all members
- Allow IRLM to decide which child locks to propagate to the lock structure when there is inter-DB2 read-write interest (see "Explicit Hierarchical Locking").

For a partitioned table space, touching one partition causes partition logical locks to be acquired for all partitions of the subject table space.

Explicit Hierarchical Locking: DB2 data sharing introduces a concept called "explicit hierarchical locking" to optimize performance by reducing the number of logical locks to be propagated. Hierarchical locking relates parent locks (table space logical locks or partition logical locks) to child locks (table, page, or row logical locks). Explicit hierarchical locking allows IRLM to relate page or row locks to their parent logical locks. Thus, IRLM can decide which locks to propagate to the lock structure in the coupling facility. The *DB2 for MVS/ESA V4 Data Sharing: Planning and Administration* explains hierarchical locking in detail in section *Explicit Hierarchical Locking*.

Drains and Claims: Drains and claims introduced in DB2 V3 to support partition independence, are used in a data sharing environment.

A drain on a page set or partition is a global logical lock. Drain logical locks always propagate to the coupling facility. This is true even for non-GBP-dependent page sets or partitions.

The first claimer obtains a share (S) logical lock against that page set or partition to serialize on the drain. Once the first claim is obtained by a member, claim counters for that page set or partition are maintained locally in each member.

Member Level Logical Locks: To optimize coupling facility access, DB2 tries to reduce the number of locks propagated to the coupling facility. If a resource is used for multiple transactions by one member, only the most restrictive logical lock needs to be propagated. All other locks are compatible and do not affect concurrent access to the table space from other members.

For example, assume two transactions are being executed by one member. The first transaction has an intent exclusive (IX) lock and the other has an IS lock on a table space. The member needs only to propagate the IX (more restrictive) logical lock for that table space to the lock structure.

4.1.2 Page set Physical Locks

Page set physical locks are used to track inter-DB2 read-write interest of different members for a page set or partition. Each member acquires one (and only one) page set physical lock for each open table space or index, or for any partition of a partitioned table space or partitioned index. Because the physical locks are owned by the member, they are independent of the transactions.

Always Propagated to Lock Structure: Because page set physical locks are used to track inter-DB2 read-write interest on a specific page set, they are always propagated to the lock structure. Even for a group with one member only, all page set physical locks are propagated to the coupling facility.

For a partitioned page set, the partition physical locks are acquired for every open partition. For a nonpartitioned page set, the lock is acquired at the page set level even for a multidataset page set.

Member Level Physical Locks: Page set physical locks are propagated at member level. Each member propagates only one page set physical lock per table space to the coupling facility. See "Member Level Logical Locks" on page 20 for a description and an example.

Page set Physical Lock Mode Changes: The inter-DB2 interest of a page set physical lock for a given page set or partition by a given member is determined by the following activities on the page set or partition:

- Physical open, read-only (RO) interest
- Physical close, no interest
- First update or pseudo-open, read-write (RW) interest
- RO-switching or pseudo-close, RO interest

Switching of a page set or partition between RO and RW interest affects the mode of page set physical locks and the GBP dependency. This switching is caused by RO-switching and by updates to the page set or partition. Installation parameter (PCLOSET, PCLOSEN) options allow a specification of time interval and number of checkpoints to initiate the check for RO-switching.

Low values for the PCLOSET and PCLOSEN parameters may cause frequent switching in and out of GBP dependence. On the other hand setting these parameters too high could mean GBP-dependent processing for a longer period of time than is really needed. Every situation must be analyzed individually, based on workload and access profiles, to establish a good balance between RO-switching and the data sharing overhead.

Table 1 shows when the mode of a page set physical lock changes and when the GBP-dependent transitions occur. Table 1 indicates the modes of the page set physical lock obtained in all possible cases of request combinations.

Table 1. Page set Physical Lock Modes for GBP-Dependence

Interest of Requesting DB2	Interest of Other Members	Page set Physical Lock Mode	GBP Dependent	Check Page Validity by Requester
Read-Only	NONE, Read-Only	Shared (S)	NO	NO
Read-Only	Read-Write	Intent Share (IS)	YES	YES
Read-Write	NONE	Exclusive (X)	NO	NO
Read-Write	Read-Only	Share with Intent Exclusive (SIX)	YES	NO/YES
Read-write	Read-Write	Intent Exclusive (IX)	YES	YES

Note: Page validity checking for SIX mode depends on the previous mode of the page set physical lock. When the change is from X mode to SIX mode, page validity is not checked; changing from IX mode to SIX mode, page validity is checked.

4.1.3 Page Physical Locks

Page physical locks are used to ensure the physical consistency of a page across members of a group. Members with read interest are assured the most current consistent version of the page. Page physical locks are used only when the page set or partition is GBP-dependent, when subpage concurrency (for example, a type 1 index subpage or a portion of a data page because row locking is in effect) is possible, or when no page logical lock is protecting the page. For example, if row locking is used and a row is updated, an X-mode page physical lock is acquired to ensure that no other member touches the page until the update is recorded. The page physical lock protects the page while its structure is being modified. The X-mode page physical lock is released when the page is written to the GBP. This can be before or during the commit of the transaction that caused the page to be modified.

X-Mode Page Physical Locks: The following are examples for X-mode page physical locks:

- Update of a GBP-dependent data page, while row locking is in effect
- Update of a GBP-dependent space map page
- Update of a GBP-dependent type 2 index page
- Update of certain DB2 directory pages.

S-Mode Page Physical Locks: S-mode page physical locks are required only with:

- Repeatable read (RR) index scan with LOCKSIZE ROW for data pages and type 2 index leaf pages
- Referential integrity (RI) checking.

Sometimes S-mode page physical locks are used for space map pages. Reader processes using cursor stability (CS) isolation never take physical locks on index or data pages.

4.2 Lock Structure

The lock structure in the coupling facility has two components:

1. A lock table, which is used to detect contentions.
2. A modified resource list which saves information on modify locks, which can be either logical or physical locks. These locks may become retained locks.

Figure 3 shows these components and their entries. Both components use approximately the same amount of space.

The lock table has a binary number of entries, the example illustrated in Figure 3 shows exactly 128 entries, which is 2 to the power of 7 (2**7). A real lock structure needs many more lock table entries and modified resource list entries than shown in Figure 3.

More details on these lock structure components are provided in the following subsections.

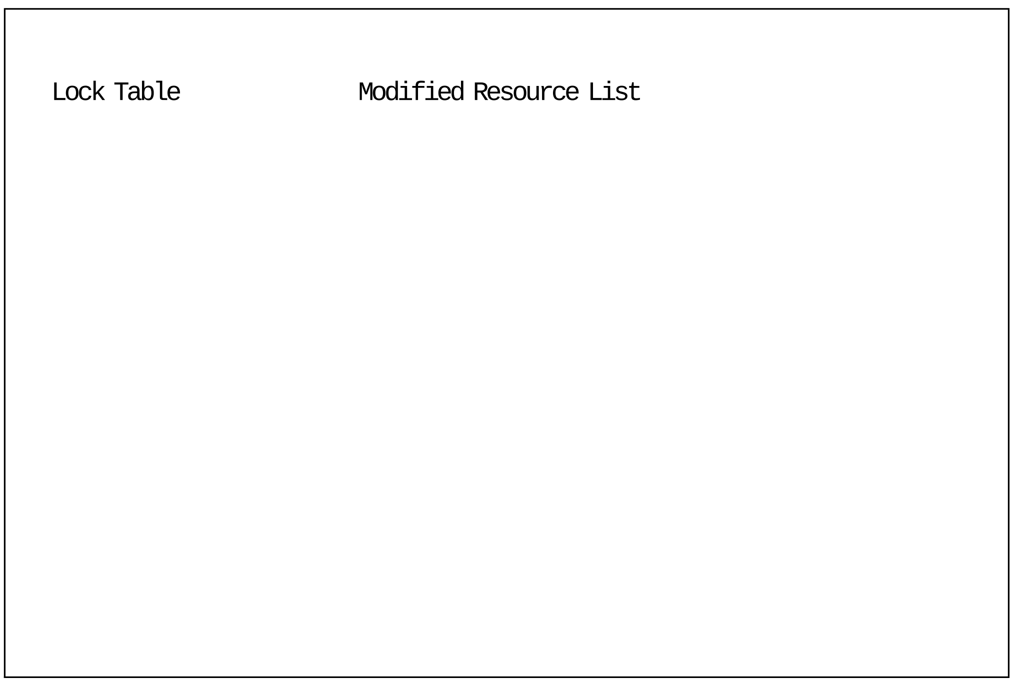

Figure 3. Lock Structure. Lock table and modified resource list have similar size.

4.2.1 Lock Table

The lock table, also called the “hash table,” is approximately one-half of the lock structure of the coupling facility. The lock table contains fixed-size entries, called “lock table entries.” A lock-table entry records the first system to have exclusive interest in the resource for which the lock table entry is made and what other systems have shared interest in the lock table entry for any resource. The width of a lock table entry can be 2, 4, or 8 bytes for up to 32 members. The purpose of a binary width is to enable binary addressing.

Figure 4 shows a lock table entry with 2-byte width. The 2-byte-wide entry can support groups with up to six members. Figure 5 on page 25 shows a lock table entry with a 4-byte width. The 4-byte-wide entry can support groups with up to 22 members. Each increase in the size of a lock table hash entry from 2 bytes to 4 bytes or from 4 bytes to 8 bytes, is accompanied by a halving of the number of entries the lock structure can hold and, therefore, an increase in the potential for contentions.

A calculated hash value is used to assign a resource to a lock table entry. The hash value is calculated by DB2 from the name of the resource to be locked. The hash value is a 32-bit value passed through IRLM to XES. XES uses a modulus based on the number of lock table entries available for the 32 bit hash value to assign a resource name to a particular lock table entry. The lock table entry registers S-mode interest with one bit for each member of the group holding the lock. Only one member can register an X-mode interest.

The example in Figure 4 shows a lock table entry for a group with less than seven members. The fact that there is an explicit identifier in the exclusive-interest byte shows that there is at least one member with X-mode interest.

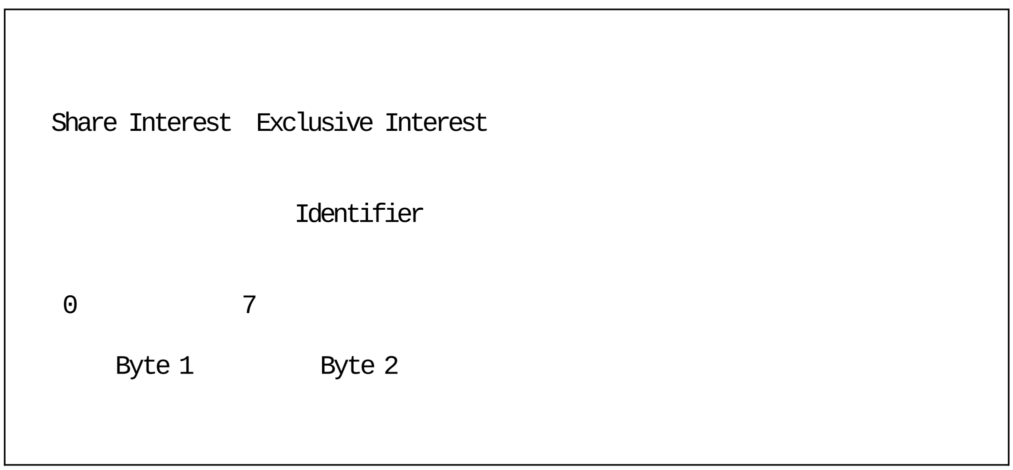

Figure 4. Two-Byte Lock Table Entry with X-Mode Interest. The 2-byte width supports six members.

The example shown on Figure 5 on page 25 shows S-mode interest for members 2, 7, 9, 16 on a group with fewer than 23 members.

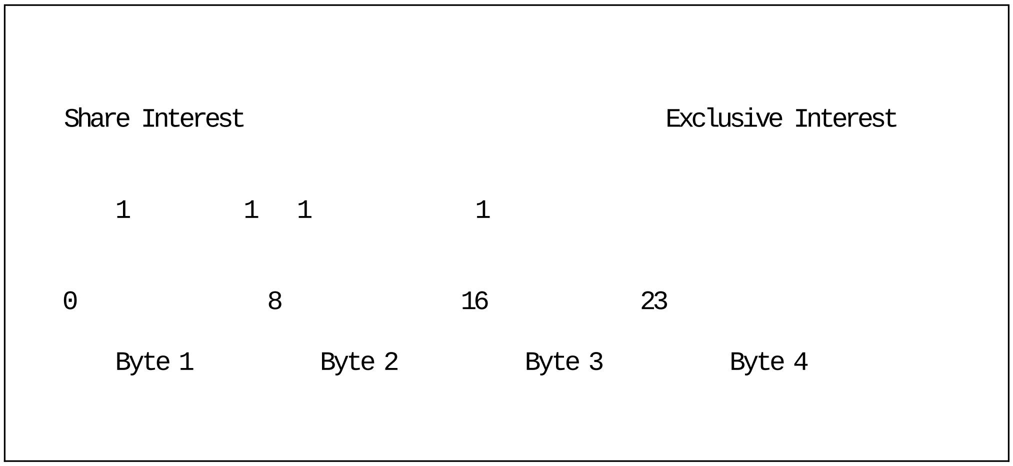

Figure 5. Four-Byte Lock Table Entry with S-Mode Interest. The 4-byte width supports 22 members.

4.2.2 Modified Resource List

The modified resource list is the other part (approximately half) of the lock structure. The list, which is used for recovery purposes, contains modified resource-list entries.

The modified resource list is where an IRLM associated with a DB2 member of a data sharing group can save run-time information about the modify locks and retained locks owned by the DB2 member.

A modify lock is an active X-type (X,IX,SIX) physical lock or logical lock owned by a DB2 member or a transaction. A retained lock is a modify lock that persists across the restart of a DB2 member.

Of the modified resource list in the lock structure, 10% is reserved for "must complete" transactions that are performing commit phase 2 or rollback processing. As the modified resource list fills up and the 90% regular space becomes exhausted, transactions requesting modify locks can be aborted.

4.2.3 Lock Structure Size

The lock structure resides in a coupling facility. The user controls the size of the complete lock structure. This size is defined in the MVS coupling facility resource management (CFRM) policy for the lock structure. The ratio of storage dedicated to each component of the lock structure (lock table and modified resource list) is currently one to one. This ratio is fixed and is determined by the IRLM.

At allocation time, and if the amount of storage defined in the CFRM policy is an exact power of 2, then both components of the structure have the same size. If the amount of storage defined is not an exact power of 2, the lock table has its storage allocation rounded down to an exact power of 2, and the rest is used for the modified resource list.

The size of the lock table is critical; it must be an exact power of 2 bytes. Because there always are fewer lock table entries than possible hash values (2 to the power of 32) and because no hash algorithm is perfect, different resources may hash to and share one lock table entry. The hashing of multiple resource names to a single lock table entry introduces the possibility of false contention. See also section 4.3, “Global Lock Manager” on page 26 and “False Contentions” on page 27.

There is a trade-off between the locking structure size and minimizing false contention. The frequency of false contention is inversely proportional to the size of the lock table (lock structure) size. If the lock table of the locking structure is too small, multiple resource names may be represented by a single lock table entry. On the other hand, if the lock table of the locking structure is larger, false contention can be reduced and global locking can be more efficient. However, overallocation of the locking structure may not be cost justified.

Hence a balance must be achieved between allocating an appropriate lock structure (lock table) size and achieving a reasonable degree of false contention. This balance is highly workload dependent. Each workload is a special case and results may vary from customer to customer. For example, a larger lock structure is required with increasing multiprogramming levels, longer-duration units of work, and heavy update intensity.

4.3 Global Lock Manager

The process of managing and handling all lock requests that hash to one lock table entry is called “global lock management.” The member who receives this function becomes the global lock manager for the lock table entry.

When there is contention on a lock with X-mode interest in the lock table, the XES component of the member holding the exclusive lock becomes the global lock manager for that entry in the lock table. The global lock manager is sometimes also called the owner of the lock table entry or the owner of the hash class. The identification of the global lock manager is inserted into the lock table entry.

A global lock manager is not created merely because different resources are hashed to the same lock table entry. It is created only when the requesting interest is in contention with the recorded interest.

The global lock manager knows the names of all the resources hashing to the lock table entry. This enables the global lock manager to resolve or confirm contentions on behalf of other XES components hashing to the lock table entry. The global lock manager must handle all subsequent requests from other XESs until the release of the lock table entry. When there is no more exclusive interest for the lock table entry, the global lock manager releases it.

4.4 Contentions

Three types of contention are defined in this section. These are listed in order of increasing time needed for their resolution:

- False contention
- XES contention
- Real contention.

Contention occurs when different members need to use the same resource at the same time. Real contention occurs when the lock requests are incompatible. When different resources with incompatible locks hash to the same entry in the lock table, a false contention occurs; It is the easiest to resolve. The third type, XES contentions, are contentions that can only be resolved by an IRLM. This is further explained below.

False Contentions: False contentions occur when the hashing algorithm for the lock table provides the same hash value for two different resources. The different resources then share that one lock table entry.

False contention can occur only when the lock table entry is managed by a global lock manager or when the lock request causes global management to be initiated.

The XES component requesting the lock needs to know the owning resource name to resolve this apparent contention. That information resides in the XES component that is global lock manager for the lock table entry. If the global lock manager is not the requesting XES, communication between XES components is needed to resolve the false contentions.

XES Contentions: XES contentions are the third type of contention. The MVS XES component is aware of only two lock modes, share and exclusive. IRLM locking supports many additional lock modes. When the MVS XES component detects a contention because of incompatible lock modes for the same resource, that contention is not necessarily a real contention by IRLM standards. For example, the IRLM finds the IX-mode to be compatible with the IS-mode. For the MVS XES component, however, these are not IX-mode and IS-mode, but X-mode and S-mode which are incompatible. To see if a real contention exists, MVS XES must give control to the IRLM contention exit associated with the global lock manager. The IRLM contention exit must determine if the contention is real or not, that is, if the locks are incompatible. If the contention is not real, it is called "XES contention" and the requested lock can be granted.

Real Contentions: Real contentions are caused by normal IRLM lock incompatibility between two members. For example, two transactions may try to update the same resource at the same time.

DB2 PM reports real contentions as IRLM contentions.

Resolving Contentions: Contentions require additional XES and XCF services if the requesting member is not global lock manager, that is, the owner of the lock registered in the lock table entry. See section 4.3, "Global Lock Manager" on page 26. More details of this process are described in section 4.7, "Global Locking Flows" on page 33.

Suspensions: When a contention occurs, execution of the requester′s SQL statement is suspended until the contention is resolved. If the contention is real, the requester remains suspended until the incompatible lock is released.

You can see the relative importance of the different types of contention in the DB2 PM Statistics report. For example, Figure 23 on page 63 shows number of suspensions due to each of the three types of contention.

4.5 Lock Duration

You use the bind option RELEASE to specify the lock duration. You can specify either RELEASE(COMMIT) or RELEASE(DEALLOCATE). RELEASE(DEALLOCATE) is the recommended lock duration option for plan or package BIND. The rest of this section explains the reason for this.

DB2 considers the table space logical locks as long-duration locks. Long-duration locks are isolated in a subset of the lock table. Because this subset is a fraction of the lock table, false contentions are more likely to happen for long-duration locks.

RELEASE(COMMIT) and a low multiprogramming level for a particular page set or partition on a specific member, may result in the continual release and acquisition global table space logical locks. This increased rate of global lock acquisition may cause an increase in contentions (XES and false) and may even change GBP-dependence.

RELEASE(DEALLOCATE) in thread reuse situation (protected threads or WFI), may reduce most of this impact. Protected threads provide thread reuse for customer information control system (CICS) applications. IMS Fast Path (IFP) and Wait For Input (WFI) regions provide thread reuse for IMS applications.

Examples are illustrated in Figure 6 on page 29 and Figure 7 on page 30. Figure 6 on page 29 shows the message traffic from Member 2 when using plans or packages bound with RELEASE(COMMIT). Figure 7 on page 30 shows the same for RELEASE(DEALLOCATE) with thread reuse. Message traffic has been reduced, because in this example the table space logical locks are not released at commit time.

For simplicity, several assumptions have been introduced in these figures:

- Member 1 activity is not shown.
- Member 1 is global lock manager for the table space logical lock.
- Only one shared table space is used by both members.
- The table space is updated by all transactions.

See notes to Figure 6 on page 29 and Figure 7 on page 30:

1 Member 2 wants to acquire an IX table space logical lock on the table space. The coupling facility denies this request (XES lock contention) and informs Member 2 that Member 1 is the global lock manager for the requested lock table entry. Member 2 now requests this lock from Member 1. The request is compatible and accepted.

2 Member 2 releases the IX logical lock on the table space, because the plan was bound with RELEASE(COMMIT). Member 2 informs Member 1 that the lock is released.

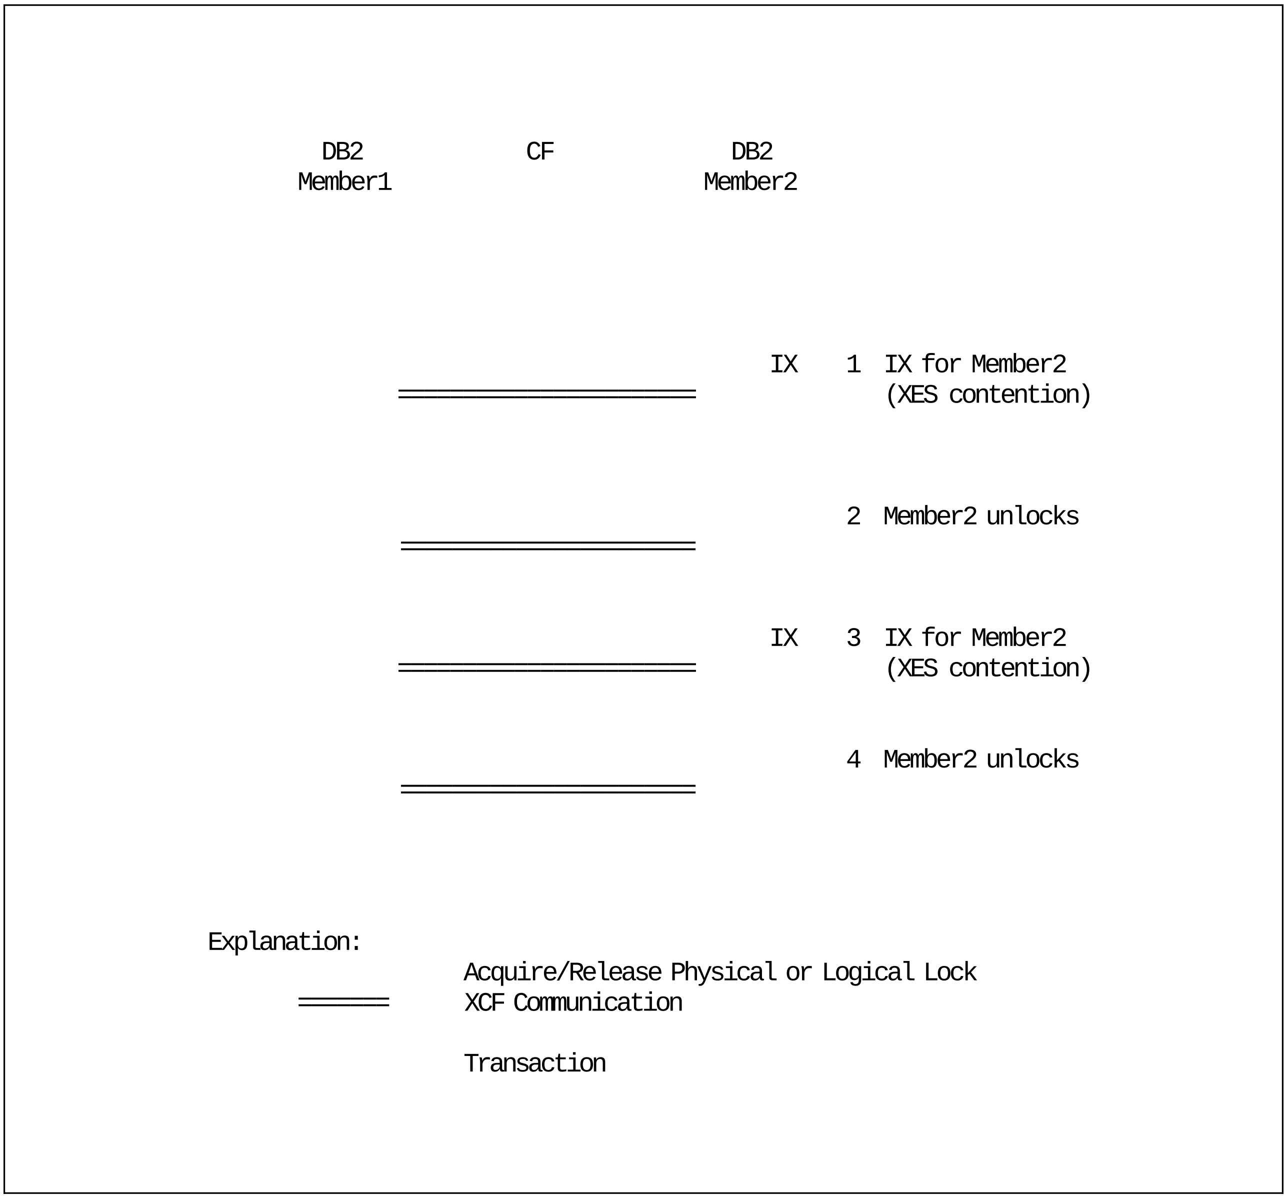

Figure 6. Global Parent Logical Locks with RELEASE(COMMIT). Member 2 interaction is shown.

3 Same as 1

4 Same as 2 .

4.6 Lock Avoidance

DB2 can avoid the use of logical locks when both of the following are true:

1. The plan or package is bound with ISOLATION(CS).
2. DB2 reads data that is guaranteed to be committed.

Lock avoidance without data sharing depends on the lowest uncommitted log record sequence number of the single DB2 for a page set or partition. In the data sharing environment, each DB2 member still maintains the CLSN for each page set or partition. However, these CLSNs are only used for non-GBP-dependent page sets or partitions. For GBP-dependent page sets or partitions, each member uses a global CLSN (GCLSN) value. Thus, to allow lock

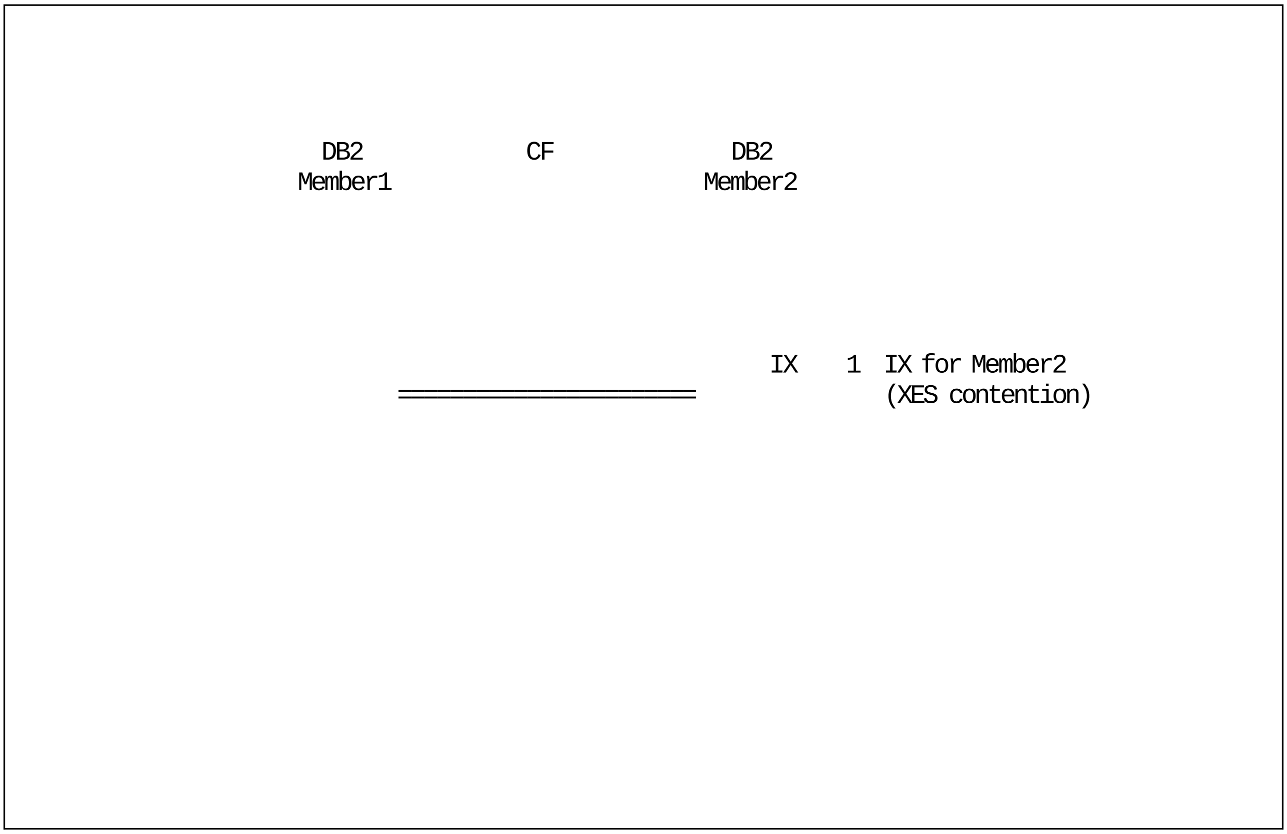

Figure 7. Global Parent Logical Locks with RELEASE(DEALLOCATE). Member 2 interaction using reusable threads is shown.

avoidance to be effective, all applications in the data sharing group are required to commit frequently.

Using lock avoidance, DB2 reads a page without taking an IRLM lock, which means that the page may contain uncommitted data. To prevent access to uncommitted data, DB2 uses a combination of techniques:

- Page latching
- Page physical locks
- The use of CLSN, in conjunction with the page version number (relative byte address, RBA or log record sequence number, LRSN in the page header of the last update to the page)
- Possible uncommit (PUNC) bit.

4.6.1 Lock Avoidance for a Data Sharing Group

Lock avoidance in a data sharing environment uses the same techniques as those in previous DB2 versions. The only new concept added for this environment is the GCLSN.

For non-GBP-dependent page sets or partitions, each member uses a local CLSN, which is the lowest uncommitted log sequence number for all active transactions.

For GBP-dependent page sets, a single GCLSN value is maintained and used for lock avoidance for the entire data sharing group. The single GCLSN value is derived from each member's CLSN values. The member's CLSN value is

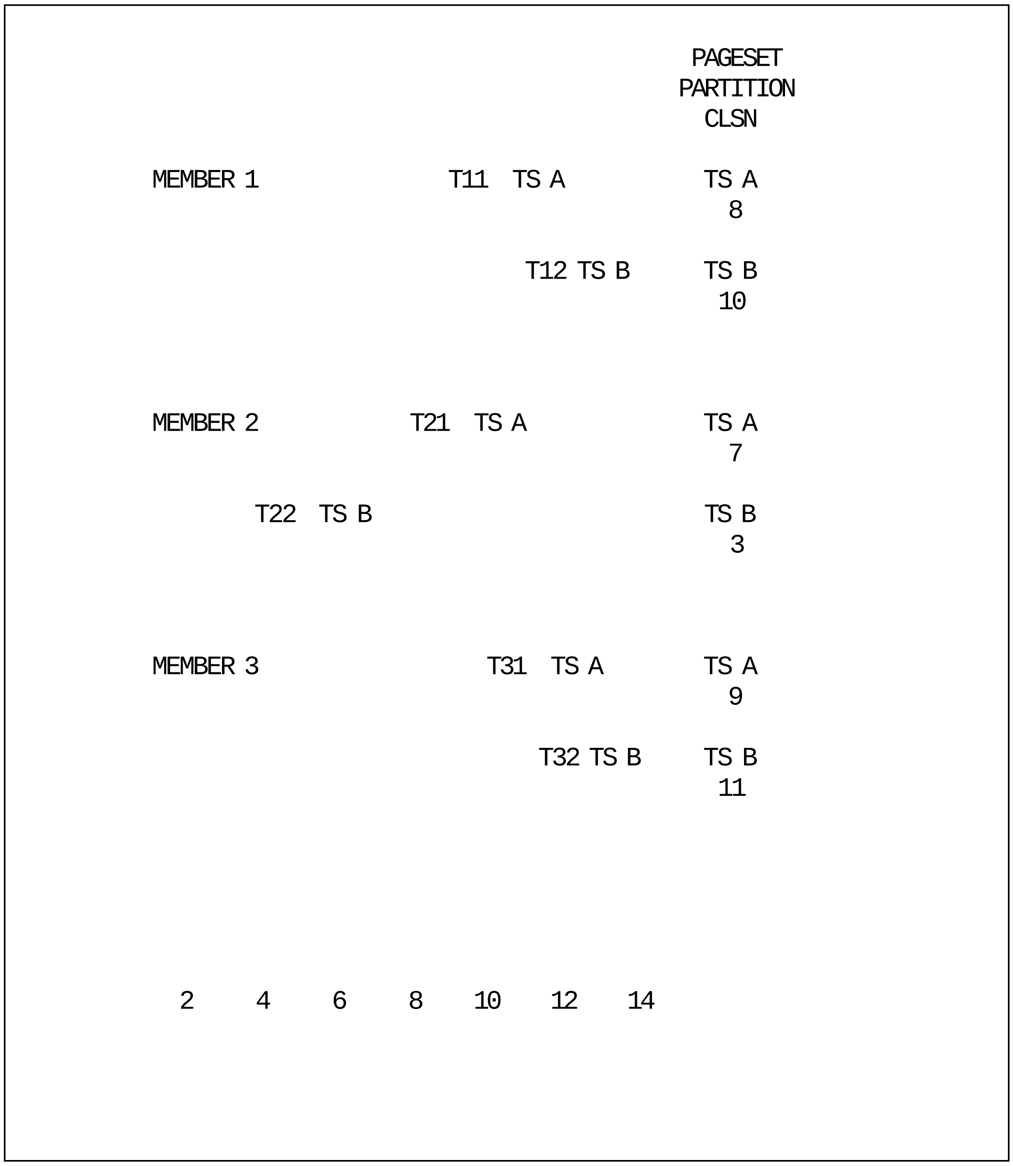

Figure 8. Local CLSN for Three DB2 Members

maintained at the DB2 subsystem level, not at the page set or partition level. It is the beginning of unit of recovery (UR) LRSN value of the oldest URs in this DB2 member.

The GCLSN is the earliest CLSN value across all members across all page sets or partitions, regardless of whether the page sets or partitions are GBP-dependent or not. Each member independently updates the GCLSN periodically, based on timer interval.

Periodically, each member pushes the latest CLSN value into the SCA and pulls from the SCA the lowest value across all members.

Figure 8 summarizes the situation for three members, all three of them accessing table spaces A and B. Member 1 has two transactions in flight, T11 and T12. T11 updates table space A at LRSN 8; T12 updates table space B at LRSN 10. Member 1 keeps the CLSN for both table spaces, in this case, table space A has CLSN=8, and table space B has CLSN=10.

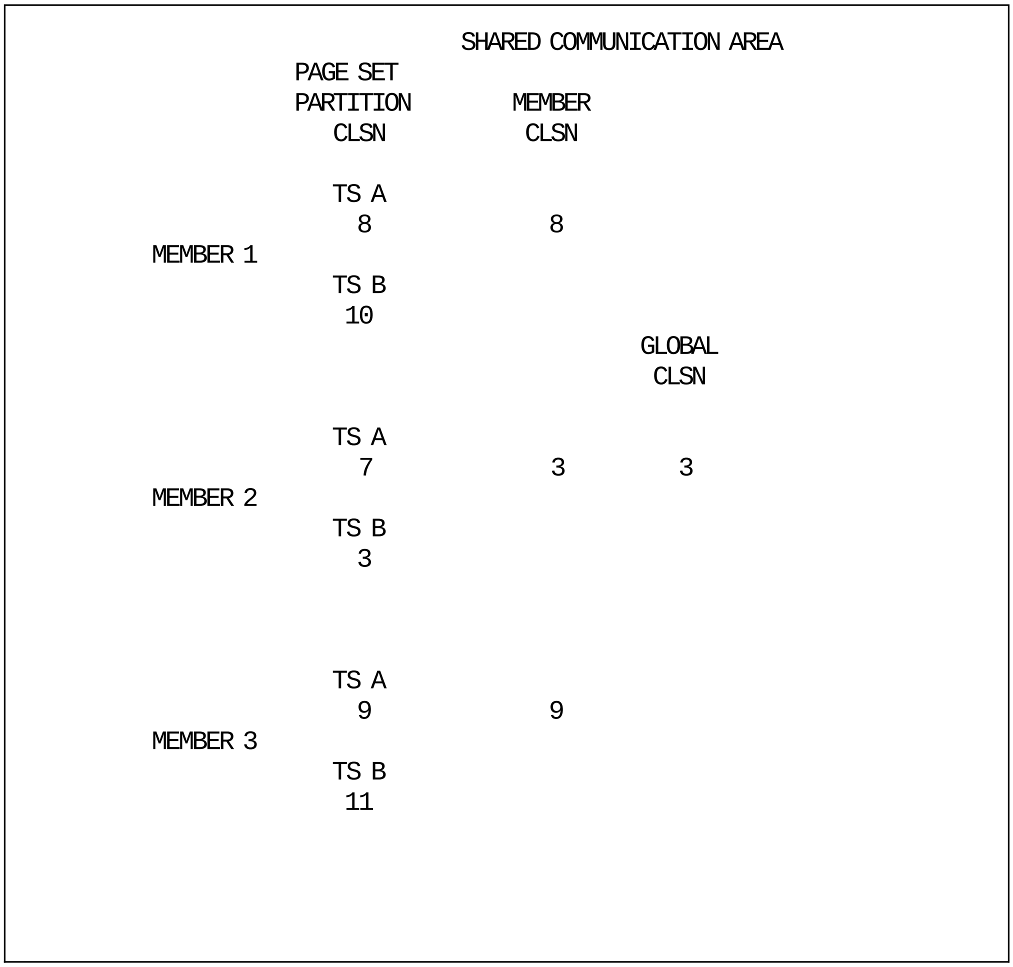

Figure 9. Global CLSN

In a data sharing environment, Member 1 cannot rely on these numbers, because other members may have transactions with a lower CLSN for these table spaces. In Figure 8, this is the case with transactions T21 and T22 for table spaces A and B, respectively.

DB2 data sharing needs to interchange the GCLSN between members in order to apply lock avoidance. Periodically, each member obtains the GCLSN from the shared communication area.

Figure 9 shows how the GCLSN is obtained and also illustrates the impact one long-running transaction can have on lock avoidance. In Figure 8 on page 31, T22 is shown running for a long time without commit. On Figure 9, the CLSN of this transaction becomes the global CLSN because it is the lowest value for all members. Because this value is low, the CLSN technique qualifies fewer pages for lock avoidance.

One long-running update process without intermediate commit points can effectively stop the GCLSN value from moving forward regardless of whether the process is updating GBP-dependent or non-GBP-dependent page sets or partitions.

The following actions are recommended to maximize lock avoidance:

- Frequent commits of all processes in all members

- Quiescing of inactive DB2 members
- Quick restart of failed members to resolve indoubt transactions.

4.6.2 Recommendations for Lock Avoidance

The following SQL and usage rules are recommended for lock avoidance:

1. Do not use ISOLATION(RR). There is no lock avoidance for an application process using RR, and the presence of RR claimers causes SQL inserters to perform next locking for transactions using CS.
2. Read-only cursors bound with ISOLATION(CS) and CURRENTDATA(NO) use lock avoidance for both qualifying and nonqualifying pages or rows.
3. Ambiguous cursors bound with ISOLATION(CS) and CURRENTDATA(NO) use lock avoidance for both qualifying and nonqualifying pages or rows.
4. Any cursor bound with ISOLATION(CS) uses lock avoidance for nonqualifying pages or rows.
5. A cursor bound with ISOLATION(UR) does not lock pages or rows.

4.7 Global Locking Flows

Figure 10 on page 34 shows the elements supporting global locking in a sysplex environment and is used to describe the global locking flow for several situations. The numbers included in Figure 10 on page 34 refer to the descriptions that follow.

If there is global lock contention between DB2 members of a DB2 data sharing group, XCF messages are used to resolve the contention. Message frequency depends upon the number of global locks acquired, the degree of lock contention, and the number of messages needed to resolve each lock contention.

The message-flow paths that can be taken to resolve a global lock request are described in the order of increasing cost. Several X-mode child logical lock requests are used as illustrative cases:

- No inter-DB2 read-write interest
- Inter-DB2 read-write interest without contention
- Inter-DB2 read-write interest with false contention
- Inter-DB2 read-write interest with XES contention
- Inter-DB2 read-write interest with real contention.

No Inter-DB2 Read-Write Interest: If there is no inter-DB2 read-write interest at the parent level, the IRLM associated with the DB2 member grants the child lock request locally, without interacting with the MVS XES component supporting the DB2 member. Therefore, when only one DB2 member has write interest or when multiple DB2 members have read-only interest, there is no additional overhead for locking at the child level in DB2 data sharing. The flow is as follows:

1. Application 1 issues an SQL UPDATE.
2. DB2 2 decides a lock is required and requests it from the IRLM.
3. IRLM 3 grants the lock locally and replies to DB2.

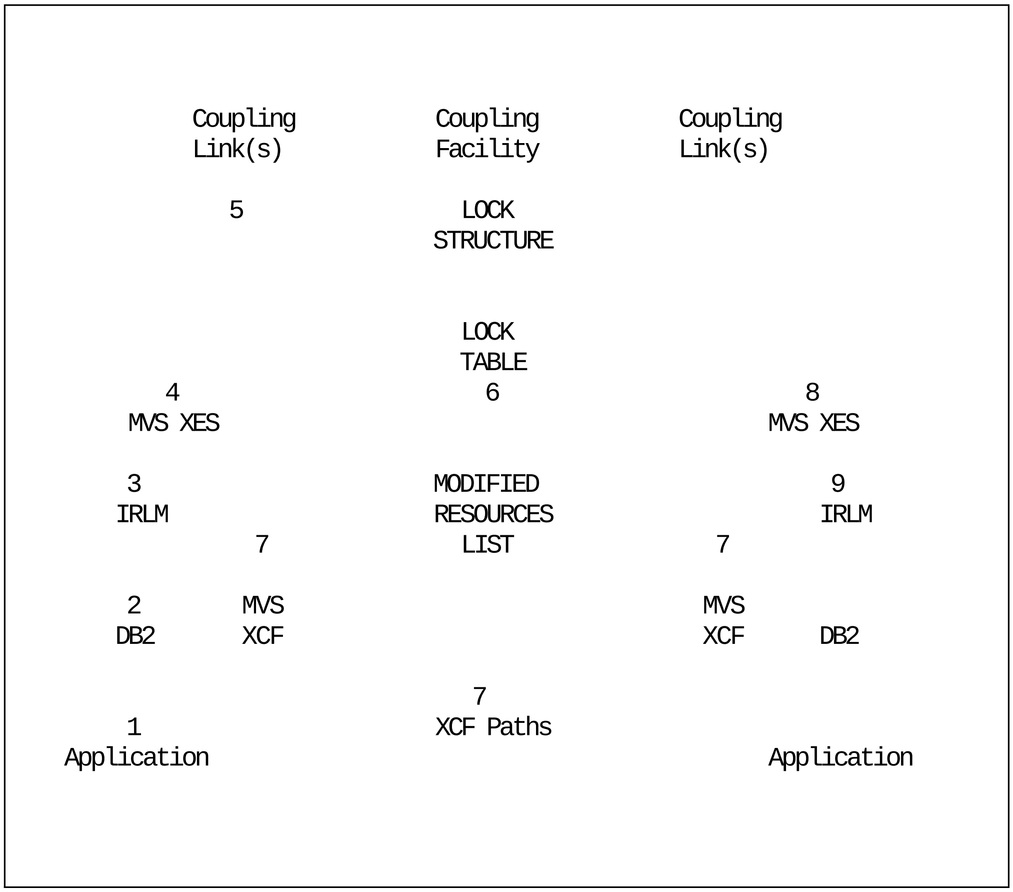

Figure 10. Locking Flows in a Sysplex Environment

4. DB2 2 executes the SQL statement and passes the result to application 1 .

Inter-DB2 Read-Write Interest Without Contention: If there is inter-DB2 read-write interest at the parent level, the IRLM associated with the DB2 member may not be able to grant the lock at the child level locally, by itself. With explicit hierarchical locking and propagation of only the most restrictive lock per member, the local IRLM may be able to handle the request without ever going to XES. Otherwise, the IRLM has to pass the lock request to the XES component supporting the DB2 member. XES interacts with the coupling facility structure. If the coupling facility structure indicates that there is no inter-system conflict on the lock table entry to which the lock request hashes, the lock is granted synchronously. This should be the most common case since the probability of contention at the child level is expected to be small.

The flow of a global logical lock with inter-DB2 read-write interest and without contention is as follows:

1. Application 1 issues a SQL UPDATE.
2. DB2 2 decides a lock is required and requests it from the IRLM.
3. IRLM 3 registers the lock and passes it to XES 4 .
4. XES 5 sends the lock to the coupling facility.
5. The coupling facility registers the lock in the lock table 6 , and no contention is found. Because this is an X lock, a modified resource list entry is made.

6. The coupling facility 5 replies to XES: lock granted.
7. XES 4 replies to IRLM.
8. IRLM 3 replies to DB2.
9. DB2 2 executes the SQL statement and passes the result to application 1 .

Inter-DB2 Read-Write Interest with False Contention: The MVS XES requester tries to register a lock in the coupling facility structure and detects a contention on the lock table entry. To resolve the contention the requester needs to interact with the global lock manager for that lock table entry. If global management has not yet been set up for this lock table entry, it is done at this time. Also, if XES knows that there is global management on the lock table entry, the request never goes to the coupling facility in the first place. XES would know that there is global management without having to go to the coupling facility if XES is itself the global lock manager, or if XES was notified when global management was set up. If XES has share interest in the lock table entry, notification should be automatic.

Until this is resolved, the DB2 transaction must be suspended. The MVS XES component of the requesting member interacts with the MVS XES component of the member that is the global lock manager of the lock table entry. The global lock manager for this example is assumed to be the other member shown on Figure 10 on page 34. Because of this, the communication uses XCF messages.

The MVS XES component that is the global lock manager determines that this is a false contention (different resource names happen to hash to the same lock table entry). The child lock request can be granted and the transaction resumed. The flow is as follows:

1. Application 1 issues an SQL UPDATE.
2. DB2 2 decides a lock is required and requests it from the IRLM.
3. IRLM 3 registers the lock and passes it to XES 4 .
4. XES 5 sends the lock to the coupling facility.
5. The coupling facility finds a contention 6 .
6. The coupling facility replies to XES 5 with a notice of contention and a global lock manager identifier.
7. XES 4 asks IRLM 3 to suspend the SQL statement 2 .
8. XES 4 use XCF 7 to query the XES global lock manager 8 .
9. The XES 8 global lock manager finds that the lock is a false contention; no real contention exists.
10. XES 8 replies to XES 4 through XCF 7 saying that no contention exists.
11. XES 4 replies to IRLM: no contention.
12. IRLM 3 replies to DB2.
13. DB2 2 resumes the SQL statement and passes the result to application 1 .

Inter-DB2 Read-Write Interest with XES Contention: The XES component knows about S-mode and X-mode locks only, while IRLM supports more lock states. The XES component of the requesting DB2 member detects a lock table entry contention on the requested lock because the same resource name is used. The MVS XES component which is the global lock manager determines from an MVS XES component perspective that the contention exists (same resource name, but incompatible lock states). Before this contention can be considered real from an IRLM perspective, an IRLM has to examine the locks. The XES global lock manager passes control to the IRLM contention exit associated with its own IRLM. The IRLM contention exit determines if the contention is true. If it is not true, (for example, if the involved request is IX to the IRLM, but maps to the MVS XES component as X), the lock request can be granted. This is called an XES contention. The DB2 transaction, which was suspended during contention resolution, can now be resumed.

The flow of a global logical lock with inter-DB2 read-write interest and with a XES contention is as follows:

1. Application 1 issues an SQL UPDATE.
2. DB2 2 decides a lock is required and requests it from the IRLM.
3. IRLM 3 registers the lock and passes it to XES 4 .
4. XES 5 sends the lock to the coupling facility.
5. The coupling facility finds a contention 6 .
6. The coupling facility replies to XES 5 with a finding of contention and global lock manager identifier.
7. XES 4 asks IRLM 3 to suspend SQL statement 2 .
8. XES 4 use XCF 7 to query the XES global lock manager 8 .
9. XES 8 global lock manager finds the lock is in contention.
10. XES 8 drives the IRLM 9 to a contention exit identifying XES contention.
11. XES 8 replies to XES 4 through XCF 7 that no real contention exists.
12. XES 4 replies to IRLM: no contention.
13. IRLM 3 replies to DB2.
14. DB2 2 resumes the SQL statement and passes the result to application 1 .

Inter-DB2 Read-Write Interest with Real Contention: This case is similar to "Inter-DB2 Read-Write Interest with XES Contention," except that the IRLM lock contention exit determines real lock contention. The child lock request cannot be granted. The DB2 transaction remains suspended until the lock state becomes compatible because of a subsequent unlock request.

The flow of a global logical lock with inter-DB2 read-write interest and with real contention is:

1. Application 1 issues an SQL UPDATE.
2. DB2 2 decides a lock is required and requests it from the IRLM.
3. IRLM 3 registers the lock and passes it to XES 4 .
4. XES 5 sends the lock to the coupling facility.
5. The coupling facility finds a contention 6 .

6. The coupling facility replies to XES 5 with a finding of contention and global lock manager identifier.
7. XES 4 asks IRLM 3 to suspend the SQL statement 2 .
8. XES 4 use XCF 7 to query the XES global lock manager 8 .
9. XES 8 global lock manager finds lock is a contention.
10. XES 8 drives IRLM 9 contention exit, identifying real contention.
11. XES 8 replies to XES 4 through XCF 7 : real contention.
12. XES 8 replies to IRLM 3 : contention.
13. The SQL statement execution 2 remains suspended until the resource is unlocked.
14. DB2 2 resumes the SQL statement and passes the result to application 1 .

Inter-DB2 Read-Write Interest with Physical Lock Negotiation: This case is a variation of the case above and is applicable only for physical locks (see 4.1.2, "Page set Physical Locks" on page 21) which are negotiable. When the IRLM associated with the MVS XES component that is the global lock manager of the lock table entry discovers contention on a physical lock (same resource name and incompatible lock state), the IRLM invokes physical lock exit code in each of the DB2 members that are holding incompatible lock states so that they can downgrade their lock state to make it compatible with the requested lock state.

Unlock Requests: If a parent object requires global lock management, the IRLM associated with a DB2 member cannot locally process an unlock request for a child lock. The unlock request must pass to the MVS XES component supporting the DB2 member. This communication is performed synchronously and control is returned to the requesting transaction.

Further processing of this unlock request is performed asynchronously:

- Interaction by the MVS XES component supporting the requesting DB2 member with the coupling facility structure.
- Possible interaction by the MVS XES component supporting the requesting DB2 member that is the global lock manager of the lock table entry through an XCF message.

Chapter 5. Measurements Analysis

This chapter describes the results of the measurements of the IBM Relational Warehouse Workload by the DB2 Performance Group at the IBM Santa Teresa Laboratory. The information presented may help you identify capacity needs for DB2 data sharing.

This chapter has four main sections:

- Measurement Summary: contains tables with data on throughput, capacity delta, and coupling facility activity.
- Data Sharing Environments: describes some migration paths and the impact on throughput.
- Performance Tests: explains the test environment, the test configuration, and the workload used.
- Performance Calculations: contains the formulas that were used and the calculations for 1-way and 2-way measurement using type 2 index and page locking.

5.1 Measurement Summary

All the measurements use plans bound with ISOLATION(CS) and CURRENTDATA(NO).

This section presents the results in tables containing the captured and the calculated performance indicators for the seven measurement runs.

Measurement Runs: We use the measurements to evaluate the performance of DB2 data sharing with the environments indicated in Table 2.

Table 2. Performance Tests at the IBM Santa Teresa Laboratory

DB2 Data Sharing Environment	**Type 2 Index**	**Type 1 Index**	**Type 2 Index, Row Locking**
	Page Locking		
1-way measurement: One DB2 member	1WP	1W1	1WR
2-way measurement: Two DB2 members	2WP	2W1	2WR
3-way measurement: Three DB2 members	3WP	–	–

We use the symbols 1WP, 2WP, ... to identify each type of measurement. The results are used to evaluate the capacity delta to process DB2 data sharing, and the improvements offered by the type 2 index when compared with the type 1 index in a data sharing environment.

Capacity Delta This is the additional capacity required for each member in a n-way data sharing environment to provide throughput equivalent to that of an environment without data sharing and without S/390 Parallel Sysplex. All measurements are made in an MVS parallel sysplex environment. Therefore, approximately 3% should be added to the calculated capacity delta to reflect the fixed sysplex multisystem MVS overhead. This fixed overhead comes from global resource serialization (GRS) processing, XCF signaling and heartbeat monitoring, job entry system (JES) checkpoints, and other miscellaneous items.

We also evaluate the impact of using row locking with data sharing. The 1WR and 2WR measurements use row locking exclusively on all tables to evaluate the effect of row locking in a data sharing world.

The sources of data for these evaluations are the reports from Resource Measurement Facility (RMF) Monitor I and the Statistics reports from DB2 Performance Monitor (DB2 PM).

5.1.1 Throughput Measurements

Table 3 on page 41 gives the results obtained and calculated values of number of transactions in different runs, during a period of 10 minutes.

The best performance for data sharing in this particular environment is achieved by using type 2 index and page locking. The type 1 index has a significant impact on throughput when migrating from 1-way to 2-way (ITR values for 1W1 to 2W1 in the table). Row locking has a higher impact than page locking when sharing data.

Run Id	CPC	CPU %Busy	ETR (Trans.)	ITR (Trans.)	CPU Seconds Per Transaction
Type 2 Index and Page Locking					
1WP	CPC 1	72.53	21493	29633	0.075
2WP	CPC 1	70.08	18056	25765	0.083
	CPC 2	70.99	18191	25625	0.085
3WP	CPC 1	68.20	17669	25908	0.083
	CPC 2	69.90	17751	25395	0.086
	CPC 3	69.45	17747	25553	0.085
Type 1 Index and Page Locking					
1W1	CPC 1	73.14	21609	29545	0.076
2W1	CPC 1	76.94	16860	21913	0.106
	CPC 2	77.00	16910	21961	0.106
Type 2 Index and Row Locking					
1WR	CPC 1	73.86	21504	29115	0.077
2WR	CPC 1	78.85	17937	22748	0.096
	CPC 2	74.22	17878	24088	0.090

Table 3. Throughput Results

Note:

ETR = External Throughput Rate; the number of transactions executed in 10 minutes.

ITR = Internal Throughput Rate; the calculated number of transactions for CPU 100% busy.

In Table 3:

- CPU %Busy is obtained from the RMF CPU Activity Report
- ETR is the TOTAL COMMITS in the DB2 PM Statistics report
- ITR is calculated by dividing ETR by CPU %Busy
- CPU Seconds Per Transaction is the Class 1 CPU TIME summarized in the GRAND TOTAL in the DB2 PM Accounting Report.

It is interesting to note that, in the measurements, the average CPU time for a transaction is between 75 and 106 milliseconds. The highest value is for 2W1 measurement corresponding to Type 1 index with data sharing.

5.1.2 Capacity Delta Measurements

This section discusses capacity delta using

- Type 2 index and page locking.
- Type 1 index and page locking.
- Type 2 index and row locking.

Capacity Delta for Type 2 Index and Page Locking: Table 4 shows the calculated performance indicators for different runs.

Capacity Delta Increase in capacity requirement compared with 1-way measurement.

Coupling Efficiency (%) 100% - Capacity Delta

For more information about projections and capacity delta using DB2 data sharing for multiple members, see section 6.3.1, "Multisystem Management Cost" on page 71 and section 6.4, "Capacity Delta Projection" on page 73.

The capacity delta when migrating from 1-way data sharing to 2-way data sharing is 13.29% as an average for two members in 2-way measurement. When we move from 2-way data sharing to 3-way data sharing, there is an increase of 0.26% in the capacity delta, rising to 13.55%.

Table 4. ITR and Other Data for Type 2 Index and Page Locking

Performance Indicator	Type 2 Index and Page Locking		
	1WP	2WP	3WP
(1-way ITR) x N	29633	59266	88899
(ITR1 + ITR2 + ... + ITRN)	29633	51390	76856
Capacity Delta	–	13.29%	13.55%
Coupling Efficiency (%)	100%	86.71%	86.45%
Increased CPU Time	–	15.33%	15.67%

The Coupling Efficiency (%) indicator gives the percentage throughput for DB2 data sharing when compared with the throughput without data sharing in the equivalent configuration.

The Increased CPU time gives the percentage of additional CPU time necessary to process data sharing for this workload.

Capacity Delta for Type 1 Index and Page Locking: Table 5 shows the calculated performance indicators for different runs using type 1 index.

The capacity delta when migrating from 1-way data sharing to 2-way data sharing is 25.75%. This capacity delta is more compared to 13.29% when migrating from 1-way data sharing to 2-way data sharing using type 2 index.

Table 5. ITR and Other Data for Type 1 Index and Page Locking

Performance Indicator	Type 1 Index and Page Locking		2WP
	1W1	2W1	
(1-way ITR) x N	29545	59090	59266
(ITR1 + ITR2 + ... + ITRN)	29545	43874	51390
Capacity Delta	–	25.75%	13.29%
Coupling Efficiency (%)	100%	74.25%	86.71%
Increased CPU Time	–	34.68%	15.33%

Capacity Delta for Type 2 Index and Row Locking: Table 6 shows the calculated performance indicators for different runs using row locking.

The capacity delta when migrating from 1-way data sharing to 2-way data sharing is 19.57%. This increase in capacity delta for 2-way data sharing using row locking is greater than the result of 13.29% for page locking in 2-way data sharing

Table 6. ITR and Other Data for Type 2 Index and Row Locking

Performance Indicator	Type 2 Index and Row Locking		2WP
	1WR	2WR	
(1-way ITR) x N	29115	58230	59266
(ITR1 + ITR2 + ... + ITRN)	29115	46836	51390
Capacity Delta	-	19.57%	13.29%
Coupling Efficiency (%)	100%	80.43%	86.71%
Increased CPU Time	-	24.33%	15.33%

Note:

The measurements 1WR and 2WR use row locking exclusively.

5.1.3 Coupling Facility Activity

The accesses to the coupling facility by DB2 are mainly because of XES requests for locking and use of the GBP.

Coupling Facility Activity for Type 2 Index and Page Locking: Table 7 on page 44 shows the calculated accesses to the coupling facility in the different runs using type 2 index.

The total number of accesses to the coupling facility at 3 accesses per second for 1-way execution is negligible when compared with 2188 accesses per second for 2-way execution.

Table 7. Coupling Facility Activity for Type 2 Index and Page Locking			
Coupling Facility Activity	**Type 2 Index and Page Locking**		
	1WP	**2WP**	**3WP**
Coupling Facility Accesses per Second			
XES Lock Requests per Second	3	1096	1695
GBP Accesses per Second	0	1092	1665
Total CF Accesses per Second	3	2188	3360
Coupling Facility Accesses per Commit			
(ETR1 + ETR2 + ... + ETRN)	21493	36247	53167
XES Lock Requests per Commit	0.08	18.14	19.13
GBP Accesses per Commit	0	18.09	18.79

For more information see section 8.1.3.1, "Evaluation of Locking Activity and Data Sharing Locking" on page 99.

Coupling Facility Activity for Type 1 Index and Page Locking: Comparison of Table 8 with Table 7 shows that:

- The number of coupling facility accesses is zero or negligible for 1-way measurement.
- The number of GBP requests per commit is very similar to that for 2-way measurement.
- For 2-way, however, the number of XES lock requests (62.30) using type 1 index is more than three times the number of XES lock requests (18.14) using type 2 index.

Table 8. Coupling Facility Activity for Type 1 Index and Page Locking		
Coupling Facility Activity	**Type 1 Index and Page Locking**	
	1W1	**2W1**
Coupling Facility Accesses per Second		
XES Lock Requests per Second	0	3506
GBP Accesses per Second	0	1004
Total CF Accesses per Second	0	4510
Coupling Facility Accesses per Commit		
(ETR1 + ETR2 + ... + ETRN)	21446	33770
XES Lock Requests per Commit	0	62.30
GBP Accesses per Commit	0	17.85

Coupling Facility Activity for Type 2 Index and Row Locking: Comparison of Table 9 on page 45 with Table 7 shows that:

- The number of coupling facility accesses is negligible for 1-way measurement.
- The number of GBP requests per commit is very similar to that for 2-way measurement.

- But for 2-way, the number of XES lock requests (40.04) using row locking is more than double the number of XES lock requests (18.14) using page locking.

Table 9. Coupling Facility Activity for Type 2 Index and Row Locking

Coupling Facility Activity	Type 2 Index and Row Locking	
	1WR	**2WR**
Coupling Facility Accesses per Second		
XES Lock Requests per Second	3	2389
GBP Accesses per Second	0	1053
Total CF Accesses per Second	3	3442
Coupling Facility Accesses per Commit		
(ETR1 + ETR2 + ... + ETRN)	21504	35815
XES Lock Requests per Commit	0.08	40.04
GBP Accesses per Commit	0	17.66

5.2 Data Sharing Environments.

This section makes ITR comparisons related to some of the possible changes in DB2 data sharing environments. These numbers can be used as a first reference, considering the specific conditions used in our measurements:

- IBM Relational Warehouse Workload
- Homogeneous S/390 microprocessor environment
- Degree of data sharing equal to 100%
- Random scheduling of transactions and keys across members of the DB2 data sharing group.

The measurements using row locking (1WR and 2WR) use this lock size exclusively. The probability of having this radical change for all tables in a DB2 production environment should be low.

Some Migration Paths in Data Sharing Environments Looking at nine migration paths in Figure 11 on page 46 we can see the throughput impact changing from one particular environment to other using IBM Relational Warehouse Workload (IRWW). We start with 1-way measurement using type 1 index and page locking.

5.2.1 Migrating within 1-way Environments

Path 1 1W1 to 1WP.

If we are using type 1 index and want to use type 2 index, the benefit is an increase of 0.3% in performance because of less locking activity. We strongly recommend using only the type 2 index in the data sharing environment. Therefore, consider altering indexes to type 2 before migrating to data sharing in DB2 Version 4.

Path 2 1W1 to 1WR

This is a possible case in a 1-way environment if we want to reduce the concurrency contention problems by taking advantage of row locking. But, before specifying row locking for a table space, we need

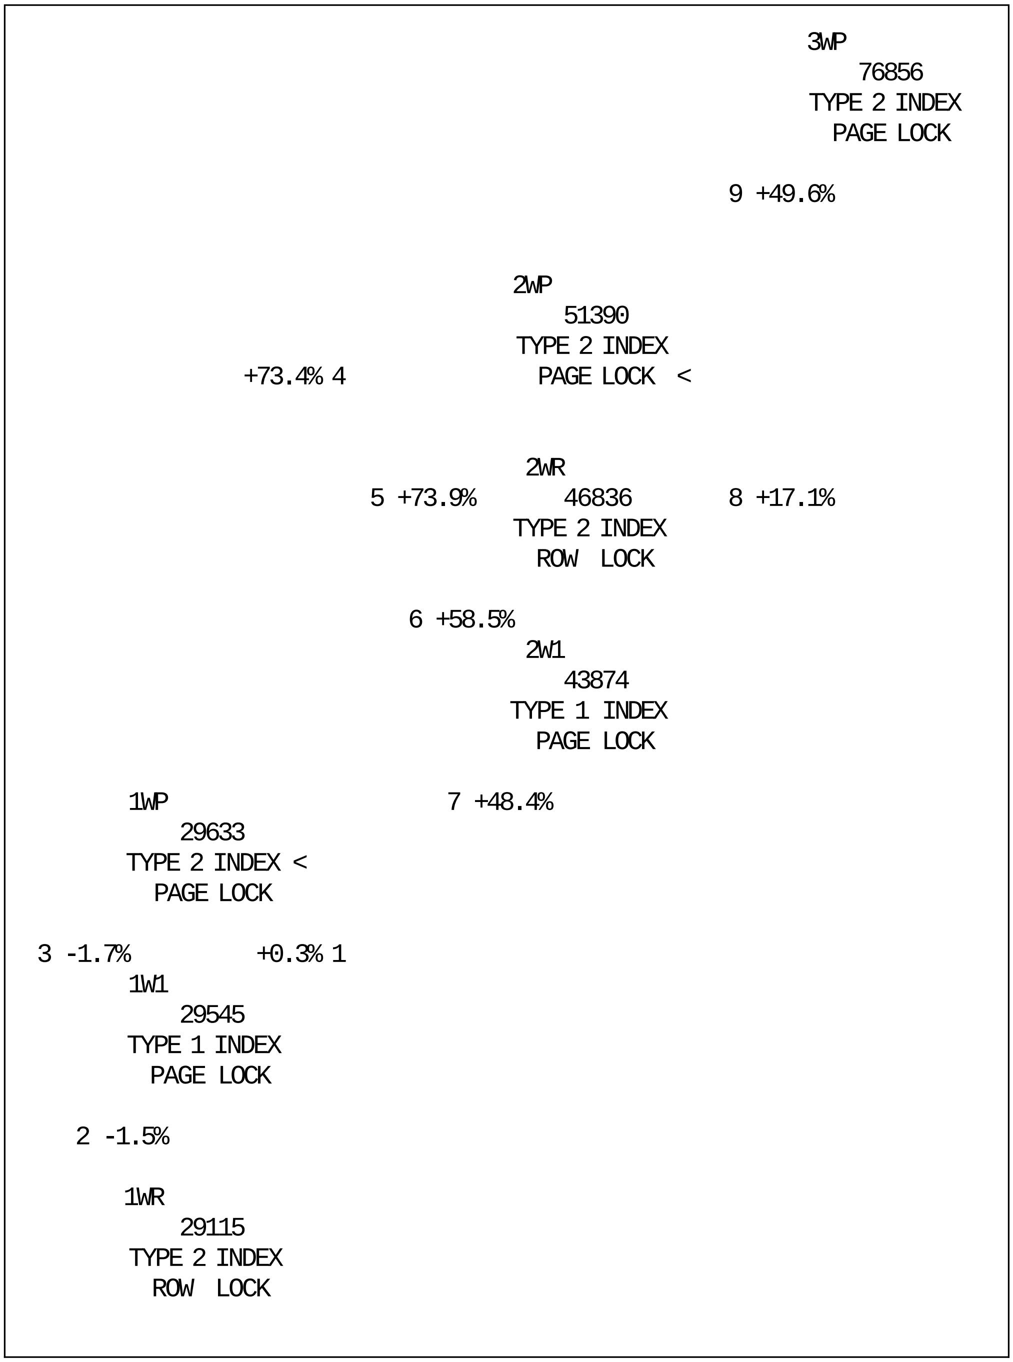

Figure 11. ITR Percentages for Different Migration Paths Using the IRWW.

to define all the indexes on the tables in the table spaces as type 2 indexes. The cost in throughput for this change is 1.5% less capacity to process transactions.

Good candidates for row locking are those tables in which we want to artificially decrease the number of rows per page to solve data concurrency problems. Other good candidates are those tables whose indexes have been defined with more subpages to decrease the number of entries per index subpage, thus avoiding index concurrency problems.

Path 3 1WP to 1WR

This case is very similar to the case of 1W1 to 1WR. The only difference is that we compare the throughput with the capacity after altering the indexes to type2. This step can be considered an intermediate step for path 2 ; the cost is also similar; 1.7% instead of 1.5%.

5.2.2 Migrating to 2-way Environments

Path 4 1WP to 2WP

This is the recommended path for most of the tables in one installation when changing from 1-way to 2-way processing to share data between two DB2 members. The additional throughput obtained with the second processor 9672-R61 is 73.4%. This percentage is less than 100% because additional CPU time is necessary to control the inter-DB2 read-write interest as discussed in section 3.2, "Group Buffer Pools" on page 14 and section 3.3, "Global Locking" on page 17.

Path 5 1W1 to 2WP

If we go from 1-way with type 1 index to 2-way with type 2 index, we have an increase in throughput of 73.9%.

Path 6 1W1 to 2WR

This case seems to us to be very hypothetical: It would start 1-way, converting the indexes to type2, and altering all table spaces to use LOCKSIZE ROW in 2-way processing. There is a 58.5% increase in throughput going from 1-way to 2-way.

Path 7 1W1 to 2W1

We do not recommend this path. The recommended option is to use the type 2 index and page locking for data sharing between two or more members because you then use data pages locks only, enable query CP parallelism and allow uncommitted read isolation (see section 7.6.2, "Type 2 Indexes" on page 87).

There is a 48.4% increase in throughput when going from 1-way to 2-way data sharing

Path 8 2W1 to 2WP

If we use path 7 , we can then convert index to type2 in order to obtain additional capacity.

The additional 17.1% of throughput is achieved because type 2 index requires fewer lock requests to the coupling facility in a data sharing environment. The difference in the number of XES locks for the scenarios before and after the migration can be found in "Coupling Facility Activity for Type 2 Index and Page Locking" on page 43 and "Coupling Facility Activity for Type 1 Index and Page Locking" on page 44.

Additional throughput for 2-way processing. is only 0.3% for the same change in the 1-way environment.

5.2.3 Migrating from 2-way to 3-way Environment

Path 9 2WP to 3WP

Once we are using data sharing with type 2 index and we need additional capacity, we plan to go from 2-way to 3-way configuration.

Surprisingly, we see a growth of almost 50% (49.6%) when compared with the throughput of 2-way data sharing This result is possible because the cost of data sharing is kept at about the same level in terms of global lock contentions, XES lock requests, and GBP accesses.

5.3 IBM Relational Warehouse Workload Performance Tests

We use the IRWW to study DB2 data sharing performance in support of DB2 design and development.

5.3.1 IRWW Test Environment

Figure 12 on page 49 shows the environment created to run the IRWW workload in 3-way data sharing. The figure shows the transactions arriving at three different IMS transaction manager (TM) systems each in separate 9672-R61 processor. The Teleprocessing Network Simulator (TPNS):

- Runs in a separate 9672 processor
- Generates transactions randomly based on a TPNS script
- Controls transaction arrival rate to achieve a given CPU %busy.

The transactions are scheduled to be executed in IMS Fast Path regions. One of the transactions, T5, starts one batch-type transaction, T6, which is executed in a wait-for-input (WFI) message processing region (MPR). Each IMS system has 66 regions to process transactions.

We use three sets of tables, one each for 1-way, 2-way, or 3-way measurements. These three sets or replicates of tables ease the load and refresh process of the databases. We use a fixed size for the total database.

TPNS script randomly generates transactions against each set of tables for each member of the data sharing group.

In this environment:

- Each IMS TM has 66 regions to process transactions.
- 22 regions access one of the three sets of tables.
- Each set of tables is shared for read and write between the three IMS TM systems.
- Just as the tables are replicated, so are the transactions.
- A single application program instance can access only a single set of tables.

```
     IMS:
     ====
          on MVS1                 on MVS2                on MVS3

            IMS                     IMS                    IMS
            I1TM                    I2TM                   I3TM



        66 REGIONS              66 REGIONS             66 REGIONS

               IFP                     IFP                    IFP
                1                       2                      3



               WFI                     WFI                    WFI
                1                       2                      3



     DB2:
     ====

             DB1G                    DB2G                  DB3G

                 CPC 1                        CPC 2                       CPC 3



 SYSPLEX
  DASD
 =======



          REPLICATE A            REPLICATE B            REPLICATE C
             TABLES                 TABLES                 TABLES
        +---------------+      +---------------+      +---------------+
        |     A.TABA    |      |     B.TABA    |      |     C.TABA    |
        |     A.TABB    |      |     B.TABB    |      |     C.TABB    |
        |        †      |      |        †      |      |        †      |
        |        †      |      |        †      |      |        †      |
        |     A.TABI    |      |     B.TABI    |      |     C.TABI    |
        +---------------+      +---------------+      +---------------+
```

IFP IMS Fast Path Region

WFI IMS Wait For Input Message Processing Region

The replicated tables are different tables in the DB2 data sharing group.

The transactions are generated randomly by TPNS (Teleprocessing Network Simulator).

Note: CPC 1 (First CPC), CPC 2 (Second CPC), and CPC 3 (Third CPC) are synonymous with members DB1G, DB2G, and DB3G respectively.

Figure 12. Test Environment for 3WP Execution (3-way)

5.3.2 Configuration and Definitions for the Performance Test

The hardware for the measurements consists of the following:

- One 9672-R61 processor (6 CMOS processors) for the 1-way run; two 9672-R61 processors for the 2-way run and three 9672-R61 processors for the 3-way run.

  Each CPC in 9672 has 1 giga byte (GB) of central storage.

- Two 9674-C01 Coupling Facility (6-way) boxes with 2GB of central storage: the first Coupling Facility box has the lock structure and the second Coupling Facility box has the SCA and GBP.

  For configuration requiring better availability, we recommend putting the lock and SCA structures in one coupling facility and have the space available for rebuild on the second coupling facility. In the IRWW measurement runs the SCA structure was put on the second coupling facility.to isolate the lock structure for performance study.

- One 9672-R61 with 1GB of central storage to run TPNS.

- DASD: mix of 3390, 3380 and RAMAC disks

  - 310 x 3390
  - 150 x 3380
  - 60 x 3380 RAMAC.

  The access to disks from CPC in the parallel sysplex environment is through the following control units:

  - 25 x 3990-3
  - 4 x 3990-2
  - 4 x 9394-3.

  The RAMAC disks are connected to 3990 control units.

Each CPC, has MVS/ESA SP Version 5 Release 2 with IMS V4.1 and DB2 V4 (see Figure 12 on page 49), and uses two links to each of two Coupling Facility boxes (4 links from each CPC).

Coupling Facility Policies Table 10 on page 51 shows the defined sizes and allocations for structures in the Coupling Facility. Sizes are in K bytes where K is 1000.

Table 10. Size and Allocation for Coupling Facility Structures			
Structure Name	**Coupling Facility Name**	**Defined Size**	**Actual Size**
DSNDB0G_LOCK1	LF01	64000 K	64000 K
DSNDB0G_SCA	LF02	10000 K	10240 K
DSNDB0G_GBP0	LF02	2640 K	2816 K
DSNDB0G_GBP1	LF02	14200 K	14336 K
DSNDB0G_GBP3	LF02	88000 K	88064 K
DSNDB0G_GBP4	LF02	20000 K	20224 K
DSNDB0G_GBP5	LF02	75000 K	75008 K
DSNDB0G_GBP6	LF02	39600 K	39680 K
DSNDB0G_GBP7	LF02	22000 K	22016 K
DSNDB0G_GBP8	LF02	133000 K	133120 K

Virtual Buffer Pool and Global Buffer Pool Settings Table 11 shows the defined sizes and allocations for buffers in virtual buffer pools for each DB2 member and GBPs in the Coupling Facility.

Table 11. Size and Allocation for Virtual Buffer Pool and Global Buffer Pool					
Buffer pool	**Virtual Pages**	**GBP Data Pages**	**GBP Directory Entries**	**Directory to Data Ratio**	**Allocated Size (KB)**
BP0	1000	620	1247	2.0	2816
BP1	5500	3253	6192	1.9	14336
BP2	31250	–	–	–	–
BP3	18750	66000	19000	3.5	88064
BP4	3125	7019	4679	1.5	20224
BP5	6250	16000	49000	3.0	75008
BP6	12500	8277	33000	4.0	39680
BP7	1000	4678	16000	3.5	22016
BP8	25000	29000	89000	3.1	133120

Note:

Class Castout Threshold = 10% for GBP0, GBP3, GBP4, GBP5, GBP6, GBP7, GBP8 and 90% for GBP1.

Group Buffer Pool Castout Threshold = 50% for GBP0, GBP3, GBP4, GBP5, GBP6, GBP7, GBP8 and 90% for GBP1.

Group Buffer Pool Checkpoint Interval = 60 minutes for all group buffer pools.

For more information on GBP Castout. see section 8.4.1, “Group Buffer Pool Class Castout Threshold” on page 130.

5.3.3 Description of the IRWW Workload

Table 12 shows an average profile of all transactions, obtained from the DB2 PM Accounting Report. The input to the DB2 PM Accounting report is the DB2 accounting trace records.

Particular focus is placed on studying lock avoidance, lock contention, hot data, and so on, to make sure that potential problem areas are covered. We believe that the IRWW workload is a better way of testing DB2 data sharing because:

- The typical transactions are complex. However, the transactions are not uniform in complexity, some are highly update intensive, others are highly read intensive.

  Very little time is spent in application code and in the transaction manager, which increases the effect of the DB2 contribution to CPU time.
- The workload has multithread capability.
- The workload has the ability to cause logical and physical contention.
- The workload has a broad spectrum of SQL functionality.
- The workload has many application and database design features typical of customer applications.
- The workload can vary the level of stress on the native machines by changing the think time in TPNS.

Table 12. Transaction Profile (Average Number for All Transactions)

Profile:	**Average Per Transaction**
SQL DML	
SELECT	4.85
INSERT	2.90
UPDATE	3.74
DELETE	0.21
OPEN	4.60
FETCH	9.30
CLOSE	2.65
DML-ALL	28.24
Buffer Pool - TOT4K	
GETPAGES	72.40
BUFFER UPDATES	14.77
SYNCHRONOUS READ	7.62
SEQUENTIAL PREFETCH	1.26
LIST PREFETCH	0.26
DYNAMIC PREFETCH	0.80
PAGES READ ASYNCHR.	37.36

5.3.4 Transactions defined for IRWW

Table 13 shows RO if the transactions only read data from DB2 tables, or R/W if the transactions update tables through reads and writes. The row labeled *Number of Transactions* is the observed number executed in a period of 10 minutes by the first CPC of the 2WP measurement.

Table 13. Transaction Mix

Transaction Attribute	Transaction Identification						
	T1	**T2**	**T3**	**T4**	**T6**	**T7**	**T8**
Read/Update SQL	R/W	R/W	RO	RO	R/W	R/W	RO
% Transaction Mix - Specified	22%	22%	24.5%	4%	2%	0.5%	25%
Number of Transactions	4082	3996	4399	741	251	94	4533
Transaction Mix - Measured	22.56%	22.08%	24.30%	4.09%	1.38%	0.52%	25.05%

Note:

RO Read-Only Transactions: T3, T4 and T8.

R/W Update transactions: T1, T2, T6 and T7.

Transactions vs. Tables Matrix: Table 14 shows the transactions from the IRWW workload and the read (R), write (W) or read/write (R/W) SQL activity against tables.

Table 14. Activities of Transactions against Tables

Table Accessed	Transaction Identification						
	T1	**T2**	**T3**	**T4**	**T6**	**T7**	**T8**
TABA	R	R/W	–	–	–	–	–
TABB	R/W	R/W	–	R	–	–	–
TABC	R	R/W	R	–	R/W	–	–
TABD	R/W	–	–	R	–	–	–
TABE	R	–	–	–	–	R/W	R
TABF	–	W	–	–	–	–	–
TABG	W	–	–	–	R/W	–	–
TABH	W	–	R	–	R/W	–	–
TABI	W	–	R	R	R/W	–	–

5.4 Performance Calculations for Data Sharing Measurements

The objective of this section is to obtain indicators for:

- Throughput
- Capacity delta for DB2 data sharing
- Activity in the coupling facility
- Contention on the locks for access to shared tables.

We can summarize the results of this workload as follows:

- The measurements show that capacity can be scaled well using type 2 index and page locking where the capacity delta is 13.29% for 2-way and 13.55% for 3-way with 100% data sharing.
- The usage of type 1 index results in significantly more overhead for 2-way data sharing (25.75%).
- Row locking results in a capacity delta of about 2% for 1-way and 17% for 2-way data sharing compared with page locking.

The details for the calculation are shown for 1WP and 2WP measurements in section 5.4.4, "Performance Indicators for 1WP (1-way, Type 2 Index and Page Locking)" on page 57. For other measurements, the calculations are repetitive with the same formulas and similar descriptions.

5.4.1 Abbreviations Used in Calculations

1WP
: 1-way measurement, type 2 index and page locking

2WP
: 2-way measurement, type 2 index and page locking

3WP
: 3-way measurement, type 2 index and page locking

1W1
: 1-way measurement, type 1 index and page locking

2W1
: 2-way measurement, type 1 index and page locking

1WR
: 1-way measurement, type 2 index and row locking

2WR
: 2-way measurement, type 2 index and row locking

ETR
: External throughput rate: number of transactions executed in the 10 minute interval at a CPU utilization of 70%

ITR Internal throughput rate: number of transactions normalized for 100% CPU utilization.

XLR
: Number of XES lock requests to coupling facility

TC%
: Total contention percentage: the percentage of suspends divided by the number of XES lock requests.

5.4.2 Calculation formulas for Each Member

The calculation formulas for each member are:

CPU %Busy: Average percentage of CPU capacity used by all processors in the CPC as indicated in the RMF CPU activity report.

External Throughput Rate: Total of number of transactions executed as indicated in the total number of commits in the DB2 PM Statistics report.

```
ETR = TOTAL COMMITS
```

Formula to Calculate Internal Throughput Rate

```
            ETR
ITR = ------------- * 100
         CPU %BUSY
```

Formula to Calculate Total XES Lock Requests (XLR) The formula does not include all XES lock requests. ASYNC.XES requests can be ignored as they should be very low.

```
Total XES Lock Requests (XLR) = (SYNCH.XES_LOCK_REQUESTS   +
                                 SYNCH.XES_CHANGE_REQUESTS +
                                 SYNCH.XES_UNLOCK_REQUESTS +
                                 ASYNCH.XES_RESOURCES      +
                                 SUSPENDS_IRLM_GLOBAL_CONT +
                                 SUSPENDS_XES_GLOBAL_CONT  +
                                 SUSPENDS_FALSE_CONTENTION  )
```

The fields used in the formula are in the Data Sharing Locking section in the DB2 PM Statistics report.

Formula to Calculate Total Global Lock Suspensions

```
Total Global Lock Suspensions (GLS) = (SUSPENDS_IRLM_GLOBAL_CONT +
                                       SUSPENDS_XES_GLOBAL_CONT  +
                                       SUSPENDS_FALSE_CONTENTION  )
```

The fields used in the formula are in the Data Sharing Locking section in the DB2 PM Statistics report.

Formula to Calculate Total GBP Accesses (GBPA)

```
Total GBP Accesses (GBPA) = (SYN.READS(XI)_DATA_RETURNED +
                             SYN.READS(XI)_R/W_INTEREST  +
                             SYN.READS(XI)_NO_R/W_INTER. +
                             SYN.READS(NF)_DATA_RETURNED +
                             SYN.READS(NF)_R/W_INTEREST  +
                             SYN.READS(NF)_NO_R/W_INTER. +
                             ASYNC.READS_DATA_RETURNED   +
                             ASYNC.READS_R/W_INTEREST    +
                             ASYNC.READS_NO_R/W_INTEREST +
                             CLEAN PAGES_SYNC.WRITTEN    +
                             CHANGED PAGES_SYNC.WRITTEN  +
                             CLEAN PAGES_ASYNC.WRITTEN   +
                             CHANGED PAGES_ASYNC.WRITTEN +
                             PAGES_CASTOUT               +
                             OTHER_REQUESTS               )
```

The fields used in the formula are in the Group TOT4K section in the DB2 PM Statistics report.

5.4.3 Calculation formulas for Each Run

The performance indicators to evaluate the impact of data sharing are calculated using the following formulas.

Formula to Calculate Data Sharing Capacity Delta: Data sharing capacity delta because of additional capacity required for equivalent throughput on N-way data sharing is calculated using the following formula:

```
                    N * (1-way ITR) - (ITR1 + ITR2 + ... + ITRN)
Capacity Delta = ---------------------------------------------- * 100
                                 N * (1-way ITR)
```

Formula to Calculate Coupling Efficiency

```
Coupling Efficiency (%) = 100% - Capacity Delta
```

Formula to Calculate Increased CPU Time for N-way data sharing

```
                       N * (1-way ITR) - (ITR1 + ITR2 + ... + ITRN)
Increased CPU Time = ---------------------------------------------- * 100
                                (ITR1 + ITR2 + ... + ITRN)
```

Formula to Calculate Total Contention Percentage

```
                                  (GLS1 + GLS2 + ... + GLSN)
Total Contention Percentage = ------------------------------ * 100
                                  (XLR1 + XLR2 + ... + XLRN)
```

We sum the values for each member in the measurement.

Formula to Calculate Total False Contention Percentage

```
                                         Sum(SUSPENDS_FALSE_CONTENTION)
Total False Contention Percentage = ---------------------------------- * 100
                                          (XLR1 + XLR2 + ... + XLRN)
```

We sum the values for each member in the measurement.

Formula to Calculate Number of XES Lock Requests per Commit

```
                                (XLR1 + XLR2 + ... + XLRN)
XES Lock Requests per Commit = ------------------------------
                                (ETR1 + ETR2 + ... + ETRN)
```

Formula to Calculate Number of GBP Accesses per Commit

```
                          (GBPA1 + GBPA2 + ... + GBPAN)
GBP Accesses per Commit = ---------------------------------
                           (ETR1 + ETR2 + ... + ETRN)
```

These formulas which are used in the following sections allow you to evaluate the throughput when using the DB2 data sharing and to identify the factors that affect the additional cost in these environments.

For information on tuning, see Chapter 8, "Instrumentation and Tuning" on page 91.

5.4.4 Performance Indicators for 1WP (1-way, Type 2 Index and Page Locking)

The CPU consumption of the 1-way processing with type 2 index and page lockng in the data sharing environment is approximately the same as the corresponding processing without data sharing in the same MVS sysplex. The key factor for this result is because we use IMS IFP and WFI regions, and plans bound with RELEASE(DEALLOCATE) to maximize thread reuse and retention of locks. Therefore, XES/CF global locking traffic is negligible during the measured period of 10 minutes.

Figure 13 on page 58 shows the configuration used for a 1-way run.

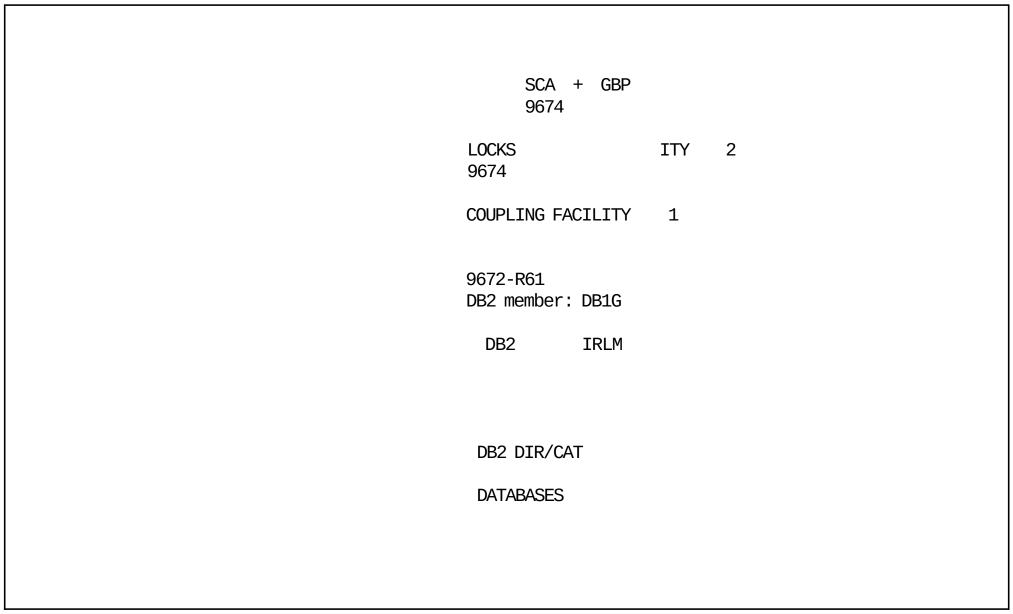

Figure 13. Configuration for 1WP Measurement

Obtaining CPU %Busy for 1WP: Figure 14 shows an extract from the RMF CPU Activity report with the TOTAL/AVERAGE MVS BUSY TIME PERC A : the average CPU-busy time percentage for the six processors for 1-way execution is 72.53%. This percentage is the average over approximately 10 minutes (9 minutes and 37 seconds).

C P U A C T I V I T Y

MVS/ESA SP5.2.0 | SYSTEM ID 961 | DATE 09/09/95 | INTERVAL 09.36.745
RPT VERSION 5.2.0 | TIME 12.47.42 | CYCLE 1.000 SECONDS

CPU MODEL 9672 VERSION 37

CPU NUMBER	VF ONLINE	VF AFFINITY PERCENTAGE	LPAR TIME	MVS BUSY TIME PERC	CPU SERIAL NUMBER	I/O TOTAL INTERRUPT RATE	% I/O INTERRUPTS HANDLED VIA TPI
0	---	****	----	85.37	000036	153.5	2.31
1	---	****	----	80.72	100036	131.2	2.54
2	---	****	----	75.27	200036	107.0	2.50
3	---	****	----	69.22	300036	86.55	2.92
4	---	****	----	64.40	400036	67.98	3.59
5	---	****	----	60.18	500036	487.8	21.00
TOTAL/AVERAGE		****	----	A 72.53		1034	11.31

Figure 14. RMF CPU Activity Report for 1WP.

Obtaining ETR for 1WP: The ETR is the total number of commits. Figure 15 on page 59 shows an extract from the DB2 PM Statistics report for member DB1G in the 10 minute interval.

ETR = B = 21493

```
   LOCATION: DSNDB0G              DB2 PERFORMANCE MONITOR (V4)             PAGE: 1-1
      GROUP: DSNDB0G                STATISTICS REPORT - LONG         REQUESTED FROM: 09/09/95 12:47
     MEMBER: DB1G                                                                TO: 09/09/95 12:57
  SUBSYSTEM: DB1G                                                     INTERVAL FROM: 09/09/95 12:47
DB2 VERSION: V4                              SCOPE: MEMBER                       TO: 09/09/95 12:57

---- HIGHLIGHTS ------------------------------------------------------------------------------------
INTERVAL START  : 09/09/95 12:47:19.27  SAMPLING START: 12:47:19.27  TOTAL THREADS      :     0.00
INTERVAL END    : 09/09/95 12:57:19.33  SAMPLING END  : 12:57:19.33  TOTAL COMMITS    B : 21493.00
INTERVAL ELAPSED:        10:00.063955   OUTAGE ELAPSED:    0.000000  DATA SHARING MEMBER:      N/A
```

Figure 15. DB2 PM Statistics Report for 1WP with Total Commits. Extracted from Figure 73 on page 150.

Calculating ITR for 1WP: We calculate ITR as

$$ITR = \frac{B}{A} * 100 = \frac{21493}{72.53} * 100 = 29633$$

Obtaining Total Number of XES Lock Requests for 1WP: Figure 16 shows an extract from the DB2 PM Statistics report with the lock requests that we have to sum

$$\text{XES Lock Requests} = D + E + F + G + H + I + J$$

$$= 857 + 0 + 856 + 0 + 0 + 0 + 0 = 1713$$

```
           DB2 PERFORMANCE MONITOR (V4)                          PAGE: 1-4
             STATISTICS REPORT - LONG                   REQUESTED FROM: 09/09/95 12:47:18.00
                                                                    TO: 09/09/95 12:57:20.00
       ---------------------------------------------------------------------------
        SAMPLING START: 09/09/95 12:47:19.27    TOTAL THREADS      :     0.00
        SAMPLING END  : 09/09/95 12:57:19.33    TOTAL COMMITS      : 21493.00

                               DATA SHARING LOCKING         QUANTITY  /MINUTE  /THREAD  /COMMIT
                               ---------------------------  --------  -------  -------  -------
                               SYNCH.XES - LOCK REQUESTS    D 857.00    85.69      N/C     0.04
                               SYNCH.XES - CHANGE REQUESTS  E   0.00     0.00      N/C     0.00
                               SYNCH.XES - UNLOCK REQUESTS  F 856.00    85.59      N/C     0.04
                               ASYNCH.XES - RESOURCES       G   0.00     0.00      N/C     0.00

                               SUSPENDS - IRLM GLOBAL CONT  H   0.00     0.00      N/C     0.00
                               SUSPENDS - XES GLOBAL CONT.  I   0.00     0.00      N/C     0.00
                               SUSPENDS - FALSE CONTENTION  J   0.00     0.00      N/C     0.00
```

Figure 16. DB2 PM Statistics Report for 1WP with Lock Requests. Extracted from Figure 75 on page 152.

Obtaining Total Number of GBP Accesses for 1WP: The GBP is not used when there is only one member in the data sharing group because there is no inter-DB2 read-write interest. the DB2 PM Statistics report in this case does not present information about GBP accesses.

Obtaining Total Number of Global Lock Suspensions for 1WP: Figure 17 on page 60 shows an extract from the DB2 PM Statistics report.

$$\text{Global Lock Suspensions} = L + M + N$$

$$= 0 + 0 + 0 = 0$$

```
                    DB2 PERFORMANCE MONITOR (V4)                          PAGE: 1-4
                      STATISTICS REPORT - LONG               REQUESTED FROM: 09/09/95 12:47:18.00
                                                                         TO: 09/09/95 12:57:20.00
            ----------------------------------------------------------------------------
                                  DATA SHARING LOCKING          QUANTITY  /MINUTE  /THREAD  /COMMIT
                                  ---------------------------  --------  -------  -------  -------
                                  SUSPENDS - IRLM GLOBAL CONT  L   0.00     0.00      N/C     0.00
                                  SUSPENDS - XES GLOBAL CONT.  M   0.00     0.00      N/C     0.00
                                  SUSPENDS - FALSE CONTENTION  N   0.00     0.00      N/C     0.00
```

Figure 17. DB2 PM Statistics Report for 1WP with Global Lock Suspensions. Extracted from Figure 75 on page 152.

5.4.5 Performance Indicators for 2WP (2-way, Type 2 Index and Page Locking)

These indicators are compared with the 1-way measurements to evaluate the impact of having two members with inter-DB2 read-write interest and accessing the coupling facility for global locks and GBP.

Figure 18 shows the configuration used for a 2-way run.

```
                                      SCA  +  GBP
                                      9674

                                LOCKS                ITY    2
                                9674

                                COUPLING FACILITY    1


                  9672-R61                          9672-R61
                  DB2 member: DB1G                  DB2 member: DB2G

                    DB2       IRLM                    DB2       IRLM


                      First CPC    |                |    Second CPC

                                 DB2 DIR/CAT

                                 DATABASES
```

Figure 18. Configuration for 2WP Measurement

Obtaining CPU %Busy for 2WP (First CPC): Figure 19 shows that the first CPC is A 70.08% busy. This percentage is the average over approximately 10 minutes (9 minutes and 37 seconds).

```
                                     C P U   A C T I V I T Y

              MVS/ESA          SYSTEM ID 961          DATE 09/09/95          INTERVAL 09.37.205
              SP5.2.0          RPT VERSION 5.2.0      TIME 08.41.39          CYCLE 1.000 SECONDS
  CPU MODEL    9672  VERSION  37
  CPU     VF        VF AFFINITY    LPAR BUSY MVS BUSY     CPU SERIAL  I/O TOTAL           % I/O INTERRUPTS
  NUMBER  ONLINE    PERCENTAGE     TIME PERC TIME PERC     NUMBER     INTERRUPT RATE       HANDLED VIA TPI
  0       ---        ****          ----      84.03         000036      124.7                  1.73
  1       ---        ****          ----      78.80         100036      105.0                  1.87
  2       ---        ****          ----      72.58         200036      85.34                  1.94
  3       ---        ****          ----      66.57         300036      66.20                  2.37
  4       ---        ****          ----      61.44         400036      51.20                  2.83
  5       ---        ****          ----      57.08         500036      337.8                 17.43
  TOTAL/AVERAGE      ****          ----    A 70.08                     770.3                  8.78
```

Figure 19. RMF CPU Activity Report for 2WP (First CPC)

Obtaining ETR for 2WP (First CPC): Figure 20 shows an extract from the DB2 PM Statistics report for member DB1G in the 10 minute interval.

```
ETR1 =  B  = 18056
```

```
   LOCATION: DSNDB0G             DB2 PERFORMANCE MONITOR (V4)               PAGE: 1-1
      GROUP: DSNDB0G               STATISTICS REPORT - LONG        REQUESTED FROM: 09/09/95 08:41
     MEMBER: DB1G                                                              TO: 09/09/95 08:51
---- HIGHLIGHTS ----------------------------------------------------------------------------------
INTERVAL START  : 09/09/95 08:41:16.36   SAMPLING START: 08:41:16.36   TOTAL THREADS      :    0.00
INTERVAL END    : 09/09/95 08:51:16.73   SAMPLING END  : 08:51:16.73   TOTAL COMMITS   B  :18056.00
INTERVAL ELAPSED:         10:00.377445   OUTAGE ELAPSED:     0.000000   DATA SHARING MEMBER:     N/A
```

Figure 20. DB2 PM Statistics Report with Total Commits for 2WP (First CPC). Extracted from Figure 63 on page 138.

Calculating ITR for 2WP (First CPC)

```
          B        18056
ITR1 = ------- = --------- * 100 = 25765
          A        70.08
```

Obtaining CPU %Busy for 2WP (Second CPC): Figure 21 shows that the second CPC is D 70.99% busy. This percentage is the average over approximately 10 minutes (9 minutes and 36 seconds).

```
                                      C P U   A C T I V I T Y

            MVS/ESA             SYSTEM ID 962          DATE 09/09/95          INTERVAL 09.35.604
            SP5.2.0             RPT VERSION 5.2.0      TIME 08.41.41          CYCLE 1.000 SECONDS
CPU MODEL    9672  VERSION  37
CPU     VF      VF AFFINITY   LPAR      MVS BUSY     CPU SERIAL  I/O TOTAL         % I/O INTERRUPTS
NUMBER  ONLINE  PERCENTAGE    TIME      TIME PERC    NUMBER      INTERRUPT RATE     HANDLED VIA TPI
0       ---      ****         ----      84.64         000035      121.8                1.69
1       ---      ****         ----      79.49         100035      103.2                1.91
2       ---      ****         ----      73.55         200035      84.17                2.01
3       ---      ****         ----      67.46         300035      65.41                2.36
4       ---      ****         ----      62.39         400035      50.44                2.93
5       ---      ****         ----      58.44         500035      344.5               17.50
TOTAL/AVERAGE    ****         ----    D 70.99                     769.5                8.97
```

Figure 21. RMF CPU Activity Report for 2WP (Second CPC)

Obtaining ETR for 2WP (Second CPC): Figure 22 shows an extract from the DB2 PM Statistics report for member DB2G in the 10 minute interval.

```
ETR2 =  E  = 18191
```

```
   LOCATION: DSNDB0G              DB2 PERFORMANCE MONITOR (V4)               PAGE: 1-1
      GROUP: DSNDB0G               STATISTICS REPORT - LONG        REQUESTED FROM: 09/09/95 08:41
     MEMBER: DB2G                                                              TO: 09/09/95 08:51
---- HIGHLIGHTS ----------------------------------------------------------------------------------
INTERVAL START  : 09/09/95 08:41:16.41   SAMPLING START: 08:41:16.41   TOTAL THREADS      :    0.00
INTERVAL END    : 09/09/95 08:51:16.85   SAMPLING END  : 08:51:16.85   TOTAL COMMITS   E  : 18191.00
INTERVAL ELAPSED:         10:00.441800   OUTAGE ELAPSED:     0.000000   DATA SHARING MEMBER:     N/A
```

Figure 22. DB2 PM Statistics Report with Total Commits for 2WP (Second CPC). Extracted from Figure 64 on page 139.

Calculating ITR for 2WP (Second CPC)

```
         E       18191
ITR2 = ------- = --------- * 100 = 25625
         D       70.99
```

Calculating the Data Sharing Capacity Delta for 2WP (2-way, Type 2 Index and Page Locking)***:*** We calculate the capacity delta of 2-way when compared with 1-way as:

```
                  N * (1-way ITR) - (ITR1 + ITR2)
Capacity Delta = ----------------------------------- * 100
                         N * (1-way ITR)

                  2 * (29633) - (25765 + 25625)
               = -------------------------------- * 100 = 13.29%
                          2 * (29633)
```

The values used in the formula are the ITR values calculated in "Calculating ITR for 1WP" on page 59, "Calculating ITR for 2WP (First CPC)" on page 61, and "Calculating ITR for 2WP (Second CPC)."

Calculating the Coupling Efficiency for 2WP: We calculate the coupling efficiency as:

```
Coupling Efficiency (%) = 100 - 13.29 =  86.71 %
```

Calculating Increased CPU Time for 2WP: We calculate the increased CPU time for 2-way when compared to 1-way as:

```
                      N * (1-way ITR) - (ITR1 + ITR2)
Increased CPU Time = ----------------------------------- * 100
                             (ITR1 + ITR2 )

                      2 * (29633) - (25765 + 25625)
                   = -------------------------------- * 100 = 15.33%
                              25765 + 25625)
```

The values used in the formula are the ITR values calculated in "Calculating ITR for 1WP" on page 59, "Calculating ITR for 2WP (First CPC)" on page 61, and "Calculating ITR for 2WP (Second CPC)."

Obtaining Total Number of XES Lock Requests for 2WP (First CPC): Figure 23 on page 63 shows an extract from the DB2 PM Statistics report.

```
XES Lock Requests for Member DB1G (XLR1) =

     A   +   B   +    C   + D + E + F + G

  128900 + 64952 + 119600 +  0  + 621 +  0  + 1016 = 315089
```

```
         DB2 PERFORMANCE MONITOR (V4)                              PAGE: 1-4
           STATISTICS REPORT - LONG                   REQUESTED FROM: 09/09/95 08:41:15.00
                                                                   TO: 09/09/95 08:51:17.00
 ---------------------------------------------------------------------------
 SAMPLING START: 09/09/95 08:41:16.36    TOTAL THREADS      :     0.00
 SAMPLING END  : 09/09/95 08:51:16.73    TOTAL COMMITS      : 18056.00

                           DATA SHARING LOCKING          QUANTITY  /MINUTE  /THREAD  /COMMIT
                           ---------------------------   --------  -------  -------  -------
                           SYNCH.XES - LOCK REQUESTS   A   128.9K    12.9K      N/C     7.14
                           SYNCH.XES - CHANGE REQUESTS B 64952.00  6491.12      N/C     3.60
                           SYNCH.XES - UNLOCK REQUESTS C   119.6K    12.0K      N/C     6.62
                           ASYNCH.XES - RESOURCES      D     0.00     0.00      N/C     0.00

                           SUSPENDS - IRLM GLOBAL CONT E   621.00    62.06      N/C     0.03
                           SUSPENDS - XES GLOBAL CONT. F     0.00     0.00      N/C     0.00
                           SUSPENDS - FALSE CONTENTION G  1016.00   101.54      N/C     0.06
```

Figure 23. DB2 PM Statistics Report with Lock Requests for 2WP (First CPC). Extracted from Figure 65 on page 140.

Obtaining Total Number of XES Lock Requests for 2WP (Second CPC): Figure 24 shows an extract from the DB2 PM Statistics report.

```
XES Lock Requests for Member DB2G (XLR2) =

    H   +  I  +   J  + K + L + M + N

  143900 + 68893 + 127900 +  0  + 643 +  0  + 1085 = 342421
```

We observe in the reports using type 1 index and page locking (2W1) that the number of SYNCH.XES_LOCK_REQUESTS and SYNC.XES_UNLOCK REQUESTS are more than three times higher than the the corresponding numbers for the 2WP obtained here.

```
         DB2 PERFORMANCE MONITOR (V4)                              PAGE: 1-4
           STATISTICS REPORT - LONG                   REQUESTED FROM: 09/09/95 08:41:15.00
                                                                   TO: 09/09/95 08:51:17.00
 ---------------------------------------------------------------------------
 SAMPLING START: 09/09/95 08:41:16.41    TOTAL THREADS      :     0.00
 SAMPLING END  : 09/09/95 08:51:16.85    TOTAL COMMITS      : 18191.00

                           DATA SHARING LOCKING          QUANTITY  /MINUTE  /THREAD  /COMMIT
                           ---------------------------   --------  -------  -------  -------
                           SYNCH.XES - LOCK REQUESTS   H   143.9K    14.4K      N/C     7.91
                           SYNCH.XES - CHANGE REQUESTS I 68893.00  6884.23      N/C     3.79
                           SYNCH.XES - UNLOCK REQUESTS J   127.9K    12.8K      N/C     7.03
                           ASYNCH.XES - RESOURCES      K     0.00     0.00      N/C     0.00

                           SUSPENDS - IRLM GLOBAL CONT L   643.00    64.25      N/C     0.04
                           SUSPENDS - XES GLOBAL CONT. M     0.00     0.00      N/C     0.00
                           SUSPENDS - FALSE CONTENTION N  1085.00   108.42      N/C     0.06
```

Figure 24. DB2 PM Statistics Report with Lock Requests for 2WP (Second CPC). Extracted from Figure 66 on page 141.

Calculating the XES Lock Requests per Second for 2WP: We calculate this indicator as a total rate for 2-way processing:

```
                                  (XLR1 + XLR2)
XES Lock Requests per Second = -------------------
                                10 minutes * 60

                                (315089 + 342421)
                             = -------------------- = 1096
                                       600
```

These indicators give the first clue to the number of of locks propagated to the coupling facility. If the number of locks is high enough to impact the performance, there is a need to better understand the locking activity. See section 8.1, "Analyzing the Use of the Lock Structure" on page 94.

Obtaining GBP Accesses for 2WP (First CPC): We sum all fields in the QUANTITY column in Figure 25

```
GBP Accesses for Member DB1G (GBPA1) =

    A  +  B  +   C  +   D  +   E  +  F  +   G + H +

    I  +  J  +   K  +   L  +   M  +  N  +   O

=  22279 + 10  +   0  + 27537 + 81112 +  0  +   0 + 0
    0  +  0  + 110700 +   0  +  56  + 61901 +  28564

=  332159
```

```
GROUP TOT4K                   QUANTITY  /MINUTE  /THREAD  /COMMIT
---------------------------   --------  -------  -------  -------
SYN.READS(XI)-DATA RETURNED A 22279.00  2226.50      N/C     1.23
SYN.READS(XI)-R/W INTEREST  B    10.00     1.00      N/C     0.00
SYN.READS(XI)-NO R/W INTER. C     0.00     0.00      N/C     0.00
SYN.READS(NF)-DATA RETURNED D 27537.00  2751.97      N/C     1.53
SYN.READS(NF)-R/W INTEREST  E 81112.00  8106.10      N/C     4.49
SYN.READS(NF)-NO R/W INTER. F     0.00     0.00      N/C     0.00
ASYNC.READS-DATA RETURNED   G     0.00     0.00      N/C     0.00
ASYNC.READS-R/W INTEREST    H     0.00     0.00      N/C     0.00
ASYNC.READS-NO R/W INTEREST I     0.00     0.00      N/C     0.00

CLEAN PAGES SYNC.WRITTEN    J     0.00     0.00      N/C     0.00
CHANGED PAGES SYNC.WRITTEN  K   110.7K    11.1K      N/C     6.13
CLEAN PAGES ASYNC.WRITTEN   L     0.00     0.00      N/C     0.00
CHANGED PAGES ASYNC.WRITTEN M    56.00     5.60      N/C     0.00

PAGES CASTOUT               N 61901.00  6186.21      N/C     3.43

-...............................................................
OTHER REQUESTS              O 8564.00  2854.60      N/C     1.58
```

Figure 25. DB2 PM Statistics Report with GBP Accesses for 2WP (First CPC). Extracted from Figure 71 on page 146.

Obtaining GBP Accesses for 2WP (Second CPC).: We sum all fields in the QUANTITY column in Figure 26 on page 65

```
GBP Accesses for Member DB2G (GBPA2) =

    A  +  B  +   C  +   D  +   E  +  F  +   G + H +

    I  +  J  +   K  +   L  +   M  +  N  +   O

=  22291 + 15  +   0  + 29591 + 88606 +  0  +   0 + 0
    0  +  0  + 117700 +   0  +  40  + 45299 +  19845

=  323387
```

```
GROUP TOT4K                   QUANTITY  /MINUTE  /THREAD  /COMMIT
---------------------------   --------  -------  -------  -------
SYN.READS(XI)-DATA RETURNED A 22291.00  2227.46      N/C     1.23
SYN.READS(XI)-R/W INTEREST  B    15.00     1.50      N/C     0.00
SYN.READS(XI)-NO R/W INTER. C     0.00     0.00      N/C     0.00
SYN.READS(NF)-DATA RETURNED D 29591.00  2956.92      N/C     1.63
SYN.READS(NF)-R/W INTEREST  E 88606.00  8854.08      N/C     4.87
SYN.READS(NF)-NO R/W INTER. F     0.00     0.00      N/C     0.00
ASYNC.READS-DATA RETURNED   G     0.00     0.00      N/C     0.00
ASYNC.READS-R/W INTEREST    H     0.00     0.00      N/C     0.00
ASYNC.READS-NO R/W INTEREST I     0.00     0.00      N/C     0.00

CLEAN PAGES SYNC.WRITTEN    J     0.00     0.00      N/C     0.00
CHANGED PAGES SYNC.WRITTEN  K   117.7K    11.8K      N/C     6.47
CLEAN PAGES ASYNC.WRITTEN   L     0.00     0.00      N/C     0.00
CHANGED PAGES ASYNC.WRITTEN M    40.00     4.00      N/C     0.00

PAGES CASTOUT               N 45299.00  4526.57      N/C     2.49
-...............................................................
OTHER REQUESTS              O 19845.00  1983.04      N/C     1.09
```

Figure 26. DB2 PM Statistics Report with GBP Accesses for 2WP (Second CPC). Extracted from Figure 72 on page 147.

Calculating the GBP Accesses per Second for 2WP: We calculate this indicator as a total rate for 2-way processing:

```
                                (GBPA1 + GBPA2)
GBP Accesses per Second = -------------------
                                10 minutes * 60

                                (332159 + 323387)
                            = --------------------- = 1092
                                      600
```

These indicators give us the first clue to the number of accesses to GBP in the coupling facility. If the number of GBP accesses is high enough to impact the performance, there is a need to better understand the GBP usage. See section 8.2, "Analyzing the Use of the Group Buffer Pool" on page 112.

Calculating the Total CF Accesses per Second for 2WP:: We calculate this indicator as total rate for 2-way processing:

```
Total CF Accesses per Second = XES Lock Requests per Second +
                                GBP Accesses per Second

                             = 1096 + 1092 = 2188
```

Obtaining Total Number of Global Lock Suspensions for 2WP (First CPC): Figure 27 on page 66 shows an extract from the DB2 PM Statistics report.

```
Global Lock Suspensions for Member DB1G (GLS1) =  H  +  I  +  J

                                               = 621 +  0 +  1016 = 1637
```

DATA SHARING LOCKING	QUANTITY	/MINUTE	/THREAD	/COMMIT
SUSPENDS - IRLM GLOBAL CONT H	621.00	62.06	N/C	0.03
SUSPENDS - XES GLOBAL CONT. I	0.00	0.00	N/C	0.00
SUSPENDS - FALSE CONTENTION J	1016.00	101.54	N/C	0.06

Figure 27. DB2 PM Statistics Report with Global Lock Suspensions for 2WP (First CPC). Extracted from Figure 65 on page 140.

Obtaining Total Number of Global Lock Suspensions for 2WP (Second CPC): Figure 28 shows an extract from the DB2 PM Statistics report.

```
Global Lock Suspensions for Member DB2G (GLS2) = L + M + N

                                             = 643 + 0 + 1085 = 1728
```

DATA SHARING LOCKING	QUANTITY	/MINUTE	/THREAD	/COMMIT
SUSPENDS - IRLM GLOBAL CONT L	643.00	64.25	N/C	0.04
SUSPENDS - XES GLOBAL CONT. M	0.00	0.00	N/C	0.00
SUSPENDS - FALSE CONTENTION N	1085.00	108.42	N/C	0.06

Figure 28. DB2 PM Statistics Report with Global Lock Suspensions for 2WP (Second CPC). Extracted from Figure 66 on page 141.

Calculating Total Contention Percentage for 2WP: We calculate this indicator as:

```
                                  (GLS1 + GLS2)
Total Contention Percentage = ----------------- * 100
                                  (XLR1 + XLR2)

                                    (1637 + 1728)
                            = --------------------- * 100 = 0.51 %
                                 (315089 + 342421)
```

This contention percentage is low (less than 1%) and results in a better performance than the estimated percentage (analyzed in section 6.6, "Processor Capacity" on page 76).

Calculating Total False Contention Percentage for 2WP: We calculate this indicator as:

```
Total False Contention Percentage =

     (SUSPENDS - FALSE CONTENTION  J  for first CPC +
      SUSPENDS - FALSE CONTENTION  N  for second CPC)
   ---------------------------------------------------- * 100
                        (XLR1 + XLR2)

         (1016 + 1085)
   = --------------------- * 100 = 0.32%
      (315089 + 342421)
```

The percentage of false contentions (0.32%) is acceptable because it is less than 1%."

Note: With a further enhancement to a hashing algorithm, we have reduced false contention ten-fold to less than 0.03%. This may not necessarily be true in all cases.

Calculating Number of XES Lock Requests per Commit for 2WP: We calculate this indicator as

$$\text{XES Lock Requests per Commit} = \frac{(XLR1 + XLR2)}{(ETR1 + ETR2)}$$

$$= \frac{(315089 + 342421)}{(18056 + 18191)} = 18.14$$

Calculating Number of GBP Accesses per Commit for 2WP: We calculate this indicator as:

$$\text{GBP Accesses per Commit} = \frac{(GBPA1 + GBPA2)}{(ETR1 + ETR2)}$$

$$= \frac{(332159 + 323387)}{(18056 + 18191)} = 18.09$$

The last two performance indicators are tabulated in Table 7 on page 44 for comparison with corresponding indicators for type 1 index and with corresponding indicators for row locking.

Chapter 6. The IBM Relational Warehouse Workload in Parallel Sysplex

The purpose of this chapter is to analyze the performance profile of the IBM Relational Warehouse Workload with DB2 data sharing.

The assumptions underlying the following formulas and graphs are:

- An algorithm can be derived to evaluate data sharing performance.
- The behavior of DB2 data sharing in a parallel sysplex is substantially similar to that in other sharing environments, so that the same types of algorithm can be used.
- Specific conditions for the formulas developed are:
  - IBM Relational Warehouse Workload
  - Homogeneous S/390 microprocessor environment
  - Page locking and type 2 indexes.

Before applying these calculations to other workloads and environments, it may become necessary to replace the cost factors in the basic formula with your own measured values.

6.1 Cost of Data Sharing

Ideally, a DB2 subsystem using many identical central processors would generate a throughput equal to the number of central processors times the transaction rate achieved by each central processor, with the response time of each transaction remaining unchanged. However, because of increased processing cost per transaction and increased DB2 and IRLM system overhead, throughput may be reduced and the response times may increase. The additional processing cost is mostly due to the introduction of global locking and intersystem data coherency control.

The design and implementation of DB2 data sharing has many positive factors that can help to mitigate the adverse effects mentioned. For example, the design has been highly optimized to detect a global lock conflict efficiently, to avoid the extra processing cost associated with global locking unless it is necessary, and to use the coupling facility to efficiently maintain data coherency across multiple DB2 subsystems.

The overhead of DB2 data sharing can be compared with the overheads of creating and maintaining multiple copies of a database under different DB2 subsystems to increase the transaction capacity for the same common data. There are several drawbacks to this multiple-copy approach:

- Additional DASD storage is needed.
- CP and I/O overhead are incurred to make copies.
- Operational procedures for refreshing a copy must be set up.
- The copied data becomes obsolete as soon as the original data is changed.

With DB2 data sharing, on the other hand, multiple DB2 subsystems can access a single copy of the data, simplifying data management and saving DASD storage.

The data sharing overhead is the additional processing capacity required in a data sharing environment to achieve a throughput equivalent to that in a non-data-sharing environment. The overhead of DB2 data sharing is calculated in four different ways for an N-way sharing group:

- Capacity delta
- Coupling efficiency
- Processor capacity
- Total throughput.

Formulas for all these are developed in the following sections.

The DB2 data sharing cost is a function of the additional processing required to provide for inter-process concurrency control and data coherency in a data sharing environment. This cost varies based on the degree of data sharing. The factors that influence the DB2 data sharing cost are:

- Locking
  - Type of index
  - Lock size
  - ISOLATION option
  - CURRENTDATA option
  - RELEASE option
  - Commit frequency
  - Lock structure size.
- Hardware and Software Configuration
  - Relative engine speed of CPs running MVS and coupling facility control code (CFCC)
  - Coupling facility structure sizes
  - Coupling link configuration
  - Hardware and software levels and maintenance
  - Number of members in the group.
- Work load characteristics and effects
  - Effectiveness of lock avoidance
  - Thread reuse
  - Amount of real lock contention
  - Amount of false lock contention
  - Amount of XES lock contention
  - Processor and DASD device contention
  - Dynamics of workload during the day.

Most of the above factors are analyzed in Chapter 7, “Planning and Design Guidelines” on page 79.

6.2 Degree of Data Sharing

The degree of data sharing is an abstraction. Briefly defined, the "degree of data sharing" is the percentage of the total processing power that is dedicated to process applications using shared data.

Formally defined, the degree of data sharing covers three factors that interact in determining the data sharing overhead required for DB2 data sharing:

1. The percentage of total CPU processing capacity across the systems for applications accessing shared DB2 data with inter-DB2 read-write interest (VTAM, Transaction Manager, Application Program, Non-DB2 work, DB2 work, MVS, and so on).
2. Within the subset of overall CP processing capacity defined above, the ratio of SQL processing against shared data, to other processing (other SQL, other DBMS, transaction manager, application, and so on).
3. The data intensity of the SQL workload in terms of driving coupling facility access rates (accesses per second).

Section 7.1, "Degree of Data Sharing" on page 79 shows an example of degree of data sharing. The performance tests shown in this document have 100% degree of data sharing, that is, every processor has a purely data sharing workload.

6.3 Performance Projection

In this section we develop a formula for estimating processor capacity requirements for moving the IBM Relational Warehouse Workload from a single system to a parallel sysplex environment.

Earlier performance measurements in a sysplex environment have demonstrated that there are three cost factors that affect the performance overhead in a sysplex environment. The three factors are:

1 a multisystem management cost
2 a cost for the second member joining the sysplex, also called initial cost
3 a cost for each additional member, also called additional cost

For more details see *System/390 MVS Parallel Sysplex Performance*.

6.3.1 Multisystem Management Cost

The multisystem management cost (factor 1) is the difference in capacity between a single system and a 1-way sysplex. This is a fixed cost, typically between 3% and 4% with the first generation of CMOS processors when migrating from nonsysplex to sysplex configuration. The cost has multiple components, such as XCF processing, GRS processing, shared JES checkpoint, and miscellaneous items.

Because the DB2 performance tests compare 1-way data sharing against 2-way and 3-way this component of the cost does not show in the measurements. This value is estimated at 3% for DB2 data sharing and must be added to the measured cost.

Refer to *System/390 MVS Parallel Sysplex Performance*, section 1.3.1, "Multisystem Management Costs" for more details.

6.3.2 Initial Data Sharing Cost

The cost for the second member joining the sysplex, or the initial data sharing cost, (factor 2) is the difference in capacity between two 1-way sysplexes and one 2-way sysplex, expressed as a percentage of the optimal capacity. “Formula to Calculate Data Sharing Capacity Delta” on page 56 shows the general (n-way) formula. To establish this cost, we apply this formula to a 2-way sysplex.

6.3.3 Additional Data Sharing Cost

The cost for each additional member joining the sysplex (factor 3) is the difference in capacity between an n-way and an (n+1)-way sysplex, also expressed as a percentage of the optimal capacity, for *n* greater than 1.

6.3.4 Capacity Delta Percentage

The capacity delta is defined in “Capacity Delta” on page 39. With the above factors, the capacity delta percentage for a sysplex with *n* members and a 100% degree of data sharing can be expressed in the following formula:

```
Capacity Delta(%) =  1  +  2 + ( 3  * (n - 2))

                     1 = Multisystem management cost
                     2 = Initial data sharing cost
                     3 = Additional data sharing cost
```

The factor (n-2) is used because the additional data sharing cost applies only from Member 3 onward.

Because the degree of data sharing only influences the DB2 capacity, the capacity delta for a sysplex with *n* members and *d* degree of data sharing can be expressed with the following formula:

```
Capacity Delta(%) = 1  + d * ( 2 + ( 3  * (n - 2))) / 100
```

This is a general formula, valid for any type of sysplex environment. Factor 1 is fixed. The performance tests run at the IBM Santa Teresa Laboratory are used to obtain the two unknown factors (2 and 3). This will be used for a performance projection of a DB2 data sharing environment. The factors are:

The multisystem management cost. This is estimated as 3%.

```
Multisystem Management =  1  = 3
```

The initial data sharing cost is the measured capacity delta from Run 2WP over Run 1WP. In “Calculating the Data Sharing Capacity Delta for 2WP” on page 62, this is calculated as 13.29%.

```
Initial data sharing    =  2  = 13.3
```

The additional data sharing cost is the difference in cost between Run 3WP and Run 2WP. (See Table 4 on page 42). From our measurement calculations, this is 13.55 - 13.29 = 0.26.

```
Additional data sharing =  3  =  0.3
```

This results in the following formula:

```
Capacity Delta(%) = 3 + d * (13.3 + (0.3 * (n - 2))) / 100
```

6.3.5 Capacity Delta Formula

The formula is rounded up to provide a more conservative estimate. The new formula is:

IBM Relational Warehouse Workload Capacity Delta

```
Capacity Delta(%) = 3 + d * (14 + (0.5 * (n - 2))) / 100

            n = Number of members
            d = Degree of data sharing in percent
```

Note: This formula applies to the IBM Relational Warehouse Workload with page locking and type 2 index.

6.4 Capacity Delta Projection

The formula shown in section 6.3.5, “Capacity Delta Formula” is used to estimate the capacity delta in an n-way data sharing group with d degree of data sharing:

```
Capacity Delta(%) = 3 + d * (14 + (0.5 * (n - 2))) / 100
```

For example, a 16 member group, with each member using 50% of its capacity for data sharing, has a capacity delta of 13.5%:

```
n = 16
d = 50

Capacity Delta(%) = 3 + 50 * (14 + (0.5 * (16 - 2)))/100
       = 13.5%
```

Figure 29 on page 74 shows the estimated capacity delta for DB2 data sharing as calculated with the above algorithm for up to 16 members. The X-axis shows the ideal case: No data sharing, no capacity delta. The 50% line shows a degree of data sharing of 50%. The 100% line shows a degree of data sharing of 100%.

Figure 30 on page 74 shows a finer level of detail for a portion of Figure 29 on page 74, in the range from one to four members. Figure 30 on page 74 shows the estimated capacity delta for a degree of data sharing of 50% and 100%. The figure also includes the measured values for 2-way and 3-way data sharing with page locking and type 2 indexes (2WP and 3WP). The dotted line extrapolates these measurements up to four members.

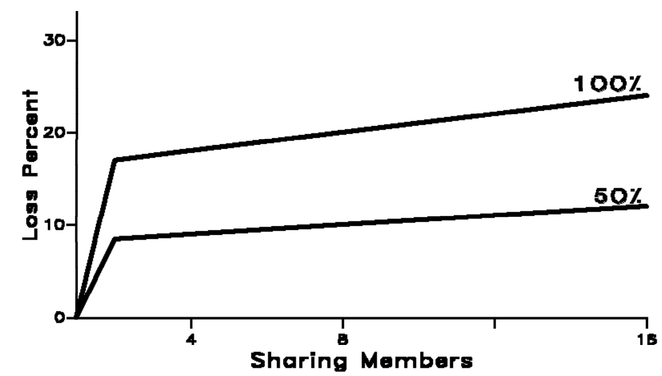

Figure 29. Projected Capacity Delta. The projections are for 50% and 100% data sharing

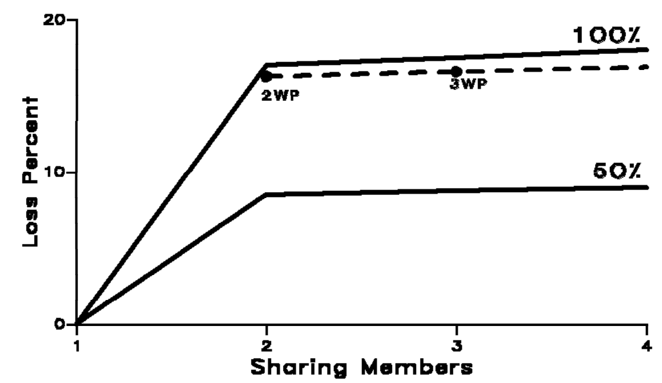

Figure 30. Measured Capacity Delta. (Solid lines are projections based on formula; dotted line is based on measurements and extrapolated to four members)

6.5 Coupling Efficiency

The coupling efficiency of an n-way data sharing group is defined as the percent of the total processor capacity that is productive in an n-way data sharing group.

The following algorithm can be used to estimate the coupling efficiency of an n-way data sharing group, in number of processors:

```
Efficiency(%) = 100 - Capacity Delta(%)

Capacity Delta(%) = Capacity delta for n members with
          a d degree of data sharing
```

For example: a 16 member group, each member using 50% of its capacity for data sharing has an efficiency of 86.5%:

```
n = 16
d = 50

Capacity Delta(%) = 13.5% (calculated on page 73)

Efficiency(%) = 100 - 13.5
              = 86.5%
```

Figure 31 shows the efficiency projected in the range from 1 to 16 members.

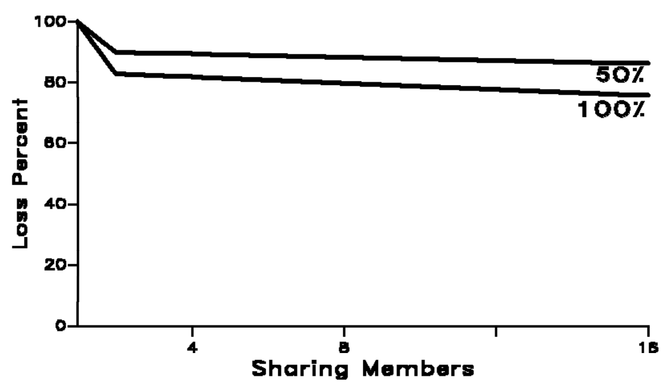

Figure 31. Projected Coupling Efficiency For 50% and 100% Data Sharing

6.6 Processor Capacity

The processor capacity of an n-way data sharing group is defined as the number of single processors that provide the same throughput as the n-way data sharing group.

The following algorithm can be used to estimate the processor capacity of an n-way data sharing group, in number of processors:

```
Processor Capacity = n * Efficiency(%) / 100

n = Number of members
Efficiency(%) = Coupling Efficiency
```

For example: a 16-member group, each member using 50% of its capacity for data sharing, has a capacity of 13.8 single processors:

```
n = 16
d = 50

Capacity Delta(%) = 13.5% (calculated on page 73)

Processor Capacity = 16 * 86.5 / 100
                   = 13.8
```

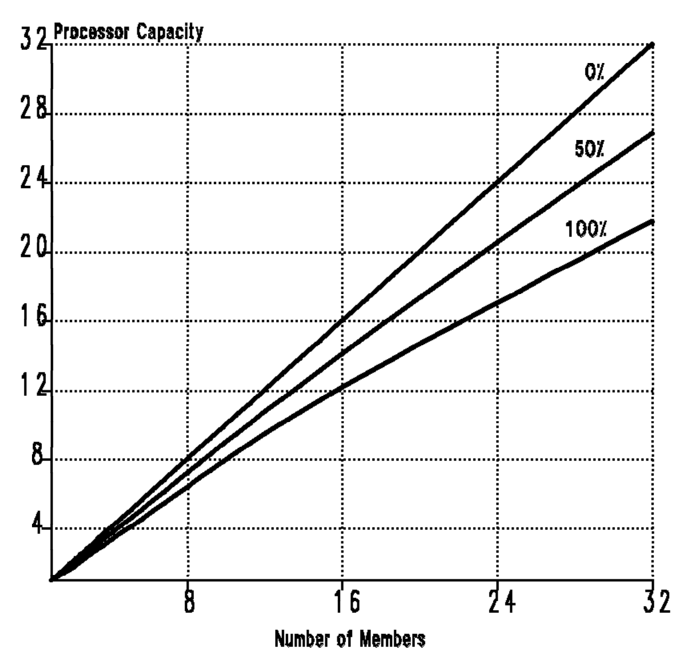

Figure 32. Estimated Processor Capacity For 0%, 50%, and 100% Data Sharing

Figure 32 shows the estimated processor capacity for DB2 data sharing as calculated with the above algorithm. The 0% line shows the ideal case, no loss due to sysplex and no loss due to DB2 data sharing. For example, a 16-member data sharing group has the capacity of 16 processors. The 50% line shows the estimate for a degree of data sharing of 50%. The 100% line shows the estimate for a degree of data sharing of 100%. For example, Figure 32 on page 76 shows that with 100% data sharing, a 16 member data sharing group has the capacity of 12 processors.

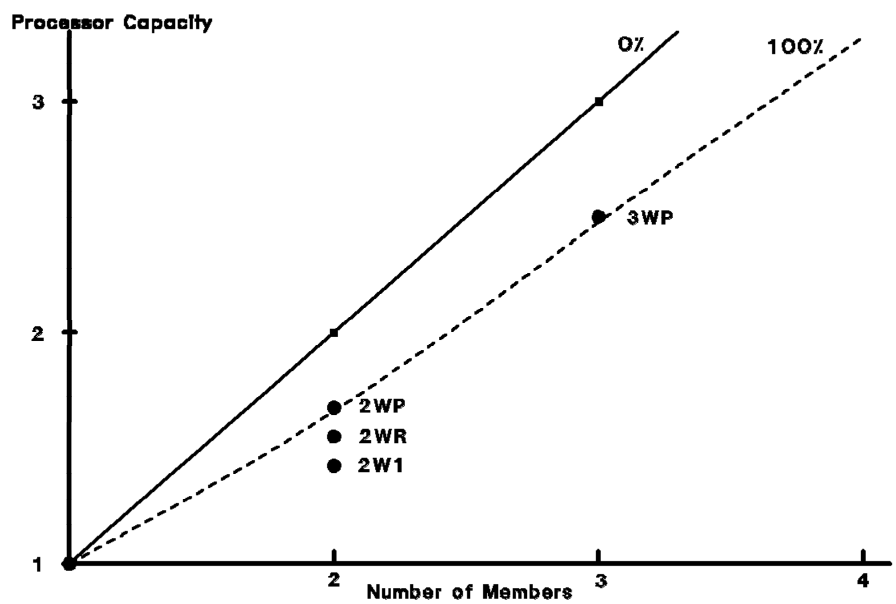

Figure 33. Measurements Plotted Against Estimated Processor Capacity

Figure 33 shows a finer level of detail for a portion of Figure 32 on page 76, in the range from one to four members. Figure 33 shows the estimated capacity for a degree of 0% and 100%. The figure also includes all the measured values. In Figure 33, the points on the 0% line show the ideal case. All measurement points must be below this line (or, theoretically, could be on this line). The dotted line shows the estimated processor capacity for 100% data sharing. The measurements 2WP and 3WP are just above the dotted line. This means, they show a better performance than the estimate line.

6.7 Total Throughput

The total throughput of an n-way data sharing group is defined as the number of transactions that can be processed by the n-way data sharing group. It is the processor capacity multiplied by the throughput of a single processor.

The following algorithm can be used to estimate the total throughput of an n-way data sharing group:

```
Throughput  = t * n * (100 - Capacity Delta%) / 100

n = Number of members
t = Single member throughput
```

Chapter 7. Planning and Design Guidelines

In this chapter, we describe the factors to consider when planning or designing for a data sharing environment. The purpose of this chapter is to enable you to obtain the maximum benefit from data sharing with DB2 V4.

The *DB2 for MVS/ESA V4 Data Sharing: Planning and Administration* Chapter 6 contains additional guidelines in the section *Improving Resource Utilization and Response Time*.

7.1 Degree of Data Sharing

The degree of data sharing is the percentage of the total processing power that is dedicated to process applications using shared data. The data sharing cost affects only this part of the total processing power. This section describes some of the factors that affect the cost of data sharing and the relative performance between environments with and without data sharing. We provide this information to allow you to make the best choices when planning for a data sharing environment.

Figure 34 on page 80 illustrates the degree of data sharing. The figure identifies four different workloads:

- Spare Capacity

  The amount of CP not utilized.

- DB2 Shared Workload

  The amount of CP dedicated to the DB2 applications which are being, or are going to be shared. It includes the online monitor, the application itself, and DB2 processing on behalf of the application. In Figure 34 on page 80, this represents approximately 30% of the total capacity.

- Other DB2 Workload

  The amount of CP dedicated to DB2 applications, which are not shared. It includes the same components as above.

- Non-DB2 Workload

  The amount of CP dedicated to all other applications.

When the workload shown in Figure 34 on page 80 moves to a parallel sysplex, the capacity delta due to data sharing overhead must be calculated with a degree of data sharing of 30%. It is only this part of the workload which is affected by the cost of data sharing.

```
      9021

    ••••••••
    ••••••••
    ••••••••
    ••••••••
    ••••••••
    ********
    ********
    ********
    ********
    ********
    #########
    #########
    #########
    #########
    #########



   #########   DB2 Shared Workload
   ********    Other DB2 Workload
   ••••••••    Non-DB2 Workload
               Spare Capacity
```

Figure 34. Degree of Data Sharing

7.2 Application Design

The parallel sysplex environment creates new challenges for designing well-behaved and well-designed applications.

7.2.1 Exploiting Parallelism

DB2 V4 has improved the query I/O parallelism introduced in DB2 V3. Executing on a parallel sysplex allows applications to exploit these features. Partitioned table spaces and applications running in parallel can improve overall throughput. Batch programs running in parallel against separate ranges or separate partitions can improve elapsed time. The *DB2 for MVS/ESA V4 Data Sharing: Planning and Administration* describes several ways to exploit parallelism in the parallel sysplex environment.

7.2.2 Application Profile

The number of table spaces with inter-DB2 read-write interest affects the degree of data sharing and the performance of a data sharing environment. The number of table spaces where there is inter-DB2 read-write interest determines child logical-lock propagation. The number of page sets or partitions having inter-DB2 read-write interest determines extent of the need for GBP-dependent functions.

You can isolate a given application and its private data on a given DB2 subsystem, to reduce the number of page sets or partitions with inter-DB2 read-write interest, and to reduce the number of page sets or partitions that are GBP-dependent. This provides better performance for that specific application. For affinity routing, to reduce the impact of GBP-dependent page sets, you would need to isolate all processing (across all applications, or utilities, or both) against a particular table space or partition on a single DB2 member for a given time period.

Affinity routing compromises availability and workload-balancing benefits of sysplex. Affinity routing may be very useful for parallel batch execution on one member.

7.2.3 Coupling Facility Access

The SQL workload of the application and the dynamics of the workload scheduling determine the coupling facility load. The coupling facility load can be measured in number of lock requests, number of group buffer pool requests, and number of shared communication area requests per second (see Table 7 on page 44 for examples). These numbers provide a means to describe and compare different SQL workloads in terms of coupling facility access load. These numbers also provide a way of measuring and observing the benefit of design and tuning actions for a given SQL workload.

Applications can be optimized for coupling facility accesses. The number of lock requests and the number of group buffer pool requests is a function of the SQL statements.

Techniques such as lock avoidance (see section 7.6.6, "Lock Avoidance" on page 89) may be used to reduce the global lock requests.

Optimizing SQL queries and analysis of access path may reduce the group buffer pool requests.

7.3 Processor Configuration

To make the best use of all processors, DB2 data sharing assigns functional system tasks to different processors. For example, one member may become the global lock manager for a table space, while another member may become a GBP castout class owner for the same table space.

This design assures a distribution of the system tasks among several processors and creates a requester/server relationship between the different members.

We use the phrase "data sharing server" in this book, for a member who is either a global lock manager for a lock table entry, a castout owner, or a GBP structure owner.

7.3.1 Relative Processor Speed

The coupling facility is another processor that acts as a server in a parallel sysplex. Most coupling facility accesses are performed synchronously to avoid the overhead of suspend and resume processing of the application or DB2 system task and avoid intermember message passing. When coupling facility accesses are made synchronously, the task on the MVS sender uses CP instruction cycles waiting for the response. This CP processing is charged back to the application task or the DB2 system service unit.

For example, assume an IRLM requests a global logical lock and the request is processed synchronously in the coupling facility. This example is more costly for an IRLM on a 9021 processor than on a 9672 processor, because the 9021 can process approximately 4 to 6 times more instructions than a 9672-Rx1 in the same time.

7.3.2 Homogeneous Configurations

DB2 V4 data sharing works best on homogeneous configurations, where all processors have a similar processing capacity (MIPS). To be specific, it is the relative model speed of the CPC running MVS versus the speed of the CPC running coupling facility. On homogeneous processors models of 9672 and 9674 the initial data sharing capacity delta (or overhead) for the IBM Relational Warehouse Workload is of the order of 17%, see 6.3, "Performance Projection" on page 71, while in an environment with 9021 processors and a 9674 coupling facility values are estimated to be in the order of 28%.

7.3.3 Mixed Configurations

Tuning and optimizing a mixed configuration can be difficult. Assume that a data sharing server resides on a slow processor. When a fast processor needs services from this data sharing server (a global lock manager for example), its response time is affected by the response time of the slower processor.

Because of the above, and in order to improve overall performance, a user may want to influence the assignment of the data sharing servers. It is not difficult at startup time to force a global lock manager onto a specific processor—for example: on the fast processor. However, this assignment may change at any time, because DB2 adapts dynamically to changes in the workload and can reassign the ownership of these functions.

The access path selection in a mixed configuration is affected by the different processor speeds. The access path for static SQL is selected at bind time. Because DB2 data sharing has a single catalog for a group, one bind creates the plan for all members. At execution time, the plan may execute on a different processor than at bind time. The access path may not be optimal if the execution processor is different from the processor used at bind time and affects CPU intensive queries which have several equally good candidate access plans. One option is to bind on the processor that is likely to be used to execute the majority of plans. Another option is to let the queries execute with dynamic SQL.

In summary, a homogeneous configuration with balanced workloads is much easier to tune and optimize than a mixed configuration.

7.4 Workload Distribution

Section 7.3.3, "Mixed Configurations" explains why DB2 V4 data sharing works better on homogeneous processors. The same reasons apply for a balanced workload. The effect of an unbalanced workload on the members is similar to that of a mixed configuration of processors. A processor that is 100% busy has long response times. If this busy processor has to act as a data sharing server, for example as global lock manager, the long response times affect a lock requester that has CP capacity to spare. A processor that takes on a high proportion of the workload is also more likely to become a data sharing server for the other members, thus affecting their response times even more.

The current recommendation is to offload a significant portion of your DB2 workload from a 9021 CP to 9672 CMOS processors and to distribute the workload evenly between these processors. A significant portion of your workload means at least 50% (hopefully 100%). The spare capacity obtained on the 9021 can be used for:

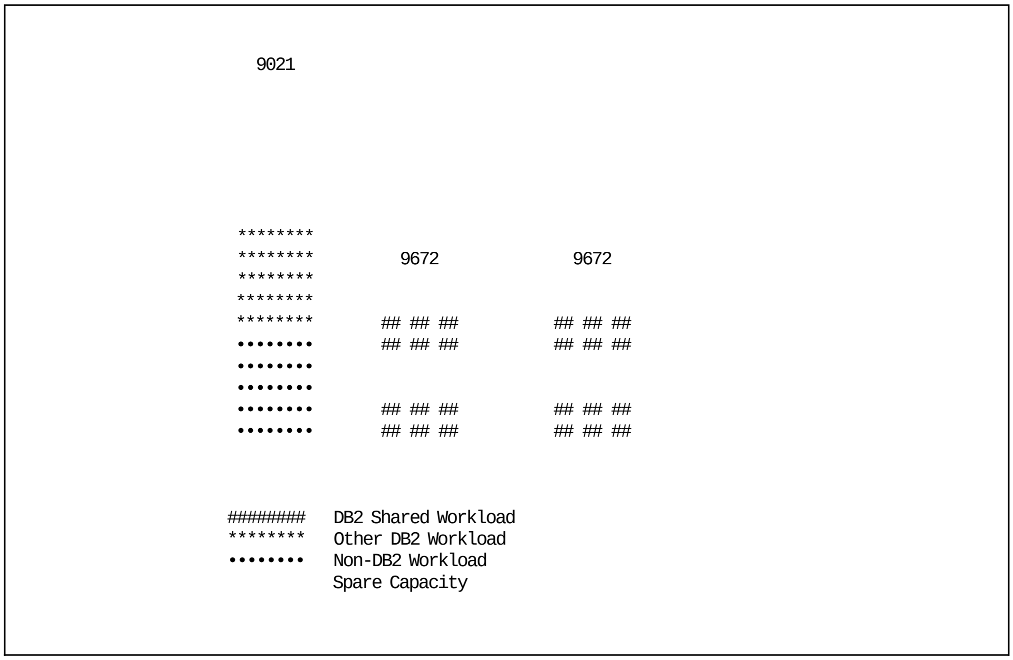

Figure 35. Example of Workload Distribution

- Non-DB2 work
- DB2 non-data-sharing work
- CP-bound batch and query DB2 data sharing work with low coupling facility access.

To illustrate this, Figure 35 shows a beneficial way to distribute the workload of Figure 34 on page 80 in a data sharing environment. On Figure 35 all data sharing is run on CMOS processors (homogeneous processors), and the workload is evenly balanced between these processors. With this type of distribution, the capacity delta is charged to the 9672 processor, maximizing the price-to-performance improvement.

7.5 Sysplex

The processing of global locks and the consistency of the buffer pools in the coupling facility contribute to the capacity delta in a data sharing environment. Because most coupling facility accesses are synchronous to the requester, their processing time adds to the application and DB2 processing time.

Two coupling facilities are recommended to ensure availability. The capacity of each should be adequate to handle the combined workload—their primary workload plus the workload of the alternative coupling facility. Each of the coupling facilities should have enough resources to cope with the failure of the other coupling facility. The coupling facility resources are

- Processor storage
- Coupling links
- Processors.

7.5.1 Processor Storage

The coupling facility uses storage for GBPs, the lock structure, and the SCA. The amount of processor storage dedicated to these structures directly affects performance and contention.

The lock structure should be allocated in exactly a power of two megabytes to obtain predictable results. The reason for this is that the two parts that make up the lock structure (the lock table and the record list) need an equal amount of storage and the size of the lock table must be an exact power of 2.

For DB2 data sharing, the following structures may require processor storage in the coupling facility:

- One SCA
- One lock structure
- Several GBPs.

See sections 8.1.5, “Allocating a Lock Structure” on page 109 and 8.3.4, “Group Buffer Pool Data Pages” on page 128, for details on sizing these structures.

Multiple GBPs can be spread around multiple coupling facilities, but no single GBP can straddle multiple coupling facilities.

The *DB2 for MVS/ESA V4 Data Sharing: Planning and Administration* provides information for planning and estimating the storage for all these structures.

For XCF message traffic, additional processor storage may be required.

7.5.2 Coupling Links

Number, utilization, and type of coupling links can affect the response time. Utilization here means how much traffic there is on the link in terms of coupling facility accesses per second. Link speed and link length affect response time.

Number of Links: Each coupling link has two subchannels. Most interactions with the coupling facility are synchronous. When a GBP access has to wait for a subchannel, it is converted into an asynchronous request. The conversion involves a suspension and resumption of the original task with the corresponding overhead. To reduce this overhead, coupling facilities should be configured with two links to each processor.

Traffic on Links: Keep normal link traffic at a low level. System checkpoint, GBP checkpoint, and GBP threshold processing cause bursts of coupling facility accesses and may cause queuing for subchannels.

Link Type and Length: Single-mode fiber transfer can go up to 50MB/s with a maximum distance of 1 km (0.6 mi) or, with RPQ, 3 km (1.9 mi). Multimode can transfer up to 100 MB/s with a maximum distance of 3 km (1.9 mi) or, with RPQ, 10 km (6.2 mi). Coupling-facility response time at the MVS host degrades with distance. The degradation is more significant with faster CP speed on the sender side.

7.5.3 Processors

Coupling-facility utilization and number of engines are related to coupling facility internal response time. It is recommended that the coupling facility be no more than 50% busy on average. Variability between 30% and 70% busy is acceptable. Do not run consistently above 70% as response times degrade. CPU consumption on the sender side also increases and, hence the capacity delta increases. Uniprocessors are not recommended for use as coupling facilities because of a steeper response time curve at utilization rates above 50%. This is illustrated in Figure 36, which shows the relative response time for a coupling facility with one, two, and six processors in the range of 30% to 70% processor busy.

See *System/390 MVS Parallel Sysplex Performance* for an algorithm to calculate the number of coupling-facility processors needed for a certain size of sysplex.

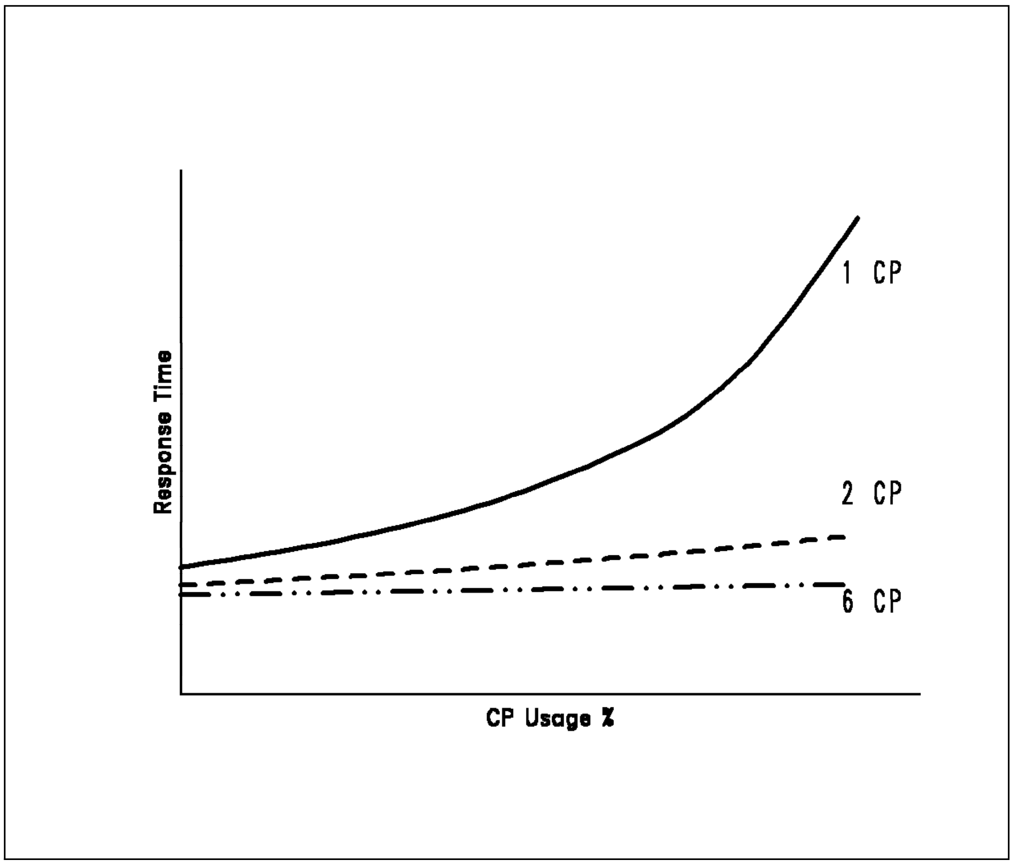

Figure 36. Coupling Facility Response Time for a Varying Number of Processors

7.5.4 Number of Members

The number of members participating in data sharing affects the usage of the lock table in the lock structure. The entries in the lock table need to change width to accommodate an increasing number of members.

The MAXUSRS parameter of the first IRLM to start in a group decides the width of the lock table entries at startup time. To avoid a rebuild of the lock structure, it is recommended that MAXUSRS specify the real maximum number of members that can be active in the group.

When MAXUSRS=7 or less, the entry size is 2 bytes long. The entry size does not change, as long as fewer than seven members are running. When the seventh member is started, the entry size changes to 4 bytes, resulting in a dynamic rebuild of the hash table. After the rebuild, there are fewer hash entries, with a higher probability of false contentions.

7.5.5 XCF Path

A channel-to-channel ESCON connection provides the best performing XCF path. A good alternative is the use of coupling facility connection for XCF paths. One benefit of using XCF paths through a coupling facility is that system administration work is simplified as the number of members in the group increases.

If the coupling facility is used to implement XCF paths, its storage and processing requirements must be considered while configuring the coupling facility.

7.6 Locking

In a data sharing environment, some global locks must propagate to the coupling facility. Thus, global locks contribute to the capacity delta for a data sharing environment. Reducing the number of global locks propagated provides the following benefits:

- Shorter transaction response times
- Reduction in the capacity delta required for data sharing
- Higher throughput.

Reduction in the number of lock contentions is important. The performance impact of lock contentions is significant because of the additional message traffic needed between the members to resolve the contention.

7.6.1 Locking Optimizations

Over the latest versions, DB2 has optimized locking to improve performance and reduce contention. DB2 V3 introduced data page lock avoidance. DB2 V4 introduces data-only locking and UR isolation as it improves lock avoidance DB2 V4 data sharing introduces physical locks for GBP-dependent objects. Figure 37 on page 87 illustrates the relative reductions in locks for the mix of transactions in the IRWW. The figure shows this impact on different DB2 versions. The values and the relationship between values must be considered as examples, because they are workload dependent and every application shows a different behavior. The IRWW uses ISOLATION(CS), and CURRENTDATA(NO). This maximizes the lock reduction obtainable with lock avoidance.

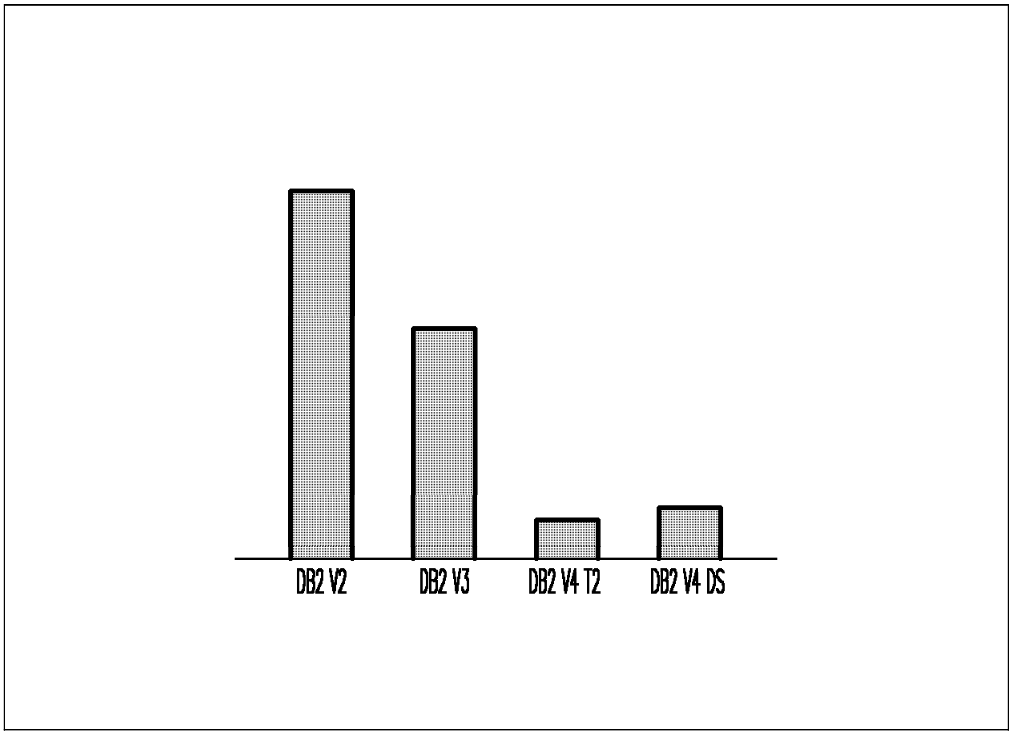

Figure 37. Number of Locks. Same workload on different DB2 versions (T2=type 2 index, DS=data sharing and type 2 index

7.6.2 Type 2 Indexes

The use of type 2 indexes for a data sharing environment is highly recommended. The difference in measurements of 2WP, with capacity delta of 13.29% and 2W1, with capacity delta of 25.75%, illustrate the advantage of type 2 indexes (see Table 4 on page 42 and Table 5 on page 43). The performance benefits from type 2 indexes are:

- Data-only locking, which reduces the total number of locks acquired by an SQL statement.
- CP parallelism for complex queries, because on CMOS processors complex queries may become CP bound.
- Uncommitted read isolation, which does not use any logical locks. This is applicable to user queries that can tolerate reading potentially uncommitted data updates.

Both the *IBM DATABASE 2 for MVS/ESA Version 4 Administration Guide* and the *DB2 V4 Non-Data-Sharing Performance Topics* contain more information on the benefits of type 2 indexes.

7.6.3 Lock Response Time

The following factors affect lock response time. They are listed in order of importance:

1. Processor configuration of coupling facility
   - Processor speed
   - Number of processors
   - Percent CP busy
2. Coupling link speed
3. Processor speed of requester
4. Type of request
5. XCF path speed
6. Coupling link length (when above 1 km or 0.6 mi).

7.6.4 Lock Size

The use of LOCKSIZE PAGE is highly recommended for a data sharing environment. LOCKSIZE ROW must be carefully evaluated and should not be used indiscriminately. The impact of row locking varies considerably from application to application, because it depends on the application profile, number of rows per page, and so on.

The difference in measurements of 2WP, with capacity delta of 13.29% and 2WR, with capacity delta of 19.57%, illustrate the advantage of page locking over row locking (see Table 4 on page 42 and Table 6 on page 43).

Applications with sequential access to rows with row locking may take multiple-row logical locks for each page accessed. However, applications with random access to data with row locking have as many locks as they would have, if they used page locking.

Lock avoidance (see 7.6.6, "Lock Avoidance" on page 89) is very important for row locking. In a data sharing environment, the use of LOCKSIZE ROW may result in more locks than LOCKSIZE PAGE, for update processes. The reason is that DB2 needs page physical locks to ensure data page coherency across multiple members (see section 4.1.3, "Page Physical Locks" on page 22). Page physical locks are not required in a non-data-sharing environment.

The additional overhead in data sharing with row locking is caused by page physical lock contention and negotiation. This occurs with very high page-access frequencies. At high page-access frequencies, applications serialize behind the page P-lock. Indiscriminate use of row locking can drive the global locking rate up significantly and thus lead to high overhead. Let us consider the two ends of a spectrum. Given 100% lock avoidance and read-only SQL, the performance of row locking approaches that of page locking. On the other hand, a data-intensive application process updating many rows per page takes many more locks; these locks may be propagated and cause performance to suffer.

In summary, *use page locking as a design default*. Previous design suggestions to reduce the number of rows per page (simulating row locking) still hold. For example, it is still useful to increase PCTFREE to reduce the number of rows per page and lower the page access frequency.

See *DB2 for MVS/ESA V4 Release Guide* for further information on when to use, or not to use row locking.

7.6.5 Lock Duration

RELEASE(DEALLOCATE) is the recommended lock duration option for plan or package BIND. RELEASE(DEALLOCATE) should be the design default. If any degree of thread reuse can be achieved in practice, then this selection helps to amortize the cost of plan allocation and deallocation over multiple units of work. Even at low multiprogramming levels, RELEASE(DEALLOCATE) has the potential to reduce data sharing cost for a given member against a data object. The reason for this is explained in section 4.5, "Lock Duration" on page 28.

7.6.6 Lock Avoidance

Without data sharing, lock avoidance depends on the highest committed log record sequence number (CLSN) of the single DB2. With data sharing lock avoidance depends on the highest committed log record sequence number of the group (global CLSN, or GCLSN).

Frequent commits of all processes in each member enables the GCLSN value to move forward and improve lock avoidance in a data sharing environment for GBP-dependent page sets or partitions. One long-running batch job without commits in one member retains its low LRSN for a long time, thus impairing lock avoidance for all members.

Specifying CURRENTDATA(NO) is of utmost importance for lock avoidance, hence for data sharing performance. CURRENTDATA(NO) must be explicitely specified, because the DB2 V4 default is CURRENTDATA(YES).

We strongly recommend that you specify CURRENTDATA(NO) for the majority of plans and packages. In a few cases CURRENTDATA(YES) may be required, as for rows read using implicit read-only cursor (ORDER BY) to encourage access via clustering index, followed by UPDATE or DELETE of the row in a noncursor operation.

7.6.7 Lock Contentions

It is important to keep down the number of global lock contentions. The purpose is to avoid the need for global lock management of the lock table entry and the associated message traffic.

The total number of lock requests propagated from IRLM to XES includes both synchronous and asynchronous requests. The vast majority of lock requests are synchronous for high performance. Asynchronous lock propagation is used, for example, when child logical locks are propagated and intersystem read/write interest first occurs for a parent resource.

The size of the lock structure influences false contentions. False contentions can be reduced by increasing the size of the lock structure.

Real contentions depend on application and database design, transaction elapsed time, transaction sequence and transaction mix. In a data sharing environment, real contentions may increase because of an increase in elapsed time caused by:

- XES and false contentions
- Higher DASD contentions
- Lower MIPS rate of CMOS processors.

The formula to evaluate global lock contention percentage is shown in "Formula to Calculate Total Contention Percentage" on page 56.

As a rule of thumb, less than 2% of lock requests should result in contention (real or false). Processes causing more than 2% contention must be investigated, to establish if the contentions are real or false.

If many contentions are false, an increase in lock structure size may help to reduce them.

If most of the contentions are real, sound application and database design are helpful. These are some typical actions that reduce real contentions:

- Optimize transactions to reduce elapsed time.
- Delay updates until just before commit.
- Use a consistent order of table accesses and updates.
- Increase PCTFREE and run REORG frequently.
- Evaluate LOCKSIZE ROW.

More details can be found in the *DB2 for MVS/ESA V4 Administration Guide*.

7.6.8 Locking Recommendations

We recommend the following to minimize lock requests:

- Use ISOLATION(CS) bind option.
- Use ISOLATION(UR) bind option, when appropriate.
- Use CURRENTDATA(NO) bind option.
- Use RELEASE(DEALLOCATE) bind option.
- Use page locks.
- Use type 2 indexes.
- Promote thread reuse as much as possible.
- Ensure sufficient commit frequency, for example, every 30 seconds.

Chapter 8. Instrumentation and Tuning

DB2 data sharing allows databases to be shared by multiple DB2 subsystems across multiple CPCs with full read/write capability. Data sharing comes with a cost which depends on the level of sharing achieved. The costs for DB2 data sharing depends on the degree of data sharing. DB2 adapts to the dynamics of the workload balancing across the group to minimize data sharing costs.

On the following pages, we outline where to start and what to look for when tuning a DB2 data sharing system. We analyze several structures, listed in order of the declining influence the structures can have on the performance of your system. The list represents a sequence of analysis we recommend if you are uncertain about the quality of a system′s performance:

- Lock structure
  This structure is responsible for DB2 data sharing group locking activity. A disturbance here can cause major performance degradation.
- Group buffer pool
  Shared data pass through the GBP to be available to other sharing members.
- Shared communication area
  The SCA maintains a consistent status across all data sharing members, including each member′s status and the database status.

Several tools allow you to access information about the structures we examine. We have not presented every tool that can be used; however, those presented are widely used and are sufficient to provide you with the information you need to tune your system.

- DB2 Performance Monitor (DB2 PM)
  DB2 PM is used to monitor group and member activities, online and in batch mode.
- DB2 DISPLAY commands
  DB2 DISPLAY commands provide information on the console and can be used to trigger activities through automated command processing when needed. DB2 DISPLAY commands have a snapshot characteristic; they reflect a particular moment in time that can be unique and may not be characteristic for your system.

  The DB2 DISPLAY commands can also be executed in a time-sharing option (TSO) batch environment. The output then appears on the job log rather than on the console.

  You can also use the instrumentation facility interface (IFI) and get the results of the DB2 DISPLAY command within your program.
- XCF DISPLAY commands
  XCF DISPLAY commands provide information on the console to show you how MVS XCF sees the DB2 data sharing group. It essentially lists member identifiers (IDs), structures used by the members, and the sizes and usage of the structures inside the coupling facility.
- Resource Measurement Facility (RMF)
  RMF lists operating system information at regular intervals from the systems management facility (SMF) data sets and informs about critical conditions. RMF Monitor I and RMF Monitor III allow an online display in a TSO session

so you can see dynamic changes online. RMF also allows you to monitor the MVS XCF and coupling facility activities.

All the tools we use display static and dynamic data. By static data, we mean anything that can be set or changed by an authorized person and does not vary with system load. By dynamic data, we mean counters that show system behavior and allow you to analyze performance.

In the following sections, we present an analysis of the detailed information for the lock structure, GBP, and SCA, using DB2 DISPLAY commands, XES DISPLAY commands, DB2 PM Accounting and Statistics reports, and RMF reports.

Table 15 summarizes the most important items listed and the tools used.

Table 15. Source of Static Information for Locking. Objects and how to find them

Objects	**D XCF**	**-DIS (DB2)**	**DB2 PM**	**RMF**
Structure status	STR			CF Activity
Data sharing members		GROUP		
Data sharing member names		GROUP		
Member status		GROUP		
IRLM ID		GROUP		
Lock structure size	STR	GROUP		

Table 16 (Page 1 of 2). Source of Dynamic Information for Locking. Objects and how to find them

Objects	**D XCF**	**-DIS (DB2)**	**DB2 PM**	**RMF**
Used lock structure size	STR	GROUP		
Locks %used		GROUP		
Lock table entry size	STR			
Page P-lock negotiation			Statistics, Accounting	
Number of P-lock requests			Statistics, Accounting	
Lock requests CF			Statistics, Accounting	CF Report
Lock suspends			Record Trace	
Contention count	STR		Statistics	CF Report

Table 16 (Page 2 of 2). Source of Dynamic Information for Locking. Objects and how to find them				
Objects	**D XCF**	**-DIS (DB2)**	**DB2 PM**	**RMF**
Lock service time				CF Report
CF requests				XCF Activity
CF queuing				XCF Activity
XCF messaging				XCF Activity

Figure 38 shows a collection of useful XCF commands. Their use is explained on the following pages.

```
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
@@      DISPLAY XCF DATA              @@
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
D XCF
D XCF,POLICY
D XCF,CF
D XCF,GROUP,DSNDB0G,ALL
D XCF,STR
D XCF,STR,STRNAME=DSNDB0G_LOCK1
D XCF,STR,STRNAME=DSNDB0G_SCA
D XCF,STR,STRNAME=DSNDB0G_GBP0
```

Figure 38. Useful XCF Commands

To find the structures names you use the command

```
D XCF,STR
```

This results in a list similar to that in Figure 39.

```
IXC359I  10.22.33  DISPLAY XCF 169
STRNAME            ALLOCATION TIME   STATUS
 1                  2                  3
DSNDB0G_GBP0       09/09/95 09:59:32 ALLOCATED
DSNDB0G_GBP1       09/09/95 09:59:13 ALLOCATED
DSNDB0G_GBP3       09/09/95 09:59:57 ALLOCATED
DSNDB0G_GBP32K        --       --    NOT ALLOCATED
DSNDB0G_GBP4       09/09/95 10:09:36 ALLOCATED
DSNDB0G_GBP5       09/09/95 09:59:57 ALLOCATED
DSNDB0G_GBP6       09/09/95 09:59:00 ALLOCATED
DSNDB0G_GBP7       09/09/95 09:59:19 ALLOCATED
DSNDB0G_GBP8       09/09/95 10:04:16 ALLOCATED
DSNDB0G_LOCK1      09/09/95 09:23:14 ALLOCATED
DSNDB0G_SCA        09/09/95 09:23:01 ALLOCATED
```

Figure 39. XCF Command Output for D XCF,STR

Three items of information can be gathered from Figure 39:

1 The names of the structures allocated in that XCF group. MVS allocates structures after DB2 requests connection. The names are implied based on the DB2 XCF Group Name.

2 The time the structure was last opened or allocated.

3 The status of the structure. If data sharing is not used for that particular buffer pool and the GBP is defined, then the GBP is not allocated.

8.1 Analyzing the Use of the Lock Structure

The lock structure is connected by IRLM during DB2 startup. As the lock structure is the most critical resource in data sharing, pay special attention to its performance.

8.1.1 Using the XCF DISPLAY Command

The MVS XCF command D XCF,STR,STRNAME=DSNDB0G_LOCK1 gives the output shown in Figure 40.

```
D XCF,STR,STRNAME=DSNDB0G_LOCK1

  ACTIVE STRUCTURE
  ----------------
   ALLOCATION TIME: 09/09/95 09:23:14
   CFNAME         : LF01
   COUPLING FACILITY: 009674.IBM.00.000000040016
                      PARTITION: 1   CPCID: 00
   ACTUAL SIZE    : 64000 K
   STORAGE INCREMENT SIZE: 256 K
   VERSION        : ABA5BA83 A3667001
   DISPOSITION    : KEEP
   ACCESS TIME    : 0
 A MAX CONNECTIONS: 7
   # CONNECTIONS  : 3
```

Figure 40. XCF Display Structure Example. D XCF,STR,STRNAME=DSNDB0G_LOCK1

A This is the maximum number of connections allowed to the structure as indicated in the CFRM policy. This number allows you to determine how many bytes your system uses for one entry in the lock table.

1–6 members Use a 2-byte lock table entry. As soon as the seventh member joins the group, 4-bytes are used for each lock entry.

7–22 members Use a 4-byte lock table entry. This results in 50% fewer entries on the same structure size as for 1–6 members, thereby doubling the false contention probability in comparison with the 2-byte table of the same size.

23–32 members Use an 8-byte lock table entry. This results in 50% fewer entries on the same structure size as for 7–22 members, thereby doubling the false contention probability in comparison with the 4-byte table of the same size.

The IRLM procedure has a parameter MAXUSRS that defaults to 7. It is initially set at DB2 installation time. For fewer than seven active members in the data sharing group, DB2 uses a lock table with 2-byte entries. When the seventh member joins the group, IRLM automatically changes the lock table in the lock structure to a 4-byte entry size. The change occurs independent of the

MAXUSRS setting in the IRLM procedure. The number of lock table entries is halved, thereby increasing the false contention probability. To return to the utilization before the occurrence of the change, you must define the lock structure twice as large as CFRM policy size. A CFRM policy determines how and where the structure resources are allocated. See section 8.1.5, “Allocating a Lock Structure” on page 109 for more details.

The change is achieved by invoking SETXCF REBUILD. It can have a dramatic effect; it can drive up false contention during SETXCF rebuild operation itself, and increase new lock requests from follow-on SQL processing. The ratio between lock table space and record list space remains identical to the ratio before the occurrence of the change.

8.1.2 Using the DB2 DISPLAY GROUP Command

We show the output of the DB2 command “-DISPLAY GROUP” in Figure 41

```
DSN7100I *DB1G DSN7GCMD
*** BEGIN DISPLAY OF GROUP(DSNDB0G )
-----------------------------------------------------------------
DB2                                        SYSTEM    IRLM
MEMBER    ID   SUBSYS  CMDPREF   STATUS    NAME      SUBSYS IRLMPROC
--------  ---  ----    --------  --------  --------  ----   -------- A
DB1G        1  DB1G    *DB1G     ACTIVE    STLABC1   DJ1G   DB1GIRLM
DB2G        2  DB2G    *DB2G     ACTIVE    STLABC2   DJ2G   DB2GIRLM
DB3G        3  DB3G    *DB3G     QUIESCED  STLABC2   DJ3G   DB3GIRLM
-----------------------------------------------------------------
SCA   STRUCTURE SIZE:    10240 KB, STATUS= AC,   SCA IN USE: <   1 %
LOCK1 STRUCTURE SIZE: B  64000 KB,             LOCK1 IN USE: C   1 %
NUMBER  LOCK ENTRIES:    16777216, LOCK ENTRIES  IN USE:     D 10240
NUMBER  LIST ENTRIES: E    228245, LIST ENTRIES  IN USE:     F  2593
*** END DISPLAY OF GROUP(DSNDB0G )
DSN9022I *DB1G DSN7GCMD ¢DISPLAY GROUP ¢ NORMAL COMPLETION
```

Figure 41. DB2 Display Group Output. -DIS GROUP

To obtain general information about all members of a particular group, use the DISPLAY GROUP command. The command can be issued from any active member of the group.

A The member names, subsystem names of the members, the MVS system names where the members are running, the IRLM subsystem names to which the members are connected, and the procedure names of the connected IRLMs are shown.

B LOCK1 STRUCTURE SIZE: The size of the allocated coupling facility lock structure is the size you specify in your CFRM policy. It is rounded to the next higher 256 KB boundary.

C LOCK1 IN USE: The percentage of the coupling facility lock structure actually in use at the moment of the display. It includes the space taken for the modified resource list and the space taken for the lock table. Because this is a point-in-time display, do not come to the conclusion that you should release the unused part of the structure. We recommend that you define the size of the coupling facility lock structure in such a way that contention is kept to less than 2%.

D LOCK ENTRIES IN USE: This shows the number of lock entries in use at the time of the display.

E NUMBER LIST ENTRIES: This is the counter for the modified resource list. The relation between space allocation for the list entries and for the lock entries is fixed by DB2 and cannot be changed.

F LIST ENTRIES IN USE: This is also called the “modified resource list.” Each entry represents an active modify lock or a retained lock. (L-locks and P-locks). The entry for the active modify lock is removed when the modify lock is released or changed to a nonmodify lock. The entry for the retained lock is released when the retained lock is purged by the restart process of the owning member.

8.1.3 Using DB2 PM Statistics Report

DB2 PM gives detailed usage information in the Statistics report, and Figure 42. shows an extract from this report. The input to the DB2 PM Statistics report is the DB2 statistics trace records produced by executing the DB2 command:

```
-START TRACE(STAT) CLASS(1,3,4,5)
```

to start a new trace, or

```
-MODIFY TRACE(STAT) CLASS(1,3,4,5) TNO(nn)
```

to modify trace number nn already started.

```
LOCKING ACTIVITY            QUANTITY  DATA SHARING LOCKING            QUANTITY
--------------------------  --------  ---------------------------     --------
SUSPENSIONS (ALL)            3752.00  LOCK REQUESTS (P-LOCKS)     A   27573.00
SUSPENSIONS (LOCK ONLY)       375.00  UNLOCK REQUESTS (P-LOCKS)       27553.00
SUSPENSIONS (LATCH ONLY)     3377.00  CHANGE REQUESTS (P-LOCKS)   B       4.00
SUSPENSIONS (OTHER)             0.00
                                      SYNCH.XES - LOCK REQUESTS   C     143.9K
TIMEOUTS                        0.00  SYNCH.XES - CHANGE REQUESTS     68893.00
DEADLOCKS                       0.00  SYNCH.XES - UNLOCK REQUESTS       127.9K
                                      ASYNCH.XES - RESOURCES              0.00
LOCK REQUESTS             L   160.6K
UNLOCK REQUESTS             59529.00  SUSPENDS - IRLM GLOBAL CONT D     643.00
QUERY REQUESTS                  0.00  SUSPENDS - XES GLOBAL CONT. E       0.00
CHANGE REQUESTS             69149.00  SUSPENDS - FALSE CONTENTION F    1085.00
OTHER REQUESTS                  0.00  INCOMPATIBLE RETAINED LOCK          0.00

LOCK ESCALATION (SHARED)        0.00  NOTIFY MESSAGES SENT                0.00
LOCK ESCALATION (EXCLUSIVE)     0.00  NOTIFY MESSAGES RECEIVED          692.00
                                      P-LOCK/NOTIFY EXITS ENGINES        10.00
DRAIN REQUESTS                  0.00  P-LCK/NFY EX.ENGINE UNAVAIL         0.00
DRAIN REQUESTS FAILED           0.00
CLAIM REQUESTS                166.1K  PSET/PART P-LCK NEGOTIATION G       0.00
CLAIM REQUESTS FAILED           0.00  PAGE P-LOCK NEGOTIATION     H     394.00
                                      OTHER P-LOCK NEGOTIATION            0.00
                                      P-LOCK CHANGE DURING NEG.         394.00
```

Figure 42. DB2 PM Statistics Report Showing All Locking Information. Extracted from Appendix Appendix A, “DB2 PM 2-Way Statistics Report (2WP)” on page 137)

A LOCK REQUESTS (P-LOCKS): Physical locks here represent all physical locks taken by one DB2 member. Page set physical locks are used to track inter-system read-write interest at the page set or partition level to determine if that page set or partition should be GBP-dependent. Page physical lock are used to allow coherency for one page when used by different data sharing members at the same time. The P-locks are used for example, in the following cases:

- Changes to space map pages
- Changes to type 2 index leaf pages
- Row locking.

Page set physical locks and page physical locks are also discussed in 4.1.2, “Page set Physical Locks” on page 21 and 4.1.3, “Page Physical Locks” on page 22 respectively.

Tuning information: A page P-lock in X-mode is released only when the updated page in the buffer pool is written back to GBP. Currently, the page P-lock is downgraded from X-mode to S-mode only if another sharing member asks for the same page P-lock in S-mode before the updating transaction is committed. For row locking, a page P-lock in S-mode is only requested if the transaction is using RR isolation, or is doing referential integrity checking.

For page set physical locks the transition from non-GBP-dependence to GBP-dependence and back is relatively expensive. Avoid continually switching in and out quickly. GBP-dependence is expensive, but staying GBP-dependent when it is possible to switch to non-GBP-dependence can also be expensive. Tune pseudo close parameters to arrive at an optimal set up.

Consider increasing the value for PCLOSEN in DSN6SYSP (the default is 5 for the number of checkpoints taken since last update before switching occurs) or PCLOSET in DSN6SYSP (the default is 10 minutes elapsed since last update until switching occurs). Each can help you to reduce the physical lock counter by reducing the number of transitions in and out of GBP-dependence. Both values can be changed in the DSN6SYSP macro in DSNZPARM. Consider also changing the LOGLOAD on DSNTIPN panel, the value for the checkpoint frequency. The default is 50000 for the number of log records that DB2 writes between successive checkpoints.

B CHANGE REQUEST (P-LOCKS): These are change requests by a member holding the lock because of a change in interest.

C SYNCH.XES - ...: The three SYNCH.XES counters for LOCK REQUESTS, CHANGE REQUESTS, and UNLOCK REQUESTS provide the total count of all synchronous lock requests. That count reflects a good part of the data sharing load. The count reflects S-mode and X-mode physical locks and logical locks, that do not experience contention. Synchronous means that the lock requests are propagated to XES synchronously.

Locks not propagated: You can compute the number of locks not propagated as follows:

```
L-Locks to CF = SYNCH.XES LOCK REQUESTS ( C ) - P-LOCK REQUESTS ( A )
              = 143900 - 27573
              = 116327

Locks not propagated = LOCK REQUESTS ( L ) - L-Locks to CF
                     = 160600 - 116327
                     = 44273
```

Therefore, the number of logical locks to the coupling facility is reduced by (44273 / 160600) * 100, that is 28% or we would have had (44273 / 116327) * 100, that is 38% more logical locks propagated to the coupling facility. With 100% data sharing, the reduction is mostly caused by the IRLM maintaining rolled-up states for a given data resource and only propagating and maintaining the most restrictive lock mode for a given resource. Another part of the reduction is attributed to hierarchical locking.

Physical locks are needed in addition to local latches. Latches synchronize access to some resources by the same member, a function that is provided by physical lock for access to the same resource by different members of the data sharing group.

Hierarchical locking and multiple agents for the same resource on the same data sharing member are responsible for the locks saved.

Tuning information: You can reduce lock requests through lock avoidance. Lock avoidance in a data sharing environment depends on having the lowest CLSN for committed, updated pages in the group used by any member. Keep the global CLSN moving forward. The global CLSN can be kept high (to maximize read lock avoidance) by frequent commits, especially in batch jobs in a data sharing environment.

We reemphasize and recommend the following to minimize lock requests:

- Use ISOLATION(CS) bind option.
- Use ISOLATION(UR) bind option, when appropriate.
- Use CURRENTDATA(NO) bind option.
- Use RELEASE(DEALLOCATE) bind option.
- Use page locks.
- Use type 2 indexes.
- Promote thread reuse as much as possible.
- Ensure sufficient commit frequency, for example, every 30 seconds.

In one of the IRWW measurement runs. the change to “CURRENTDATA(NO)” and “RELEASE(DEALLOCATE)” improves performance by 31% including the data sharing overhead.

D SUSPENDS - IRLM GLOBAL CONT: IRLM lock states are in conflict for the same resource. Unlike local contention that is resolved by the member IRLM, global contention requires global lock management to resolve the conflict. Global lock management is an expensive mechanism compared to a lock request granted synchronously without contention. The IRWW has about 0.51% global contention for 2WP measurement. Our recommendation is to keep IRLM global contention below 2% of

```
(SYNCH.XES - LOCK REQUESTS   +
 SYNCH.XES - CHANGE REQUESTS +
 SYNCH.XES - UNLOCK REQUESTS  )
```

When migrating to data sharing or implementing a new application, first make sure that real contention is acceptable in 1-way data sharing. This way you can be sure that you are dealing with application caused contention only and can try to reduce it, without interference by data sharing related contention.

E SUSPENDS - XES GLOBAL CONT: This is contention for the same resource in MVS XES terms, not in IRLM terms. This XES contention can be resolved by IRLM, as IRLM knows all lock states. To resolve it, the contention must be communicated by IRLM to the other contending members using MVS XCF.

F SUSPENDS - FALSE CONTENTION: False contention can occur when two different resources are mapped to the same hash class. MVS XES resolves that problem by communicating with the other member involved. False contention requires MVS XCF communication and is therefore dependent on the speed of the communication link used. False contention can sometimes be relieved by enlarging the lock structure in the coupling facility, doubling it, for example. However, you should verify the result of a lock structure change by checking the DB2 PM Statistics report.

G PSET/PART P-LCK NEGOTIATION: A negotiable lock is a lock whose mode can be downgraded, by agreement among contending users, to be compatible to all. A physical lock is a negotiable lock. In comparison with CHANGE REQUESTS (P-LOCKS) B , P-LOCK negotiation is a slow process that involves two or more data sharing members using MVS XES to communicate. The P-LOCK negotiation can involve more than two members; for example, if two members are doing read-only sharing a table space and the third member intends to update this table space, then it needs to drive other members′ page set P-lock exits to downgrade the page set P-locks from S-mode to IS-mode.

H PAGE P-LOCK NEGOTIATION: This field shows that for 394 out of 27573 (A) lock requests for P-locks, page P-lock negotiation is done (1.4%).

PAGE P-LOCK NEGOTIATION / TOTAL COMMITS = 394 / 18191 = 0.0216

This means that less than 2.2% of the transactions on average lead to page P-lock negotiation. If this indicator is too high for your environment, one conclusion could be that too many rows reside on too few pages, and row locking is used. The more likely ones are space map pages and index leaf pages. For index leaf pages, one way to reduce P-lock negotiation is to reduce the number of keys in each index leaf page, that is, leave more space in each index leaf page; or to commit more frequently. Frequent commits also reduces P-lock negotiation for the space map pages. If you have to tolerate a value higher than 0, take care the percentage of affected transactions is small.

Chapter 5 in the *IBM DATABASE 2 for MVS/ESA Version 4 Administration Guide* has more details on lock contention.

8.1.3.1 Evaluation of Locking Activity and Data Sharing Locking

We compare locking activities in comparable environments using DB2 PM Statistics and DB2 PM Accounting reports. All the report extracts provide information for the IRWW 1-way and 2-way measurements.. The Statistics report includes information belonging to the subsystems, including the data sharing members in the group, while the Accounting report includes plan related information. Out of the difference we can define activities of the data sharing member that are not directly accountable to a specific plan. Figure 43 on page 100 is an extract from DB2 PM Accounting report showing information about class 3 suspensions.

Accounting Report Extract

CLASS 3 SUSP.	1-Way AV.TIME	1-Way AV.EVENT	2-Way Member 1 AV.TIME	2-Way Member 1 AV.EVENT	2-Way Member 2 AV.TIME	2-Way Member 2 AV.EVENT
LOCK/LATCH	A 0.003570	2.39	0.006334	1.90	0.007494	2.04
SYNCHRON. I/O	0.285165	11.20	0.223908	9.04	0.232183	9.35
OTHER READ I/O	0.180544	1.14	0.166700	1.17	0.162587	1.14
OTHER WRTE I/O	0.000948	0.01	0.000006	0.00	0.000013	0.00
SER.TASK SWTCH	0.022624	0.54	0.031217	0.54	0.031082	0.54
ARC.LOG(QUIES)	0.000000	0.00	0.000000	0.00	0.000000	0.00
ARC.LOG READ	0.000000	0.00	0.000000	0.00	0.000000	0.00
DRAIN LOCK	0.000000	0.00	0.000000	0.00	0.000000	0.00
CLAIM RELEASE	0.000000	0.00	0.000000	0.00	0.000000	0.00
PAGE LATCH	0.000213	0.00	0.000398	0.00	0.000228	0.00
STORED PROC.	0.000000	0.00	0.000000	0.00	0.000000	0.00
NOTIFY MSGS.	0.000000	0.00	0.000000	0.00	0.000000	0.00
GLOBAL CONT.	0.000000	0.00	0.023198	0.11	0.022649	0.11
TOTAL CLASS 3	T 0.493063	15.28	0.451760	12.76	0.456235	13.18

Figure 43. Accounting Report Extract Showing Class 3 Times. Comparison of 1-way, 2-way member 1, and 2-way member 2 measurements

A CLASS 3 SUSP. LOCK/LATCH: The average number of lock/latch suspensions and the average time of these suspensions. If you compare AVERAGE TIME (for suspensions) with TOTAL CLASS 3 (T), then the suspension time becomes insignificant for our measurements.

Figure 44 on page 101 is an extract from DB2 PM Statistics report showing information about locking activity.

Statistics Report Extract

LOCKING ACTIVITY		1-Way /COMMIT	2-Way Member 1 /COMMIT	2-Way Member 2 /COMMIT
SUSPENSIONS (ALL)	B	0.11	0.19	0.21
SUSPENSIONS (LOCK ONLY)	C	0.00	0.02	0.02
SUSPENSIONS (LATCH ONLY)	D	0.11	0.17	0.19
SUSPENSIONS (OTHER)	E	0.00	0.00	0.00
TIMEOUTS		0.00	0.00	0.00
DEADLOCKS		0.00	0.00	0.00
LOCK REQUESTS	F	6.63	8.17	8.83
UNLOCK REQUESTS	H	1.53	3.01	3.27
QUERY REQUESTS		0.00	0.00	0.00
CHANGE REQUESTS	J	3.80	3.61	3.80
OTHER REQUESTS		0.00	0.00	0.00
LOCK ESCALATION (SHARED)		0.00	0.00	0.00
LOCK ESCALATION (EXCLUSIVE)		0.00	0.00	0.00
DRAIN REQUESTS		0.00	0.00	0.00
DRAIN REQUESTS FAILED		0.00	0.00	0.00
CLAIM REQUESTS		9.10	9.18	9.13
CLAIM REQUESTS FAILED		0.00	0.00	0.00

Figure 44. Statistics Report Extract Showing Locking Activity. Comparison of 1-way, 2-way member 1, and 2-way member 2 measurements

B SUSPENSIONS (ALL): These are average number of suspensions per commit; in this case between 11 and 21 per 100 commits. The increase in suspensions for 2-way data sharing are attributed mostly to latch suspensions. As the total suspension time is only 0.7% of the total accounting class 3 time (3.57 A / 493.06 T = 0.0072), the latch suspensions can be considered insignificant, especially as they are only a small part of that 0.7%.

C SUSPENSIONS (LOCK ONLY): These are the events that are expensive in terms of elapsed time.

D SUSPENSIONS (LATCH ONLY): Latch suspensions are the majority of all suspensions in this case. As they are of very short duration and do not involve a long path length, it does not seem advisable to count them together with lock-only suspensions.

E SUSPENSIONS (OTHER): These should not occur under normal circumstances.

F LOCK REQUESTS: The increase in lock request for the 2-way measurements is attributed mainly to the increase in physical locks. We see some more lock requests in addition to page latch requests. Page physical locks are used to propagate the page latch request to the other data sharing members. In the case of IRWW, considering both CPCs, we see a 28% increase in lock requests. The lock request increase is caused by X-mode physical locks on index updates and space map updates. Lock avoidance is possible for read-only requests, if the isolation is cursor stability.

A higher lock and unlock request count in Statistics compared to Accounting is found to be due to the fact that lock and unlock requests in Statistics include P-lock and P-unlock requests, while the Accounting counters do not.

After migration to DB2 V4 with type 2 index, you will find a significant decrease in locking activity, even considering the comparably moderate increase for data sharing.

H UNLOCK REQUESTS: These are the unlock requests as seen by the member. They include plan-attributable unlocks and unlocks not being attributed to a specific plan, such as physical locks held by the data sharing member.

J CHANGE REQUESTS: Switching from 1-way to 2-way data sharing does not affect the number of change requests. Change requests depend only on the plan and are not influenced by data sharing. These requests are propagated to the coupling facility in an n-way data sharing environment; in a 1-way environment, they are dealt with only by the local IRLM. The change requests are identical in the Statistics and Accounting reports.

Figure 45 is an extract from DB2 PM Accounting report showing information about locking activity.

```
                Accounting Report Extract
                -------------------------

                        1-Way     2-Way     2-Way
                                 Member 1  Member 2
                       -------   --------  --------
LOCKING                AVERAGE   AVERAGE   AVERAGE
--------------         -------   -------   -------
TIMEOUTS                 0.00      0.00      0.00
DEADLOCKS                0.00      0.00      0.00
ESCAL.(SHARED)           0.00      0.00      0.00
ESCAL.(EXCLUS)           0.00      0.00      0.00
MAX LOCKS HELD           5.20      4.97      5.23
LOCK REQUEST    G        6.63      6.74      7.30
UNLOCK REQUEST  I        1.52      1.57      1.71
QUERY REQUEST            0.00      0.00      0.00
CHANGE REQUEST           3.80      3.61      3.80
OTHER REQUEST            0.00      0.00      0.00
LOCK SUSPENS.            0.00      0.02      0.02
LATCH SUSPENS.           0.11      0.17      0.19
OTHER SUSPENS.           0.00      0.00      0.00
TOTAL SUSPENS.  S        0.11      0.19      0.21
```

Figure 45. Accounting Report Extract Showing Locking Activity. Comparison of 1-way, 2-way member 1, and 2-way member 2 measurements

G LOCK REQUEST: Comparison of the number of lock requests in Accounting report with the number of lock requests F in Statistics report shows that the 2-way measurements are lower here. The reason for this is that F captures all lock requests, both for the plans and for the data sharing members, while G captures only those lock requests that can be directly attributed to a specific plan.

A lower lock and unlock request count in Accounting compared to Statistics is found to be due to the fact that lock and unlock requests in Statistics include P-lock and P-unlock requests, while the Accounting counters do not.

I UNLOCK REQUEST: These are the counters for the average unlock requests per plan. The 1-way measurements match with H in Statistics report, while the 2-way measurements are lower.

Unlock requests are fewer than lock requests (G). The number of unlocks reported in Accounting is not the actual number of resources unlocked but the number of calls by DB2 to IRLM. At end of a unit of work, DB2 makes a call to IRLM (UNLOCK ANY) which performs a generic unlock by work unit.

Comparison of both 2-way values for the unlock requests in the Accounting report with the unlock requests (H) in the Statistics report shows that the values in the Accounting report are lower by 1.44 and 1.56 respectively. That coincides with the differences found for F and G , the number of lock requests. The reason is that system held P-locks have to be released singly as they are not tied to a unit of work.

Figure 46 on page 104 is an extract from DB2 PM Statistics report showing information about data sharing locking activity.

```
                    Statistics Report Extract
                    ------------------------

                                  1-Way     2-Way      2-Way
                                           Member 1   Member 2
                                 -------   --------   --------
DATA SHARING LOCKING             /COMMIT   /COMMIT    /COMMIT
---------------------------      -------   -------    -------
LOCK REQUESTS (P-LOCKS)       1    0.00       1.43       1.52
UNLOCK REQUESTS (P-LOCKS)     2    0.00       1.43       1.51
CHANGE REQUESTS (P-LOCKS)     3    0.00       0.00       0.00

SYNCH.XES - LOCK REQUESTS     4    0.04       7.14       7.91
SYNCH.XES - CHANGE REQUESTS   5    0.00       3.60       3.79
SYNCH.XES - UNLOCK REQUESTS   6    0.04       6.62       7.03
ASYNCH.XES - RESOURCES        7    0.00       0.00       0.00

SUSPENDS - IRLM GLOBAL CONT   8    0.00       0.03       0.04
SUSPENDS - XES GLOBAL CONT.   9    0.00       0.00       0.00
SUSPENDS - FALSE CONTENTION   T    0.00       0.06       0.06
INCOMPATIBLE RETAINED LOCK         0.00       0.00       0.00

NOTIFY MESSAGES SENT          U    0.00       0.04       0.00
NOTIFY MESSAGES RECEIVED      V    0.00       0.00       0.04
P-LOCK/NOTIFY EXITS ENGINES         N/A        N/A        N/A
P-LCK/NFY EX.ENGINE UNAVAIL        0.00       0.00       0.00

PSET/PART P-LCK NEGOTIATION   W    0.00       0.00       0.00
PAGE P-LOCK NEGOTIATION       X    0.00       0.02       0.02
OTHER P-LOCK NEGOTIATION      Y    0.00       0.00       0.00
P-LOCK CHANGE DURING NEG.     Z    0.00       0.02       0.02
```

Figure 46. Statistics Report Extract Showing Data Sharing Locking. Comparison of 1-way, 2-way member 1, and 2-way member 2 measurements

1 LOCK REQUESTS (P-LOCKS): These are propagated system-related data sharing locks. Notice about a 17% additional physical locks for the 2-way measurements:

$$\text{Percentage of additional physical locks} = \frac{\text{LOCK REQUESTS (P-LOCKS) 1}}{\text{LOCK REQUESTS F}} * 100$$

$$= \frac{(1.43 + 1.52)}{(8.17 + 8.83)} * 100 = 17\%$$

F is taken from Figure 44 on page 101

2 UNLOCK REQUESTS (P-LOCKS): These are propagated system-related data sharing locks.

3 CHANGE REQUESTS (P-LOCKS): These are the change requests for the P-locks. During the IRWW measurement period there were no changes to the parent P-locks. Change request is not used for child P-locks.

S-mode page P-locks on index leaf pages and data pages (if row locking) are acquired only for RR readers and RI checking. Currently, the page

P-lock downgrade from X-mode to S-mode is done only if S-mode page P-lock is requested by another member when processing a RR or RI request on behalf of a transaction.

4 SYNCH.XES - LOCK REQUESTS: The lock requests include the physical locks from LOCK REQUESTS (P-LOCKS) 1 and plan-related locks, as propagated to the coupling facility lock structure.

5 SYNCH.XES - CHANGE REQUESTS: We find the values almost identical to CHANGE REQUESTS J in Figure 44 on page 101. That means that all changed locks had global scope.

6 SYNCH.XES - UNLOCK REQUESTS: This is the average number of unlocks related to transactions for the data sharing member. Compare with the plan-specific unlock request in the accounting report. The values for UNLOCK REQUESTS I in Figure 45 on page 102 (1.52, 1.57, 1.71) are very different. The reason for the difference is that the DB2 PM Accounting report shows batch unlocks per unit of work, while the DB2 PM Statistics report shows unlocks per resource. We see that for 1-way data sharing, no locks or unlocks for the plan are propagated to the coupling facility. Comparing (6.62, 7.03) unlocks with (1.57, 1.71), we find (5.05, 5.32) more unlocks per commit.

7 ASYNCH.XES - RESOURCES: These are the system-related counters per commit for asynchronous communication with the coupling facility. This counter is incremented for propagation of child locks under IRLM service request block (SRB); for example, member A has an S-lock and member B acquires an IX-lock. In such a case, member A must propagate all child S-locks to the coupling facility.

8 SUSPENDS - IRLM GLOBAL CONT: This is the average number of IRLM suspensions per commit.

9 SUSPENDS - XES GLOBAL CONT.: This is the average number of XES suspensions per commit caused by compatible locks. They lead to suspensions as XES has to find out about it first.

T SUSPENDS - FALSE CONTENTION: This is the average number of XES suspensions per commit.

U NOTIFY MESSAGES SENT: The messages are mainly related to castout. At periodic intervals the GBP structure owner queries the coupling facility GBP structure about the threshold for the GBP castout. If the threshold is reached, the GBP structure owner sends a message to the page set owner to commence castout.

V NOTIFY MESSAGES RECEIVED: The messages are mostly related to castout. It is the page set owner who receives a message from the GBP structure owner if the castout threshold is reached.

W PSET/PART P-LCK NEGOTIATION: This is the average number of negotiations for a page set or partition physical lock.

X PAGE P-LOCK NEGOTIATION: This is the average number of page physical lock negotiations executed by the system. In our case, the value reflects updates done to the space map pages and the index pages.

Y OTHER P-LOCK NEGOTIATION: These negotiations take place for such things as SKCT, SKPT, and castout P-locks.

Z P-LOCK CHANGE DURING NEG.: This is the average number of times per commit that a physical lock change request is issued during physical lock negotiation. It should be identical to E in most cases, as a physical lock negotiation changes the physical lock.

Figure 47 is an extract from DB2 PM Accounting report showing information about data sharing locking activity.

Accounting Report Extract

DATA SHARING		1-Way AVERAGE	2-Way Member 1 AVERAGE	2-Way Member 2 AVERAGE
LOCK REQ - PLOCKS		0.00	1.43	1.52
UNLOCK REQ - PLOCKS		0.00	1.41	1.49
CHANGE REQ - PLOCKS		0.00	0.00	0.00
LOCK REQ - XES	H	0.04	6.95	7.52
UNLOCK REQ - XES	I	0.04	2.31	2.47
CHANGE REQ - XES		0.00	3.60	3.79
SUSPENDS - IRLM		0.00	0.03	0.04
SUSPENDS - XES		0.00	0.00	0.00
SUSPENDS - FALSE		0.00	0.06	0.06
INCOMPATIBLE LOCKS		0.00	0.00	0.00
NOTIFY MSGS SENT		0.00	0.00	0.00

Figure 47. Accounting Report Extract Showing Data Sharing Locking. Comparison of 1-way, 2-way member 1, and 2-way member 2 measurements

H LOCK REQ - XES: This is the average number of locks propagated on behalf of a plan.

I UNLOCK REQ - XES: These are the plan-related unlocks, as seen by XES.

8.1.4 Using RMF Report

RMF reports can be obtained online using TSO, and can be generated off line as batch jobs. The relevant RMF Reports can be produced only by the RMF Post Processor from extracted SMF 7x records. RMF Monitor I and Monitor III must be active with RECORD option.

The batch listing obtained during the IRWW measurements from the *RMF Coupling Facility Activity* is reported in Appendix F.1, “RMF – Lock Structure Activities” on page 171. Figure 48 on page 107 shows an extract from that report.

```
STRUCTURE NAME = DSNDB0G_LOCK1     TYPE = LOCK
          # REQ    ------------ REQUESTS -----------
SYSTEM    TOTAL          #   % OF  -SERV TIME(MIC)-    REQUEST
NAME      AVG/SEC        REQ  ALL    AVG STD_DEV    CONTENTIONS

STLABC1     634K A SYNC  634K 100% D 223.4   51.4   E #REQ        625K
          521.0  B ASYNC   0  0.0%    0.0    0.0    F DELAYED     4538
                 C CHNGD   0  0.0%  INCLUDED IN ASYN G  -CONT      4535
                                                    H  -FALSE CONT 1825
```

Figure 48. RMF Coupling Facility Activity Report for the Lock Structure. For Coupling Facility Model 9674/R61. Extracted from Appendix F.1, "RMF – Lock Structure Activities" on page 171

A SYNC: Counts all requests to the coupling facility that operate in synchronous mode, which means, all requests that can be handled by the coupling facility without communicating with other members. As a result of a synchronous lock request, a conflict may be detected. There is a need for MVS XCF message passing to resolve the conflict.

B ASYNCH: DB2 does not make asynchronous requests to the coupling facility. This value is always 0. Note that RMF uses the word "synchronous" differently from DB2.

C CHNGD: Asynchronous requests come from a variety of sources, as some requests are issued as asynchronous commands. If MVS determines that a synchronous request is delayed (probably because the subchannels are busy), it changes the request to asynchronous. It is these changed requests that are counted in the CHNGD field of the RMF Coupling Facility Structure Activity report.

Based on our experience, we recommend that the total number of requests CHNGD be less than 10% of the total requests for a given structure.

D SERVICE TIME AVG: That is the average service time of the coupling facility in microseconds.

E #REQ: The total number of synchronous requests without contention

F #REQ DELAYED: A subset of the number of locks requested. It represents the number of requests that cannot be immediately completed. These include any requests that require additional processing to complete. This is the total number of requests being delayed, including delays because of an unavailable coupling facility link (for example, busy).

G #DELAYED CONT: A subset of the number of lock request delays (# REQ DELAYED). It represents the number of lock requests that are in fact delayed because of contention on the lock (whether real or false). Except for some unusual exceptions, DELAYED CONT is the same as #REQ DELAYED. In comparison with E (#REQ) this is a small number (G / E = 0.7%) that does not affect the IRWW workload. A tolerable upper limit for this number would be about 2% of E (#REQ)

```
625000 * 2 / 100 = 12500
```

If it is less than 1%, it is good. If it is 2%, more hash entries are desirable, if the reason is not real contention.

H DELAYED FALSE CONT: MVS XES sorts out the false contentions. You can calculate the percentage of false contention causing the delay:

```
False contention delay = H / E
                       = 1825 / 625000 = 0.3%
```

As H is included in G (#REQ CONT), you can calculate the percentage of real contention causing the delay:

```
Real contention delay = ( G - H ) / E
                      = (4535 - 1825) / 625000 = 0.4%
```

The XES definition of false contention presented in RMF differs from the DB2 definition represented in DB2 PM. The difference can occur due to contention being resolved before XCF global management signaling. This difference usually results in slightly lower RMF false contention counts than DB2 PM false contention counts. In determining the adequacy of the lock table size, the RMF counts are better.

Tuning information: False contention can be relieved by increasing the size of the lock table.

Coupling Facility I/O Queuing A look into the XCF activities helps you to determine whether your MVS XCF is properly set up. Essentially, look for the average queue length A , as can be seen in Figure 49 extracted from the RMF XCF activities report.

```
                  OUTBOUND FROM STLABC1
---------------------------------------------------------
        T FROM/TO
 TO     Y DEVICE, OR      TRANSPORT        REQ     AVG Q
 SYSTEM P STRUCTURE       CLASS            OUT     LNGTH

 STLABC2 C 0B14 TO 0B14   DEFAULT        2,456      0.00
         C 0B15 TO 0B15   DEFAULT        1,387      0.00
 STLABC3 C 0B24 TO 0B24   DEFAULT        1,448      0.00
         C 0B25 TO 0B25   DEFAULT        1,610   A  0.02
```

Figure 49. RMF Coupling Facility Report for Outbound Messages. For message traffic between members. Extracted from Appendix F.2, "RMF – XCF Activities" on page 172

A AVG Q LNGTH: Average queue length should be well below the DASD queue length shown by B (AVG Q LNGTH) in Figure 50, "I/O QUEUING ACTIVITY REPORT"

```
                   I/O  Q U E U I N G   A C T I V I T Y

          SYSTEM ID 961              DATE 09/09/95
          RPT VERSION 5.2.0          TIME 08.41.39

IOP   ACTIVITY RATE  AVG Q LNGTH     IODF = C1   CR-DATE: 07


 00      611.928      B 0.19
```

Figure 50. RMF Coupling Facility Report for I/O Queuing Activity. Extracted from Appendix F.3, "RMF – I/O Queuing Activity" on page 176

B AVG Q LNGTH: This is the average DASD queue length for all the DASD.

For the message length, see Appendix F.2, "RMF – XCF Activities" on page 172. We used 956 bytes.

8.1.5 Allocating a Lock Structure

In determining how much to allocate for the lock structure, two references are helpful:

- *SETI, Document OZSQ662530, May 1994*, available under DIAL IBM.
- *IBM DATABASE 2 for MVS/ESA Version 4 Data Sharing: Planning and Administration* , Chapter 3, "Panning and Installing DB2 Data Sharing, Estimating Storage"

We analyze the use of the lock structure a bit further. A lock structure essentially consists of the following parts:

- One lock table for locks whose entries are in powers of 2 bytes and can be expressed in binary as b′10000...′.
- A lock table for long-duration locks as part of the same physical lock table.
- A modified resource list with an entry length of 64 bytes for the DB2 resource names and additional bytes for the coupling facility′s use of the lock structure.

8.1.5.1 Analysis of the DB2 DISPLAY GROUP Command

To analyze the Figure 41 on page 95 further, we show an extract in Figure 51.

```
LOCK1 STRUCTURE SIZE:  A 64000 KB,             LOCK1 IN USE:     1 %
NUMBER  LOCK ENTRIES:  B 16777216, LOCK ENTRIES  IN USE:      10240
NUMBER  LIST ENTRIES:  C   228245, LIST ENTRIES  IN USE:       2593
```

Figure 51. DB2 DISPLAY GROUP, Lock Structure Extract. Extracted from Figure 41 on page 95

We show the sizes in hexadecimal because the length of a lock table in the lock structure always takes the form 2**n.

A Our lock structure is 64 MB.

B The number of lock entries is 16 777 216 (16 MB). The length of a lock table entry is 2 bytes, resulting in a total size of the lock table of 32 MB. We sized the coupling facility such that it would not give more than 0.5% false contentions.

C The number of list entries (modified resource list) is 228 245. They have to fit in 32 MB.

```
64 MB ( A ) - 32 MB ( B ) = 32 MB
```

Note: You should be aware that sizing the lock table to reduce contention is only relevant if you have tuned your databases and applications to minimize contention. Sizing the lock table alone does not help to reduce the global resource contention.

8.1.5.2 Analysis of the RMF Coupling Facility Report

To confirm the values found in Figure 51 on page 109, we show in Figure 52 an extract of the RMF report "Lock Structure Activities."

```
     COUPLING  FACILITY  USAGE  SUMMARY
-------------------------------------------------
           STRUCTURE SUMMARY
-------------------------------------------------
                                LIST     LOCK
        STRUCTURE       ALLOC  ENTRIES  ENTRIES
 TYPE   NAME            SIZE   TOT/CUR  TOT/CUR
 LOCK   DSNDB0G_LOCK1    64M    C 228K   B  17M
```

Figure 52. RMF Coupling Facility Activity Report Extract. Extracted from Appendix F.1, "RMF – Lock Structure Activities" on page 171

B Confirms B in Figure 51 on page 109.

C Confirms C in Figure 51 on page 109.

With this information we can build Table 17 for varying lock table sizes and their number of entries.

Table 17. Number of 2-Byte Entries in Lock Structures. Number of entries change for different sizes of lock structure

Structure Size	1 **Lock Entries**	2 **Modified Resource List Entries**
8 MB	2 M	25 K
16 MB	4 M	50 K
32 MB	8 M	100 K
64 MB	16 M	200 K
128 MB	32 M	400 K

The values in Table 17 are rounded down; they are presented as a rough guide. As we recommend using 1% of the entries available, the rounding does not influence your choice.

You can calculate the number of locks you require by this formula:

1 Lock Entries = A * B * C * D

A
Maximum locks per transaction or commit (MAX LOCKS HELD field in the LOCKING block in DB2 PM Accounting).

B
Duration of transaction in seconds (ACCOUNTING CLASS 1 ELAPSED TIME field in DB2 PM Accounting).

C
Number of transactions per second (TOTAL COMMITS in HIGHLIGHTS block in DB2 PM Statistics).

D
Workload specific contention factor.

For example, in the case of IRWW 2WP measurements, we have:

A = 46 for the first member and 47 for the second member.

B = 2.101547 seconds for the first member and 2.084017 seconds for the second member.

C = (18056 + 18191) / 600 for the average value in the statistics interval of 10 minutes.

D = For the specific IRWW workload, 16K. That is, we have used (16 x 1024) or
16384 times more entries than we have lock entries in use to achieve less than 1% contention.

Substituting these values in the formula, 1 Lock Entries is calculated as 9.8 M. This value is too low for one obvious reason. We have assumed that all transactions using the same number of locks per minute as long transactions. When you look into Table 17 on page 110 with the value of 1 then you see that for the number of locks in the lock table we could have used a 64 MB lock structure in the coupling facility for this particular run.

You can have a first estimate of the number of modified resource list entries you require by this formula:

2 Modified Resource List Entries = (A * B) + (C * D)

A
Number of concurrent threads.

B
Maximum locks per transaction or commit (MAX LOCKS HELD field in the LOCKING block in DB2 PM Accounting).

C
Sum of all buffers in all buffer pools for all members that can become GBP-dependent (BUFFERS ALLOCATED - VPOOL field in DB2 PM Statistics).

D
Probability of a buffer P-lock at a given point in time.

For example, in the case of IRWW 2WP measurements, we have:

A = 66 for the first member and 67 for the second member.
Each IMS TM has 66 regions to process transactions, for a total of 132 concurrent threads.

B = 46 for the first member and 47 for the second member.

C = 147 000 (See note below)

D = 0.5 (See note below)

Note:

As our XCF DISPLAY in Figure 39 on page 93 shows, we allow every local buffer pool except BP2 to be GBP-dependent. Summing up all virtual buffer pool sizes used and not counting BP2, we get 73100 buffers for each member. As each member has 73100 buffers that can become GBP-dependent, the total for the 2WP measurement is 147000 buffers.

For the probability of a buffer P-lock for a given point in time. We use 0.5, knowing that this is on the high side. For a first estimate, we think being on the high side with the size of the lock structure is less damaging than being on the low side. The result is not an exact number. It requires tuning.

Substituting these values in the formula, 2 Modified Resource List Entries is calculated as 79704. When you look into Table 17 on page 110 with the resulting value, then you see that, for the number of modified resource list entries in the lock table we could have used a 32 MB lock structure in the coupling facility for this particular run.

8.2 Analyzing the Use of the Group Buffer Pool

Many tools allow you to display information about the GBP, DB2 PM and RMF among them. However, there are methods that are fast and can be used to great advantage when connected to automated operating procedures. Such a connection allows you to trigger reports, inform system programming or alert operators at remote locations.

8.2.1 Using the XCF DISPLAY Command

Use the XCF display command D XCF,STR to get the names of the structures for the group buffer pools in the coupling facility. Use the XCF display command D XCF,STR,STRNAME=DSNDB0G_GBP8 to get the details of a particular structure, in this case DSNDB0G_GBP8. Figure 53 shows the details of the structure DSNDB0G_GBP8.

```
IXC360I  10.22.34  DISPLAY XCF 179
STRNAME: DSNDB0G_GBP8
 STATUS: ALLOCATED
  POLICY SIZE    : 133000 K       A
  POLICY INITSIZE: N/A
  REBUILD PERCENT: N/A
  PREFERENCE LIST: LF02
  EXCLUSION LIST IS EMPTY

 ACTIVE STRUCTURE
 ----------------
  ALLOCATION TIME: 09/09/95 10:04:16
  CFNAME         : LF02
  COUPLING FACILITY: 009674.IBM.00.000000040029
                     PARTITION: 1   CPCID: 00
  ACTUAL SIZE    : 133120 K        B
  STORAGE INCREMENT SIZE: 256 K    C
  VERSION        : ABA5C3AF 94CA1803
  DISPOSITION    : DELETE
  ACCESS TIME    : 0
  MAX CONNECTIONS: 32
  # CONNECTIONS  : 2               D

 CONNECTION NAME  ID VERSION  SYSNAME  JOBNAME  ASID STATE
 ---------------- -- -------- -------- -------- ---- ------  E
 DB2_DB1G         01 000100B0 STLABC1  DB1GDBM1 001D ACTIVE
 DB2_DB2G         02 000200AB STLABC2  DB2GDBM1 019A ACTIVE
```

Figure 53. XCF Display Showing Cache Size for Group Buffer Pool GBP8

A The structure size as defined in the CFRM policy.

B The actual allocated size, in multiples of 256 KB.

C The increment value by which the storage size can be changed.

D The active number of connections to this structure.

E The details of the active connections to this structure.

See *Setting Up a Sysplex* for more details.

8.2.2 Using the DB2 DISPLAY GROUPBUFFERPOOL Command

The DB2 DISPLAY GROUPBUFFERPOOL command provides greater detail for the GBP than the MVS XCF command. When the DB2 DISPLAY GROUPBUFFERPOOL command is used with the MDETAIL (for member detail) or with the GDETAIL (for group detail) option, the output consists of two parts, one header with static information, and detail statistics with dynamic information.

8.2.2.1 DB2 DISPLAY GROUPBUFFERPOOL Output Member Detail Static Information

We look at the member detail static information in Figure 54.

```
*DB1G DIS GBPOOL(GBP8) MDETAIL(*)
DSNB750I *DB1G DISPLAY FOR GROUP BUFFER POOL GBP8 FOLLOWS
DSNB755I *DB1G DB2 GROUP BUFFER POOL STATUS
             CONNECTED                                = YES
             CURRENT DIRECTORY TO DATA RATIO          = 3.1        A
             PENDING DIRECTORY TO DATA RATIO          = 3.1
DSNB756I *DB1G   CLASS CASTOUT THRESHOLD              = 10%        B
             GROUP BUFFER POOL CASTOUT THRESHOLD      = 50%        C
             GROUP BUFFER POOL CHECKPOINT INTERVAL    = 10 MINUTES D
             RECOVERY STATUS                          = NORMAL
DSNB757I *DB1G MVS CFRM POLICY STATUS FOR DSNDB0G_GBP8 = NORMAL
             MAX SIZE INDICATED IN POLICY             = 133000 KB  E
             ALLOCATED                                = YES
DSNB758I *DB1G     ALLOCATED SIZE                     = 133120 KB  F
              VOLATILITY STATUS                       = VOLATILE   G
DSNB759I *DB1G     NUMBER OF DIRECTORY ENTRIES        = 89461      H
              NUMBER OF DATA PAGES                    = 28854      I
              NUMBER OF CONNECTIONS                   = 2
DSNB772I *DB1G CUMULATIVE MEMBER DETAIL STATISTICS SINCE 08:41:13
SEP  9, 1995
```

Figure 54. DISPLAY GROUPBUFFERPOOL Output Member Detail Static Information

A CURRENT DIRECTORY TO DATA RATIO is the number of directory entries per data page in the GBP.

```
                                     NUMBER OF DIRECTORY ENTRIES ( H )
CURRENT DIRECTORY TO DATA RATIO = ------------------------------------
                                        NUMBER OF DATA PAGES ( I )

                                     89461
                                = --------- = 3.1
                                     28854
```

The value shown in A is set using the

```
-ALTER GROUPBUFFERPOOL RATIO(3.1)
```

command.

B CLASS CASTOUT THRESHOLD limits the space that the changed pages for a single class of data sets can occupy in the GBP. The value shown in B is set using the

-ALTER GROUPBUFFERPOOL CLASST(10)

command.

Tuning information: Lowering the class castout threshold forces a more frequent castout by the data sharing member that made this page set GBP-dependent. This can be a means of balancing castout over multiple data sharing members. When you enlarge the GBP to keep clean pages longer, remember to lower this threshold if you do not want to change the castout characteristics.

C GROUP BUFFER POOL CASTOUT THRESHOLD is the percentage of changed pages that can be kept in the GBP without castout. When the number of changed pages is higher than this threshold, a GBP castout is triggered. The value shown in C is set using the

-ALTER GROUPBUFFERPOOL GBPOOLT(50)

command.

D The GROUP BUFFER POOL CHECKPOINT INTERVAL is used by the GBP structure castout owner to write a GBP checkpoint. The checkpoint triggers writing of changed pages from the coupling facility to the page set. The GBP checkpoint is triggered at specified time intervals by the castout structure owner. The default checkpoint interval is 8 minutes. A broadcast is made through XCF to all data set castout owners to trigger the castout of updated pages to DASD. For more information on the GBP checkpoint see *IBM DATABASE 2 for MVS/ESA Version 4 Data Sharing: Planning and Administration* The value shown in D is set using the

-ALTER GROUPBUFFERPOOL GBPCHKPT(10)

command.

Tuning information: You need to balance the GBP checkpoint frequency against the overhead and reduced write efficiency. Reducing GBP checkpoint interval helps increase data availability. Data is available sooner after -START DATABASE command and DB2 restart after failure. We recommend using multiple buffer pools with isolation of most frequently accessed data. We suggest a minimum value for the checkpoint interval of 1 minute for such data. A change in the GBP checkpoint interval becomes effective only after the next checkpoint becomes due.

E MAX SIZE INDICATED IN POLICY is the maximum size of the GBP structure (the parameter SIZE in the CFRM policy). If you are on MVS/ESA SP Version 5 Release 2, then there are two sizes in the CFRM policy:

SIZE - maximum size
INITSIZE - initial allocation size.

You can use the SETXCF ALTER command to dynamically expand the size up to the maximum, without having to change CFRM policies and deallocate the structure, as you would have to do on MVS/ESA SP Version 5 Release 1 (You can also dynamically contract the size with the SETXCF ALTER command.) To change the size in MVS/ESA SP Version 5 Release 1, you first change the CFRM policy, then activate the CFRM policy in MVS, and then have DB2 allocate the size. There is no single command to do this. See *IBM DATABASE 2 for MVS/ESA Version 4 Data Sharing: Planning and Administration* Chapter 6, “Tuning Group Buffer Pools” for detailed information on how to change the size of the group buffer pool.

F ALLOCATED SIZE is the size allocated for the GBP rounded up to 256 KB boundaries. DB2 uses the size given in the CFRM policy and allocates the structure at the time of first use. For details on CFRM policy specification, see *MVS/ESA Setting Up a Sysplex*

G VOLATILITY STATUS does not affect normal operations, and does not affect restart processing. We recommend nonvolatile coupling facility so that the data in the coupling facility is preserved across a power outage. That avoids the need to do GBP recovery.

H The NUMBER OF DIRECTORY ENTRIES must be higher than the number of data entries, as stated in A .

I The NUMBER OF DATA PAGES is the number of data pages in the GBP. Change your CFRM policy if a different value is needed.

8.2.2.2 DB2 DISPLAY GROUPBUFFERPOOL Output Group Detail Static Information

The group detail static information is identical to that of the member detail static information as in Figure 54 on page 113.

8.2.2.3 Display GROUPBUFFERPOOL Output Group Detail Dynamic Information

We look at the group detail dynamic information in Figure 55.

```
 DSNB782I *DB1G INCREMENTAL GROUP DETAIL STATISTICS SINCE 08:41:11
  SEP 9, 1995
1DSNB784I *DB1G GROUP DETAIL STATISTICS
             READS
               DATA RETURNED                        = 36855   A
 DSNB785I *DB1G     DATA NOT RETURNED
                 DIRECTORY ENTRY EXISTED            = 2493    B
                 DIRECTORY ENTRY CREATED            = 69993   C
                 DIRECTORY ENTRY NOT CREATED        = 141, 0  D
 DSNB786I *DB1G   WRITES
               CHANGED PAGES                        = 93795   E
               CLEAN PAGES                          = 0       F
               FAILED DUE TO LACK OF STORAGE        = 0       G
             CHANGED PAGES SNAPSHOT VALUE           = 13325   H
 DSNB787I *DB1G   RECLAIMS
               FOR DIRECTORY ENTRIES                = 0       I
               FOR DATA ENTRIES                     = 75254   J
             CASTOUTS                               = 81933   K
 DSNB788I *DB1G   CROSS INVALIDATIONS
               DUE TO DIRECTORY RECLAIMS            = 0       L
               DUE TO WRITES                        = 10896   M
 DSNB790I *DB1G DISPLAY FOR GROUP BUFFER POOL GBP8 IS COMPLETE
```

Figure 55. DISPLAY GROUPBUFFERPOOL Output Group Detail Dynamic Information

The following fields are for READS:

A DATA RETURNED: The number of reads in which data is returned from the GBP.

B DATA NOT RETURNED DIRECTORY ENTRY EXISTED: A data page is registered in the GBP because it belongs to a GBP-dependent page set. However, the page itself is not in the GBP. The most likely reason is that the page is read by one or more members and not updated. The member reads the page from DASD.

```
                             DATA NOT RETURNED DIRECTORY ENTRY EXISTED ( B )
Missing shared page rate = -------------------------------------------- * 100
                                          DATA RETURNED ( A )
                              2493
                          = --------- * 100 = 7%
                              36855
```

Tuning information: If the missing-shared-page-rate percentage stays high for a page set, consider using GBPCACHE ALL for that particular page set.

C DATA NOT RETURNED DIRECTORY ENTRY CREATED: Neither the page nor the directory entry could be found in the coupling facility. Still, the requesting member asked for that particular page because the page set is GBP-dependent. CF creates the entry, registers the requesting member for that page and returns “not found.” For GBP-dependent page sets, this is generally a normal situation. This number represents the number of pages read in a GBP-dependent page set that were not read previously by another member. The pages may have been previously read or updated, but the directory entry does not exist. In that case the GBP would be considered too small.

D DATA NOT RETURNED DIRECTORY ENTRY NOT CREATED: Neither page nor directory entry is in the coupling facility. No directory entry is created because no other member in the group has R/W interest in the table space or partition. GBP-dependence still exists. The ″directory entry not created″ is a normal condition; we put the feature in to avoid using directory resources when they′re not necessary.

Two values separated by comma are shown for this field. The second value shows that the GBP directory does not have space and so no directory entry can be created. This condition is serious.

Tuning information:

The lack of GBP directory space can cause the lack of directory entries. When directory entries are reclaimed to handle new work, cross-invalidation must occur for all members that have those pages in their buffer pools, even when the data has not actually changed. Check the directory-to-data-page ratio. You should either increase the directory space by changing the current directory-to-data ratio (see A in Figure 54 on page 113) or reduce the data sharing level by moving some application to another time or to another data sharing member. Another option is to increase the size of the GBP increasing the residence time and the chance to find the page.

The following fields are for WRITES

E CHANGED PAGES: These writes to the coupling facility are pages changed at least once and written because of the force-at-commit policy, or asynchronously triggered when 50% of VDWQT, the data set level write threshold, is reached.

F CLEAN PAGES: These writes to the coupling facility represent unchanged pages. Only the option GBPCACHE ALL causes clean pages to be written to the coupling facility. When a member reads data from a table, and the table space is defined with GBPCACHE ALL, any page read is written to the coupling facility. This is well suited to transaction workloads with intersystem interest having a high R/W ratio, sufficient coupling facility resources, and capable of achieving a high read hit ratio. See also B .

G FAILED DUE TO LACK OF STORAGE: DB2 tries to write data pages for that GBP but cannot complete the operation because the GBP does not have enough storage. This is a condition that needs attention. On write failure, the affected DB2 member waits for only 3 seconds after triggering castout, retries up to four times, and then the page goes on the logical page list (LPL) requiring recovery. The reason is that castout processing cannot keep up with force-at-commit policy and writes triggered asynchronously when 50% of VDWQT, the data set level write threshold, is reached.

Tuning information: To alleviate the problem, you can take some of the following actions:

- Increase the size of the GBP.
- Reduce the ratio of directory entries to data entries.
- Reduce the GBP checkpoint interval, GBP threshold, and the GBP castout class threshold.
- In extreme situations, set the GBP castout class threshold to zero to trigger continuous castout processing. Every 32 changed pages created trigger castout.
- Reduce the number of pages kept in the GBP. To achieve this, these are possible actions:
  - COMMIT more frequently.
  - Increase the checkpoint frequency for the GBP. See also D in Figure 54 on page 113.
  - Reduce number of updates and inserts between commits.

H WRITES - CHANGED PAGES SNAPSHOT VALUE: This is the counter for changed pages in the GBP. It is a dynamic value, not affected by an incremental or cumulative display. If this value is constant for a longer period in your peak load time, then you have achieved a balance of castout interval and arrival of updated pages.

I RECLAIMS FOR DIRECTORY ENTRIES: Reclaims of directory entries may or may not cause cross invalidation. Reclaims of directory entries for clean pages with sending of XI signals, hurts only if the locally cached copy is to be referenced. Only directory entries with no associated data or with associate data that is clean can be reclaimed.

Tuning information: If the value for I is too high, you should increase the GBP size and with it the directory size (or at least change the current directory-to-data ratio).

See also section 8.2.3, “Using DB2 PM Statistics Report” on page 122.

J RECLAIMS FOR DATA ENTRIES: Only clean pages can be reclaimed. This is done on a least recently used (LRU) basis. Changed pages marked clean after castout and still available for reads, and also available for reclamation. The main point here is page residence time. If the GBP is sized too small, pages are reclaimed too quickly. If a page is changed by one member, the change cross invalidates the copy of that page in the local buffer pools of all other members. A reread of that page fails if the page is reclaimed by the coupling facility.

Tuning information: Try to achieve 90%-95% read hit ratio for GBP reads caused by XI.

See also section 8.2.3, "Using DB2 PM Statistics Report" on page 122.

K CASTOUTS: This is the number of pages cast out since the last display.

L CROSS INVALIDATIONS DUE TO DIRECTORY RECLAIMS: This field shows how many cross invalidations occurred because of the reclaims shown by the field RECLAIMS FOR DIRECTORY ENTRIES I

Tuning information: To reduce the number of invalidations, you can do the following:

- Increase the size of the GBP directory.

  See *IBM DATABASE 2 for MVS/ESA Version 4 Data Sharing: Planning and Administration* Chapter 6, "Tuning Group Buffer Pools" for details.
- Avoid updating a page set while an application is scanning a large table space, or move updating applications to the member where the applications doing frequent retrieval of the same data reside.
- Use smaller virtual buffer pools that register fewer pages in the GBP.
- Use CLOSE YES for some large, rarely used table spacees. This may free registered pages and their directory entries for the data sharing member that closes the table space. For CLOSE YES page sets, we initiate a physical close for the object in data sharing if the object is read-only and GBP-dependent (cached state = IS), and the object has not been accessed over PCLOSEN checkpoints or in PCLOSET minutes. By expediting the physical close, we remove the inter-DB2 interest sooner, which also reduces the load on the GBP directory.

  Getting out of GBP-dependence requires a castout to be triggered and all cross-system invalidated pages for that buffer pool in all members to be marked as "needing to be refreshed."

M CROSS INVALIDATIONS DUE TO WRITES: Number of cross-system invalidation requests sent out by the coupling facility for a changed page written to the coupling facility that was registered by other members of the data sharing group.

```
   CROSS INVALIDATION DUE TO WRITES  M           10896
-------------------------------------- * 100 = ------- * 100 = 11%
         WRITES CHANGED PAGES  E                 93795
```

gives an indication of the percentage of updates affecting other members.

8.2.2.4 Display GROUPBUFFERPOOL Output Member Detail Dynamic Information

Member detail display can help spotting problems and irregularities. It gives a good overview of GBP activities with the coupling facility. On the following pages, we go through each field of the member detail report as presented by the display command and present you with some tuning information.

```
DSNB771I *DB1G INCREMENTAL MEMBER DETAIL STATISTICS SINCE 08:41:13
SEP  9, 1995
DSNB773I *DB1G MEMBER DETAIL STATISTICS
              SYNCHRONOUS READS
                DUE TO BUFFER INVALIDATION
                  DATA RETURNED                         = 2492     A
                  DATA NOT RETURNED                     = 0, 0     B
DSNB774I *DB1G      DUE TO DATA PAGE NOT IN BUFFER POOL
                  DATA RETURNED                         = 3104     C
                  DATA NOT RETURNED                     = 36478, 0 D
DSNB775I *DB1G    PREFETCH READS
               REGISTER NAME LIST NOT AVAILABLE
                DATA RETURNED                           = 0        E
                DATA NOT RETURNED                       = 0, 0     F
DSNB789I *DB1G     REGISTER PAGE LIST                   = 4794     G
                RETRIEVE CHANGED PAGES                  = 5329     H
                RETRIEVE CLEAN PAGES                    = 7524     I
           FAILED READS DUE TO LACK OF STORAGE          = 0        J
DSNB776I *DB1G    SYNCHRONOUS WRITES
                CHANGED PAGES                           = 47042    K
                CLEAN PAGES                             = 0        L
DSNB777I *DB1G    ASYNCHRONOUS WRITES
                CHANGED PAGES                           = 1        M
                CLEAN PAGES                             = 0        N
              FAILED WRITES DUE TO LACK OF STORAGE      = 0        O
DSNB778I *DB1G    CASTOUT THRESHOLDS DETECTED
                FOR CLASSES                             = 0        P
                FOR GROUP BUFFER POOL                   = 22       Q
              CASTOUTS                                  = 49002    R
DSNB779I *DB1G    ENGINES NOT AVAILABLE
                FOR CASTOUT                             = 0        S
                FOR WRITING                             = 0        T
              OTHER INTERACTIONS                        = 19454    U
DSNB790I *DB1G DISPLAY FOR GROUP BUFFER POOL GBP8 IS COMPLETE
DSN9022I *DB1G DSNB1CMD ¢-DIS GBPOOL¢ NORMAL COMPLETION
```

Figure 56. DISPLAY GROUPBUFFERPOOL Output Member Detail Dynamic Information

A request to the coupling facility is “synchronous” when it is executed under the user’s execution unit; the user’s address space is in control. A request to the coupling facility is “asynchronous” when it is executed under DB2 V4 system execution unit, for example during a list prefetch.

Note: A “synchronous” coupling facility request is an IXL... request executed without interrupt on the calling side and returning directly to the caller, in contrast to a supervisor call (SVC) that causes an interrupt.

A SYNCHRONOUS READS DUE TO BUFFER INVALIDATION, DATA RETURNED: A page that is cross invalidated in the virtual buffer pool is reread from the GBP. The page may have been cross invalidated when changed by another member. The changed page is requested and transferred to the requester.

B SYNCHRONOUS READS DUE TO BUFFER INVALIDATION, DATA NOT RETURNED: Cross-system invalidation occurs when a page is updated for example, or when a directory entry is reclaimed from the GBP. When the page is committed, force-at-commit policy writes the page to the GBP. The requesting member has a page that is cross invalidated. When trying to obtain the page from the GBP, the member finds that the page is not there anymore. This can happen when castout is complete but the GBP needs more pages than are free. Pages are then reclaimed. This field shows two values separated by comma. If the first value is not equal to zero, it would suggest that when a buffer is read from the coupling facility because of cross-system invalidation and the page set is GBP-dependent, then a directory entry is created in the GBP. The second value shows the number of times the GBP read is issued and the directory entry is not created because there is no R/W interest.

Tuning information:

Normally, if the buffer is cross invalidated, the buffer is refreshed from the GBP. If there are lots of GBP misses when the invalidated buffer is refreshed from the GBP, it may be that the average residence time of a cached page in the GBP is too short (increase the number of GBP data pages), or that the directory entry space is not large enough to handle the page registration and caching load (increase the size of the directory entry space), or both (increase both the GBP data pages and directory entries).

C SYNCHRONOUS READS DUE TO DATA PAGE NOT IN BUFFER POOL, DATA RETURNED: The requesting member needs a page from a table space that is GBP-dependent or has GBPCACHE ALL defined. To get that page, the GBP is checked before the page set on DASD. The requested page is returned to the requesting member.

D SYNCHRONOUS READS DUE TO DATA PAGE NOT IN BUFFER POOL, DATA NOT RETURNED: The requesting member needs a page from a table space that is GBP-dependent. To get that page, the GBP is checked before the page set on DASD. The requested page is not returned to the requesting member and is subsequently read from DASD. This field shows two values separated by comma. The second value shows the number of times the GBP read is issued and the directory entry is not created because there is no R/W interest from other members.

E PREFETCH READS, DATA RETURNED: This is the number of data pages found and returned for the prefetch read from the coupling facility with a register page list not available.

F PREFETCH READS, DATA NOT RETURNED: This is the number of data pages not found nor returned for the prefetch read with a register page list from the coupling facility.

G REGISTER PAGE LIST: A register page list is identical to a register name list. The register page list is used on CFLEVEL=2, MVS/ESA SP Version 5 Release 2, and appropriate maintenance.

A register page list has two functions:

- It holds a list of data pages from one table space.
- It is used to register the pages in the list for the requesting data sharing member in the GBP directory.

The register page list is used by prefetch. As it is the normal case for sequential access, we would expect many GBP misses. The register page list operation allows efficient registration of pages.

H REGISTER PAGE LIST, RETRIEVE CHANGED PAGES: This is a count of the changed pages retrieved. The pages are found in the GBP and returned to the requesting member. These are usually pages written out with force-at-commit policy but not yet castout to the page set.

I REGISTER PAGE LIST, RETRIEVE CLEAN PAGES: These are pages in the GBP that are clean, meaning already written out to DASD or unchanged since the last read in.

J REGISTER PAGE LIST, FAILED READS DUE TO LACK OF STORAGE: The register page list request, in most cases, registers additional page for the requesting member. If GBP directory space is not sufficient, the register page list fails. The prefetch stops and the pages are then read synchronously when required.

Tuning information:

This field should ideally show a zero value. If not, the DB2 PM Accounting report ordered by plan name, provides information about the affected plans. See 8.2.3.1, "Using DB2 PM Report for Buffer Pool Statistics" on page 125 for further analysis.

K SYNCHRONOUS WRITES, CHANGED PAGES: These are changed and committed pages written out to the coupling facility at force-at-commit time from this member. See *IBM DATABASE 2 for MVS/ESA Version 4 Data Sharing: Planning and Administration* Chapter 6, "Tuning Group Buffer Pools," for details on how pages are written to the GBP.

L SYNCHRONOUS WRITES, CLEAN PAGES: These are unchanged pages written to the coupling facility for which you specify GBPCACHE ALL for a table space using this GBP. The pages are written to the coupling facility even if the page set is not GBP-dependent.

M ASYNCHRONOUS WRITES, CHANGED PAGES These are pages written to the coupling facility by a data sharing member under the system task of DB2. This happens for all system-initiated activities (SRBs) that write to the coupling facility, as for any write from virtual buffer pool in case a write threshold is reached.

Tuning information:

Write threshold can be used to avoid writing large commits. Commits may write to the coupling facility in burst mode. The longer the commit interval, the bigger a commit write burst becomes. Unwieldy writes can be circumvented by setting the horizontal or vertical deferred write threshold very low, thus forcing a steady stream of pages to the coupling facility.

N ASYNCHRONOUS WRITES, CLEAN PAGES: Unchanged pages are written out to the coupling facility only for GBPCACHE ALL. Ensure that your coupling facility has enough storage and capacity on the coupling links to sustain the data flow.

O FAILED WRITES DUE TO LACK OF STORAGE: A synchronous write operation triggered by the data sharing member fails when the GBP does not provide sufficient space for the data pages. This applies both to clean pages and to changed pages from the member. Failed writes must be avoided for proper performance as they lead to a delay of the transaction, if they happen during a commit, and may lead to serious delays for any other transaction requesting those pages.

P CASTOUT THRESHOLDS DETECTED FOR CLASSES: This is the number of times a class threshold is detected by DB2. A class threshold is detected when the pages in the class exceed, for example, 10% of the GBP specified in the

```
-ALTER GROUPBUFFERPOOL CLASST(10)
```

command. When the class castout threshold is detected, changed pages for the page sets belonging to that class are written from the coupling facility to DASD by the data sharing member that is castout owner for that page set. See 8.4.3, “Group Buffer Pool Castout Interval” on page 131 for more details.

Q CASTOUT THRESHOLDS DETECTED FOR GROUP BUFFER POOL: This represents the number of times a GBP threshold is detected by the coupling facility. A group buffer pool threshold is detected when the GBPOOLT threshold for the GBP is triggered.

R CASTOUTS: This is the total count of pages cast out in one measurement period. If the count is steady over a longer period, it can be taken for the number of different pages changed in that period.

$$\frac{\text{CASTOUTS R}}{\text{SYNCHRONOUS WRITES CHANGED PAGES K} + \text{ASYNCHRONOUS WRITES CHANGED PAGES M)}} = \frac{49002}{(47042 + 1)} = 1$$

gives you a measure for the hit ratio on changed pages, that is, the number of changes a page has during its life time in the coupling facility.

S ENGINES NOT AVAILABLE FOR CASTOUT: Castout cannot commence for the castout process scheduled. Castout waits for an engine to become available. DB2 V4 has up to 300 castout engines, depending on the castout workload and DBM1 storage available,

T ENGINES NOT AVAILABLE FOR WRITING: A write engine cannot be started, although needed for log writing or similar activities. The activity is delayed until an engine becomes available.

U OTHER INTERACTIONS: These are castout-related interactions with the coupling facility and are related to the value in the field CASTOUTS R .

8.2.3 Using DB2 PM Statistics Report

Statistics trace class 5 produces a new instrumentation facility interface identifier (IFCID) 230 record containing GBP information similar to that contained in -DIS GROUPBUFFERPOOL output, and this is included in the in the DB2 PM Statistics report. Figure 57 on page 123. shows an extract from this report. This extract is for the GBP summary for the total of all 4K buffer pools for the first CPC in the 2WP measurement, taken from Appendix A, “DB2 PM 2-Way Statistics Report (2WP)” on page 137.

```
GROUP TOT4K                          QUANTITY
---------------------------          --------
SYN.READS(XI)-DATA RETURNED A        22279.00
SYN.READS(XI)-R/W INTEREST  B           10.00
SYN.READS(XI)-NO R/W INTER. C            0.00
SYN.READS(NF)-DATA RETURNED D        27537.00
SYN.READS(NF)-R/W INTEREST  E        81112.00
SYN.READS(NF)-NO R/W INTER. F            0.00
ASYNC.READS-DATA RETURNED   G            0.00
ASYNC.READS-R/W INTEREST    H            0.00
ASYNC.READS-NO R/W INTEREST I            0.00

CLEAN PAGES SYNC.WRITTEN    K            0.00
CHANGED PAGES SYNC.WRITTEN  L          110.7K
CLEAN PAGES ASYNC.WRITTEN   M            0.00
CHANGED PAGES ASYNC.WRITTEN N           56.00

PAGES CASTOUT               P        61901.00
CASTOUT CLASS THRESHOLD     Q            0.00
GROUP BP CASTOUT THRESHOLD  R           64.00
CASTOUT ENGINE NOT AVAIL.   S            0.00
```

Figure 57. DB2 PM Statistics Report Showing Group Buffer Pool Totals

A SYN.READS(XI)-DATA RETURNED: This is the number of synchronous coupling facility read requests caused by the cross-system invalidation of the page in the local buffer pool. The page is returned from the GBP. If you increase the size of the GBP, the number of pages returned from the GBP can increase. Conversely, if you decrease the size of the GBP, the number of pages returned from the GBP can decrease, and the GBP cannot hold the pages long enough to allow them to be retrieved again.

B SYN.READS(XI)-R/W INTEREST: This the number of synchronous coupling facility read requests caused by the cross-system invalidation of the page in the local buffer pool. The page is not returned from the GBP, and a directory entry is created if one does not exist. The page is GBP-dependent at the time of the cross-system invalidation and still is at the time of the read request. The requested page has been removed from the group buffer pool for one of two reasons

- Shortage of buffer pages and consequent reclamation of this page
- shortage of directory entries and consequent removal of the page together with cross invalidation of that page in the local buffer pools of all members using that page.

C SYN.READS(XI)-NO R/W INTER.: This is the number of synchronous coupling facility read requests caused by the cross-system invalidation of a page in the local buffer pool. The page requested is not returned from the GBP, and no directory entry is created for it. The page is GBP-dependent at the time of the cross-system invalidation, however, it is not GBP-dependent at the time of the read request. The reason for not being GBP-dependent at the time of the read request can be either that the checkpoint on the writer side discovers no update activity during the last complete checkpoint interval, or only one member has the page set open.

D SYN.READS(NF)-DATA RETURNED: This is the number of pages returned from the coupling facility that could not be found in the local buffer pool of a page set that is GBP-dependent. The page is returned from the group buffer pool. If you increase the size of the GBP, the number of pages returned from the GBP can increase. Conversely, if you decrease the size of the GBP, the number of pages returned from the GBP can decrease, and the GBP cannot hold the pages long enough to allow them to be retrieved again.

E SYN.READS(NF)-R/W INTEREST: The number of synchronous coupling facility read request for pages that could not be found in the local buffer pool of a GBP-dependent page set. The page is not returned from the group buffer pool, and a directory entry is created if it does not exist.

F SYN.READS(NF)-NO R/W INTER.: A page is requested from the GBP although the page set is not GBP-dependent. This happens when either

- GBPCACHE ALL is specified for the page set and there is no read/write interest in the page set
- Only the requesting member has R/W interest in the page set; all other members have read-only interest. The writing member does not need to register a read-only page as no other member updates that page.

This section of counters show asynchronous reads from the GBP. Asynchronous reads can result from prefetch processing, and for those cases when P-lock negotiation necessitates registering pages cached in the member′s buffer pool.

G ASYNC.READS-DATA RETURNED: Same as D .

H ASYNC.READS-R/W INTEREST: Same as E .

I ASYNC.READS-NO R/W INTEREST: Same as F .

K CLEAN PAGES SYNC.WRITTEN: These are writes to the coupling facility. Only GBPCACHE ALL causes clean pages to be written to the coupling facility. When a data sharing member opens a table in a table space that is defined with GBPCACHE ALL, any page read is written to the coupling facility. This is well suited to transaction workloads with intersystem interest and a high R/W ratio, sufficient coupling facility resources and capable of achieving a high read hit ratio.

L CHANGED PAGES SYNC.WRITTEN: These writes to the coupling facility. are pages changed at least once and written because of the force-at-commit process.

M CLEAN PAGES ASYNC.WRITTEN: Clean pages are written out to the coupling facility only for GBPCACHE ALL. Make sure that your coupling facility has enough storage and capacity on the coupling links to sustain the data flow.

N CHANGED PAGES ASYNC.WRITTEN: These are pages written to the coupling facility by a data sharing member under the system task of DB2. This happens for all system-initiated tasks that write to the coupling facility, as for any write from virtual buffer pool when a write threshold is reached.

Tuning information:

Write threshold can be used to avoid writing large commits. Commits may write to the coupling facility in burst mode. The longer the commit interval, the bigger a commit write burst becomes. Unwieldy writes can be circumvented by setting the horizontal or vertical deferred write threshold very low, thus forcing a steady stream of pages to the coupling facility.

P PAGES CASTOUT: This is the number of pages castout from the coupling facility in the report interval. It is available for all GBPs and does not discriminate by castout method.

Q CASTOUT CLASS THRESHOLD: This is the number of times that the castout class threshold is triggered in the statistics interval.

R GROUP BP CASTOUT THRESHOLD: This is the number of times a GBP castout is initiated when the GBP castout threshold is triggered during the statistics interval. The CASTOUT CLASS THRESHOLD and the GROUP BP CASTOUT THRESHOLD are set with commands

```
-ALTER GROUPBUFFERPOOL CLASST(10)
and
-ALTER GROUPBUFFERPOOL GBPOOLT(50)
```

command. We use these values to achieve a constant castout rate over all data sharing members. Figure 54 on page 113 B and H have more details on the command.

S CASTOUT ENGINE NOT AVAIL.: This means that the castout class owner cannot start the castout process when asked by the GBP structure owner. The castout class owner then retries three times. When the GBP is used by more than one page set (castout class), other data sharing members also castout from that GBP. If the changed pages in the GBP cannot be cast out, the number of free and clean pages in the GBP is depleted over time.. To allow force-at-commit policy to write changed pages to the coupling facility, the coupling facility reclaims the changed pages. For the changed pages to be reclaimed, they must be cross invalidated, then put on the LPL. To remove the page from the LPL, recover the page set.

The DB2 PM accounting report contains similar information, however, no asynchronous activity is reported there. For that reason we comment on the Statistics Report - Long. The same comments apply to the accounting report.

8.2.3.1 Using DB2 PM Report for Buffer Pool Statistics

The Figure 58 on page 126 shows an extract. As you can see, there are prefetch requests 1 , however, there are no asynchronous reads 2 . Prefetch was turned off by specifying SPTH=0%. (In our case, the prefetch 1 was disabled as not suited for the workload. The prefetch would just begin when we had read the last page the transaction needed.)

```
BP5    READ OPERATIONS           QUANTITY
---------------------------      --------
SEQUENTIAL PREFETCH REQUEST 1    8460.00
SEQUENTIAL PREFETCH READS           0.00
PAGES READ VIA SEQ.PREFETCH 2       0.00
S.PRF.PAGES READ/S.PRF.READ          N/C

LIST PREFETCH REQUESTS              0.00
LIST PREFETCH READS                 0.00
PAGES READ VIA LIST PREFTCH         0.00
L.PRF.PAGES READ/L.PRF.READ          N/C

DYNAMIC PREFETCH REQUESTED          0.00
DYNAMIC PREFETCH READS              0.00
PAGES READ VIA DYN.PREFETCH         0.00
D.PRF.PAGES READ/D.PRF.READ          N/C

PREF.DISABLED-NO BUFFER     3    8460.00
```

Figure 58. DB2 PM Statistics Report Extract Showing Buffer Pool Reads. Extracted from Appendix Appendix A, "DB2 PM 2-Way Statistics Report (2WP)" on page 137

The report shown is a summary report for all plans. If you feel the failed register page list are significant for the performance of your system, see 8.3, "Sizing the Group Buffer Pool" on how to improve the situation.

8.3 Sizing the Group Buffer Pool

GBP tuning is a multiple-step process:

- Many decisions can be made fairly quickly, based on DB2 PM Statistics report, output of the DB2 DISPLAY GROUPBUFFERPOOL command, and monitoring the activities with the DB2 PM Online Monitor. These allow you to remedy most acute situations on the spot.
- Long-term observation may be needed in some critical cases. We suggest you save statistics over long periods of time and track GBP behavior periodically.

The following points can help you in your decision on what to look out for. Our suggestions are based on the assumption that you use your table spaces with GBPCACHE CHANGED (default).

8.3.1 DB2 Member Virtual Buffer Pool

In general, DB2 performance improves as more buffers are added up to a certain point. However, if the number of buffers in a member buffer pool associated with a GBP-dependent page set or partition is larger than the number of directory elements in the coupling facility structure, then problems can arise:

- Page reclamation and cross invalidation
- Page refreshes from DASD.

These problems can have a significant negative impact in a data sharing environment on the buffer pool hit ratio achieved by data sharing members.

8.3.2 Coupling Facility Structure Size

DB2 data sharing uses coupling facility structures to maintain data coherency across DB2 members of a data sharing group. Each group buffer pool must reside in a coupling facility structure. The coupling facility structure is used to:

- Store all updated pages of a GBP-dependent page set or partition.
- Register all pages of a GBP-dependent page set or partition that are cached in the buffer pool of data sharing members.

If the size of a coupling facility structure is too small, DB2 data sharing performance can degrade as follows:

- The castout threshold for changed pages is reached more frequently. DB2 castout processing is triggered more often. Castout processing becomes a continual background process. This leads to increased loading on the coupling facility structure, and increased CPU and I/O resource consumption.
- Coupling facility space management has to reclaim clean coupling facility directory elements more frequently. A directory element may or may not be associated with a page of GBP-dependent page set or partition. Since there is a good chance that the directory element being reclaimed is associated with a page that is cached in DB2 member buffer pools, the reclaiming of directory elements requires the cross-system invalidation of the copies of the page cached in DB2 member buffer pools and forces each DB2 member to refresh the page from DASD.

Both increases of I/O resource consumption can degrade transaction response times and decrease overall throughput.

8.3.3 Group Buffer Pool Directory Entries

A page must be registered for a particular member when a new page read from the GBP or DASD is cached locally. A directory entry is required for each distinct page cached locally across member buffer pools and for each changed page cached in the GBP. DB2 uses the directory to maintain coherency over all data in pages used in the system. Cross-system invalidation is used when a page is updated by one member of the group while another member is using that same page.

Length of a directory entry: You can go through the following steps to calculate the length of the directory entry:

```
Allocated GBP size in bytes = Allocated size (KB) * 1024

                                 Allocated GBP size in bytes
Size of a page including     = -------------------------------
directory entries                Number of data pages in GBP

                                 Size of a page including directory entries
Directory space allocated    = --------------------------------------------
per GBP page                                    4096

                                 Directory space allocated per GBP page
Length of a directory entry = -----------------------------------------
                                        Directory to data ratio
```

In the case of IRWW measurement for example, for group buffer pool GBP8:

- Directory to data ratio = 3.1 (A in Figure 54 on page 113)

- Allocatd size (KB) = 133120 (F in Figure 54 on page 113)
- Number of data pages in GBP = 28854 (I in Figure 54 on page 113)

The length of a directory entry is calculated as 203 bytes.

Note: The formula for the length of a directory entry provides only a rule-of-thumb value and does not provide an exact length, as a cache structure in the coupling facility has more detail than given here; however, it gives you an idea of directory entry size, including some overhead and unused space typical for the GBP you use.

Directory Entries Needed

D = SUM(V) + G

D Number of directory entries

V Number of pages in virtual buffer pool of each data sharing member

G Number of pages in GBP

The maximum number of directory entries is the total number of pages allocated in the virtual buffer pools for that group buffer pool plus the number of pages in the group buffer pool.

In the IRWW 2WP measurements, for example, the size of the virtual buffer pool for the two members is 25000. You may subtract from this value all pages that are cached more than once. The number of group buffer pool pages for our runs is 29,000, as you can see in Table 11 on page 51. Adding 2 * 25000 + 29000 gives us 79000 required directory entries. As you see in Table 11 on page 51, we used 89K, about 11% more than needed.

Directory Entries Needed

D = (25000 + 25000) + 29000 = 79000

A shortage of GBP directory entries is more serious than a shortage of GBP data pages, and results in increased values for the fields

```
SYN.READS(XI)-R/W INTEREST
SYN.READS(XI)-NO R/W INTEREST
```

in the group buffer pool block in the DB2 PM Statistics report.

8.3.4 Group Buffer Pool Data Pages

To get a first estimate for the size of the GBP, begin with the premise that the GBP must be able to accommodate all updated pages that are written to it by force-at-commit policy.

You can go through the following steps to calculate the number of GBP pages:

```
                                 Desired page residence time (minutes)
Factor for prolonged stay  F = ----------------------------------------
of pages in the GBP                    Castout interval (minutes)


                   100
GBP pages =  --------------- * Snapshot value of changed pages *  F
                GBP castout
```

```
threshold
```

In the IRWW 2WP measurements for group buffer pool GBP8:

- Desired page residence time is 11 minutes
- Castout interval is 10 minutes
- GBP castput threshold is the default value of 50%.
- Snapshot value of changed pages is 13325 (H in Figure 55 on page 115)

GBP pages is calculated as 29000

You can also use this formula to calculate the average residence time for a page in the GBP with a given size.

```
                                     Castout interval       Casout threshold
Residence time = GBP pages * -------------------- * --------------------
                                      Snapshot value               100
```

8.3.5 Group Buffer Pool Directory-to-Data Ratio

There is no direct relationship between the number of directory entries and the number of pages in the GBP. Directory entries are used to track coherency of read and write pages in the local buffer pools. GBP pages are used mostly at commit time to keep pages changed by DB2. Therefore the directory on average has more entries than the GBP data pages.

Directory to Data Ratio: Use the following formula to calculate GBP directory-to-data ratio:

```
                              Number of directory entries
Directory to data ratio = ------------------------------
                              Number of pages in the GBP
```

In the IRWW 2WP measurements for group buffer pool GBP8:

- Number of directory entries is 79000—See section 8.3.3, “Group Buffer Pool Directory Entries” on page 127.
- Number of pages in the GBP is 29000—See section 8.3.4, “Group Buffer Pool Data Pages” on page 128.

Directory to Data Ratio is calculated as 2.7

Group Buffer Pool Size: Use the following formula to calculate GBP size:

```
  A  = Directory to data ratio
  B  = Length of a directory entry
                                A  *  B
GBP size = GBP pages * (1 + ---------) * 4096
                                 4096
```

For example, substituting the values 29000 for GBP pages, 2.7 for the directory to data ratio, and 203 for the length of a directory entry the GBP size is calculated as 134 MB.

See *IBM DATABASE 2 for MVS/ESA Version 4 Data Sharing: Planning and Administration* , Chapter 6, “Tuning group buffer pools” for details on how to size the GBP.

8.4 DB2 Castout Process

Each GBP contains both changed and clean pages. A changed page in a GBP is more current in terms of data content than the copy of the page on DASD. A certain percentage of each GBP must have clean pages that can be used for page replacement. The castout process is to ensure that changed pages are written to DASD often enough that a GBP does not get filled with changed pages. Castout pages are left in the GBP, but are now marked as clean. DB2 castout process uses MVS XES services for reading and locking.

Castout process is triggered by one of the following:

- Castout class thresholds
- GBP thresholds
- GBP checkpoints triggered by a timer
- Conversion of page set or partition to non-GBP-dependent.

Castout process is a background activity performed asynchronously, on a castout class basis. The name of each page set or partition is hashed into a castout class. When a changed page is written to a GBP by the buffer manager of a data sharing member using an MVS XES service, the castout class is passed and feedback is given on the number of changed pages in the given castout class. The data sharing member tests the value against the respective thresholds and triggers castout processing to be performed by the DB2 member of the data sharing group who owns the page set or partition. Communication with the data sharing member is achieved by sending an XCF message. The data sharing member who owns the particular page set or partition is the first member with write interest when there is intersystem read/write interest.

8.4.1 Group Buffer Pool Class Castout Threshold

As with the virtual buffer pool, this value is dependent on your type of workload and can change considerably between online (during the day) and batch (overnight) environment. There can be castouts for the same GBP in different members, depending on who is the castout owner for a page set or partition. In the IRWW 2WP measurement, the number of pages castout (field PAGES CASTOUT in the Group Buffer Pool block for BP8 in the DB2 PM Statistics report) is 46414 in the first member and 31073 in the second member. The number of times the class castout threshold is triggered during the statistics interval is shown by the field CASTOUT CLASS THRESHOLD in the Group Buffer Pool block in the DB2 PM Statistics report.

8.4.2 Group Buffer Pool Castout Threshold

GBP castout threshold affects the availability of the GBP, as this influences when the buffers are removed from the GBP (next to class castout and GBP checkpoint). For GBP-dependent data sets, force-at-commit policy writes all changed pages into the GBP. The GBP castout threshold is a safety margin to protect the GBP from accidental flooding by hyperactive applications.

DB2 forces a GBP castout as the castout threshold is reached. The number of times the GBP castout threshold is triggered during the statistics interval is shown by the field GROUP BP CASTOUT THRESHOLD in the Group Buffer Pool block in the DB2 PM Statistics report.

8.4.3 Group Buffer Pool Castout Interval

In principle, there are two ways of writing changed pages from the GBP to DASD:

- Regular castout by checkpoint

  The DB2 command

  `-ALTER GROUPBUFFERPOOL GBPCHKPT(10)`

  sets the checkpoint interval to 10 minutes. Checkpoint castout works in a burst mode, leaving the system to its tasks in the meantime.

  The small the GBP, or the higher the update rate of the applications, the more frequent the checkpoint should be. That means the GBPCHKPT value should be smaller.

- Constant castout by threshold.

  You can set the checkpoint interval to a high value and still force castouts periodically by setting the appropriate values for the GBP castout and GBP class castout thresholds. For example, in the IRWW measurement runs,

  `-ALTER GROUPBUFFERPOOL CLASST(10) GBPOOLT(50)`

  command sets the GBP class castout threshold to 10% and the GBP castout threshold to 50%. This way it is possible to achieve a constant and balanced castout over all members that fit well with the workload.

8.5 Analyzing the Use of the Shared Communication Area

The SCA is not critical for performance. For reference we include here a few details.

8.5.1 Using the XCF DISPLAY Command

The XCF DISPLAY command gives static information about size, and allocation; it also gives dynamic information about the number of connections and their system names. See Figure 59 on page 132 for an example of such a display, taken from the IRWW measurement run.

```
 D XCF,STR,STRNAME=DSNDB0G_SCA    -----------------------
 IXC360I  10.22.33  DISPLAY XCF 170
 STRNAME: DSNDB0G_SCA
  STATUS: ALLOCATED
A  POLICY SIZE    : 10000 K
   POLICY INITSIZE: N/A
   REBUILD PERCENT: N/A
   PREFERENCE LIST: LF02
   EXCLUSION LIST IS EMPTY

  ACTIVE STRUCTURE
  ----------------
   ALLOCATION TIME: 09/09/95 09:23:01
   CFNAME         : LF02
   COUPLING FACILITY: 009674.IBM.00.000000040029
                      PARTITION: 1   CPCID: 00
B  ACTUAL SIZE    : 10240 K
C  STORAGE INCREMENT SIZE: 256 K
   VERSION        : ABA5BA77 E90F7200
   DISPOSITION    : KEEP
   ACCESS TIME    : 0
   MAX CONNECTIONS: 32
D  # CONNECTIONS  : 2

  CONNECTION NAME  ID VERSION  SYSNAME  JOBNAME  ASID STATE
  ---------------- -- -------- -------- -------- ---- ------
E DB2_DB1G         01 0001023D STLABC1  DB1GMSTR 0194 ACTIVE
  DB2_DB2G         02 00020186 STLABC2  DB2GMSTR 0192 ACTIVE
```

Figure 59. XCF SCA Structure Display. D XCF,STR,STRNAME=DSNDB0G_SCA

A This is the policy size as defined in the CFRM policy.

B This is the storage size allocated by DB2 V4 in the coupling facility. As the policy size is rounded up to the next higher multiple of 256 K (C), the original 10000 K is rounded to 10240 K

D This is the number of data sharing members connected to the coupling facility at the time of the display.

E These are the names of the connections from the data sharing members to the coupling facility.

8.5.2 Using RMF Coupling Facility Activity Report

The RMF activity report gives more dynamic information. Figure 60 on page 133 is taken from the *RMF Coupling Facility Activity Report*.

```
     C O U P L I N G   F A C I L I T Y   A C T I V I T Y
------------------------------------------------------------
   1          COUPLING  FACILITY  USAGE  SUMMARY
------------------------------------------------------------
STRUCTURE SUMMARY
------------------------------------------------------------

                           % OF            % OF     AVG
STRUCTURE        ALLOC       CF       #     ALL      REQ/
NAME             SIZE    STORAGE    REQ    REQ       SEC

DSNDB0G_SCA       10M      0.5%     221 D  0.0% A   0.22 B
DSNDB0G_GBP0       3M      0.1%   10294    0.8%    10.13
DSNDB0G_GBP1      14M      0.7%   30215    2.5%    29.74
DSNDB0G_GBP3      86M      4.3%  516624   42.1%   508.49
DSNDB0G_GBP4      20M      1.0%   20834    1.7%    20.51
DSNDB0G_GBP5      73M      3.7%   26185    2.1%    25.77
DSNDB0G_GBP6      39M      1.9%  118817    9.7%   116.95
DSNDB0G_GBP7      22M      1.1%   69811    5.7%    68.71
DSNDB0G_GBP8     130M      6.5%  433518   35.3%   426.69
                ------    -----  -------  ------  -------
STRUCTURE TOTALS 396M     19.9%    1227K E  100%  1207.2 C
```

Figure 60. RMF Coupling Facility Structure Usage Summary Report. For Coupling Facility Model 9674/R61. Extracted from F.4, "RMF – Coupling Facility Activity" on page 177

1 Reference point to Appendix F.4, "RMF – Coupling Facility Activity" on page 177

A This represents the percentage of activities in the SCA structure related to all activities in this coupling facility.

```
                                       # REQ SCA STRUCTURE     D
Percentage of activites in the SCA = ----------------------------- * 100
                                       # REQ STRUCTURE TOTALS  E

                                        221
                                   = --------- = 0.0%
                                       1227K
```

B This is the average number of requests per second to the coupling facility on behalf of the SCA, in this case 0.22. 1/ B (=4.54) gives you the seconds between requests or their frequency. As there are two data sharing members connected to the coupling facility, each member has one request to the coupling facility every 9 seconds.

C This is the total number of requests per second to the coupling facility for all GBPs and the SCA. As you can see, the SCA activities do not influence performance much in the IRWW measurement runs.

Figure 61 on page 134 is interesting for the average service time for the requests to the SCA. It is extracted from the *RMF Coupling Facility Activity Report*.

```
STRUCTURE NAME = DSNDB0G_SCA       TYPE = LIST
           # REQ     -------------- REQUESTS -------------
SYSTEM     TOTAL              #      % OF   -SERV TIME(MIC)-
NAME                          REQ    ALL      AVG    STD_DEV

STLABC1      221     SYNC     219    99.1%    500.7 A    59.9



 2
           # REQ     ----------- REQUESTS -----------
SYSTEM     TOTAL              #    -SERVICE TIME(MIC)-
NAME       AVG/SEC            REQ     AVG     STD_DEV

STLABC1      1228K   SYNC       1200K   265.8 B    94.6
```

Figure 61. RMF Coupling Facility Structure Activity Report. For Coupling Facility Model 9674/R61. Extracted from Figure 92 on page 179 1 *and Figure 93 on page 180* 2

1 This is an extract from coupling facility structure activity section from Figure 92 on page 179.

2 This is an extract from subchannel activity section from Figure 93 on page 180.

A REQUESTS, SERV AVG: represents the average service time for a request to the SCA. Please note that this average service time is about twice as long as the average service time for the group buffer pools. In this coupling facility all the GBPs and the SCA are allocated.

B REQUESTS, SERVICE TIME(MIC): This is the average service time for all structures in this coupling facility. This value is consistent with the expected coupling facility response time.

8.6 Analyzing the MVS XCF Communication

Check the cache size for the MVS XCF group. As can be seen in the Figure 62, RMF lists all request and how they fit into the buffer size available.

```
                     OUTBOUND FROM STLABC1
-------------------------------------------------------------


                                              ----- BUFFER -----
TO         TRANSPORT   BUFFER         REQ    %     %     %     %
SYSTEM     CLASS       LENGTH         OUT   SML   FIT   BIG   OVR
 A
STLABC2    DEFAULT      B 956       3,835  C 0   >99    <1   100
STLABC3    DEFAULT        956       3,026    0   >99    <1   100
```

Figure 62. RMF Coupling Facility Report Showing XCF Communication. For Coupling Facility Model 9674/R61. Extracted from F.2, “RMF – XCF Activities” on page 172

A The system identifier of the originating system shows the MVS XCF group you should be changing buffer size for.

B For the IRWW, the message class length is 956 bytes. Only a fraction of one percent of the messages uses a buffer size larger than provided and uses more than one buffer.

C   If there are performance problems, check that all the buffers fit (FIT in Figure 62) Too small a message length for the buffer size (SML in Figure 62) causes unneeded line traffic, and too big a message length for the buffer size (BIG in Figure 62) forces unneeded error recovery. See *Setting Up a Sysplex*, Chapter 6, "Tuning the Maximum Message Buffer Space" for more detailed information.

See *IBM DATABASE 2 for MVS/ESA Version 4 Data Sharing: Planning and Administration*, Chapter 5, "Operating with Data Sharing" for more information on SCA.

Appendix A. DB2 PM 2-Way Statistics Report (2WP)

This appendix compares DB2 PM Statistic reports for the 2-way measurements.

A.1 DB2 PM Statistics Report: 2-Way First CPC, SQL Activity

```
   LOCATION: DSNDB0G                              DB2 PERFORMANCE MONITOR (V4)                                  PAGE: 1-1
      GROUP: DSNDB0G                               STATISTICS REPORT - LONG                     REQUESTED FROM: 09/09/95 08:41:15.00
     MEMBER: DB1G                                                                                           TO: 09/09/95 08:51:17.00
  SUBSYSTEM: DB1G                                                                                INTERVAL FROM: 09/09/95 08:41:16.36
DB2 VERSION: V4                                          SCOPE: MEMBER                                      TO: 09/09/95 08:51:16.73

---- HIGHLIGHTS ----------------------------------------------------------------------------------------------------------
INTERVAL START  : 09/09/95 08:41:16.36   SAMPLING START: 09/09/95 08:41:16.36   TOTAL THREADS      :      0.00
INTERVAL END    : 09/09/95 08:51:16.73   SAMPLING END  : 09/09/95 08:51:16.73   TOTAL COMMITS      : 18056.00
INTERVAL ELAPSED:         10:00.377445   OUTAGE ELAPSED:              0.000000   DATA SHARING MEMBER:       N/A

SQL DML                        QUANTITY  /MINUTE  /THREAD  /COMMIT   SQL DCL                        QUANTITY  /MINUTE  /THREAD  /COMMIT
---------------------------    --------  -------  -------  -------   ---------------------------    --------  -------  -------  -------
SELECT                         86981.00  8692.63      N/C     4.82   LOCK TABLE                         0.00     0.00      N/C     0.00
INSERT                         53011.00  5297.77      N/C     2.94   GRANT                              0.00     0.00      N/C     0.00
UPDATE                         64530.00  6448.94      N/C     3.57   REVOKE                             0.00     0.00      N/C     0.00
DELETE                          2505.00   250.34      N/C     0.14   SET HOST VARIABLE                  0.00     0.00      N/C     0.00
                                                                     SET CURRENT SQLID                  0.00     0.00      N/C     0.00
PREPARE                            0.00     0.00      N/C     0.00   SET CURRENT DEGREE                 0.00     0.00      N/C     0.00
DESCRIBE                           0.00     0.00      N/C     0.00   SET CURRENT RULES                  0.00     0.00      N/C     0.00
DESCRIBE TABLE                     0.00     0.00      N/C     0.00
OPEN                           81417.00  8136.58      N/C     4.51   CONNECT TYPE 1                     0.00     0.00      N/C     0.00
CLOSE                          45883.00  4585.42      N/C     2.54   CONNECT TYPE 2                     0.00     0.00      N/C     0.00
FETCH                            166.8K    16.7K      N/C     9.24   RELEASE                            0.00     0.00      N/C     0.00
                                                                     SET CONNECTION                     0.00     0.00      N/C     0.00
TOTAL                            501.2K    50.1K      N/C    27.76
                                                                     TOTAL                              0.00     0.00      N/C     0.00

SQL DDL                        QUANTITY  /MINUTE  /THREAD  /COMMIT   STORED PROCEDURES              QUANTITY  /MINUTE  /THREAD  /COMMIT
---------------------------    --------  -------  -------  -------   ---------------------------    --------  -------  -------  -------
CREATE TABLE                       0.00     0.00      N/C     0.00   CALL STATEMENTS EXECUTED           0.00     0.00      N/C     0.00
CREATE INDEX                       0.00     0.00      N/C     0.00   PROCEDURE ABENDS                   0.00     0.00      N/C     0.00
CREATE VIEW                        0.00     0.00      N/C     0.00   CALL STATEMENT TIMEOUTS            0.00     0.00      N/C     0.00
CREATE SYNONYM                     0.00     0.00      N/C     0.00   CALL STATEMENT REJECTED            0.00     0.00      N/C     0.00
CREATE TABLESPACE                  0.00     0.00      N/C     0.00
CREATE DATABASE                    0.00     0.00      N/C     0.00
CREATE STOGROUP                    0.00     0.00      N/C     0.00
CREATE ALIAS                       0.00     0.00      N/C     0.00

ALTER  TABLE                       0.00     0.00      N/C     0.00
ALTER  INDEX                       0.00     0.00      N/C     0.00
ALTER  TABLESPACE                  0.00     0.00      N/C     0.00
ALTER  DATABASE                    0.00     0.00      N/C     0.00
ALTER  STOGROUP                    0.00     0.00      N/C     0.00

DROP   TABLE                       0.00     0.00      N/C     0.00
DROP   INDEX                       0.00     0.00      N/C     0.00
DROP   VIEW                        0.00     0.00      N/C     0.00
DROP   SYNONYM                     0.00     0.00      N/C     0.00
DROP   TABLESPACE                  0.00     0.00      N/C     0.00
DROP   DATABASE                    0.00     0.00      N/C     0.00
DROP   STOGROUP                    0.00     0.00      N/C     0.00
DROP   ALIAS                       0.00     0.00      N/C     0.00
DROP   PACKAGE                     0.00     0.00      N/C     0.00

COMMENT ON                         0.00     0.00      N/C     0.00
LABEL ON                           0.00     0.00      N/C     0.00

TOTAL                              0.00     0.00      N/C     0.00
```

Figure 63. Run 2WP: 2-Way First CPC, SQL Activity

A.2 DB2 PM Statistics Report: 2-Way Second CPC, SQL Activity

```
   LOCATION: DSNDB0G                               DB2 PERFORMANCE MONITOR (V4)                                      PAGE: 1-1
      GROUP: DSNDB0G                                STATISTICS REPORT - LONG                          REQUESTED FROM: 09/09/95 08:41:15.00
     MEMBER: DB2G                                                                                                 TO: 09/09/95 08:51:17.00
  SUBSYSTEM: DB2G                                                                                  INTERVAL FROM: 09/09/95 08:41:16.41
DB2 VERSION: V4                                          SCOPE: MEMBER                                            TO: 09/09/95 08:51:16.85

---- HIGHLIGHTS ----------------------------------------------------------------------------------------------------
INTERVAL START  : 09/09/95 08:41:16.41   SAMPLING START: 09/09/95 08:41:16.41   TOTAL THREADS      :     0.00
INTERVAL END    : 09/09/95 08:51:16.85   SAMPLING END  : 09/09/95 08:51:16.85   TOTAL COMMITS      : 18191.00
INTERVAL ELAPSED:          10:00.441800   OUTAGE ELAPSED:               0.000000   DATA SHARING MEMBER:      N/A

SQL DML                      QUANTITY  /MINUTE  /THREAD  /COMMIT   SQL DCL                      QUANTITY  /MINUTE  /THREAD  /COMMIT
---------------------------  --------  -------  -------  -------   ---------------------------  --------  -------  -------  -------
SELECT                       88199.00  8813.41      N/C     4.85   LOCK TABLE                       0.00     0.00      N/C     0.00
INSERT                       52689.00  5265.02      N/C     2.90   GRANT                            0.00     0.00      N/C     0.00
UPDATE                       68060.00  6800.99      N/C     3.74   REVOKE                           0.00     0.00      N/C     0.00
DELETE                        3799.00   379.62      N/C     0.21   SET HOST VARIABLE                0.00     0.00      N/C     0.00
                                                                   SET CURRENT SQLID                0.00     0.00      N/C     0.00
PREPARE                          0.00     0.00      N/C     0.00   SET CURRENT DEGREE               0.00     0.00      N/C     0.00
DESCRIBE                         0.00     0.00      N/C     0.00   SET CURRENT RULES                0.00     0.00      N/C     0.00
DESCRIBE TABLE                   0.00     0.00      N/C     0.00
OPEN                         83582.00  8352.05      N/C     4.59   CONNECT TYPE 1                   0.00     0.00      N/C     0.00
CLOSE                        48150.00  4811.46      N/C     2.65   CONNECT TYPE 2                   0.00     0.00      N/C     0.00
FETCH                          169.3K    16.9K      N/C     9.31   RELEASE                          0.00     0.00      N/C     0.00
                                                                   SET CONNECTION                   0.00     0.00      N/C     0.00
TOTAL                          513.8K    51.3K      N/C    28.24
                                                                   TOTAL                            0.00     0.00      N/C     0.00

SQL DDL                      QUANTITY  /MINUTE  /THREAD  /COMMIT   STORED PROCEDURES            QUANTITY  /MINUTE  /THREAD  /COMMIT
---------------------------  --------  -------  -------  -------   ---------------------------  --------  -------  -------  -------
CREATE TABLE                     0.00     0.00      N/C     0.00   CALL STATEMENTS EXECUTED         0.00     0.00      N/C     0.00
CREATE INDEX                     0.00     0.00      N/C     0.00   PROCEDURE ABENDS                 0.00     0.00      N/C     0.00
CREATE VIEW                      0.00     0.00      N/C     0.00   CALL STATEMENT TIMEOUTS          0.00     0.00      N/C     0.00
CREATE SYNONYM                   0.00     0.00      N/C     0.00   CALL STATEMENT REJECTED          0.00     0.00      N/C     0.00
CREATE TABLESPACE                0.00     0.00      N/C     0.00
CREATE DATABASE                  0.00     0.00      N/C     0.00
CREATE STOGROUP                  0.00     0.00      N/C     0.00
CREATE ALIAS                     0.00     0.00      N/C     0.00

ALTER  TABLE                     0.00     0.00      N/C     0.00
ALTER  INDEX                     0.00     0.00      N/C     0.00
ALTER  TABLESPACE                0.00     0.00      N/C     0.00
ALTER  DATABASE                  0.00     0.00      N/C     0.00
ALTER  STOGROUP                  0.00     0.00      N/C     0.00

DROP   TABLE                     0.00     0.00      N/C     0.00
DROP   INDEX                     0.00     0.00      N/C     0.00
DROP   VIEW                      0.00     0.00      N/C     0.00
DROP   SYNONYM                   0.00     0.00      N/C     0.00
DROP   TABLESPACE                0.00     0.00      N/C     0.00
DROP   DATABASE                  0.00     0.00      N/C     0.00
DROP   STOGROUP                  0.00     0.00      N/C     0.00
DROP   ALIAS                     0.00     0.00      N/C     0.00
DROP   PACKAGE                   0.00     0.00      N/C     0.00

COMMENT ON                       0.00     0.00      N/C     0.00
LABEL ON                         0.00     0.00      N/C     0.00

TOTAL                            0.00     0.00      N/C     0.00
```

Figure 64. Run 2WP: 2-Way Second CPC, SQL Activity

A.3 DB2 PM Statistics Report: 2-Way First CPC, Locking Activity

```
   LOCATION: DSNDB0G                          DB2 PERFORMANCE MONITOR (V4)                                    PAGE: 1-4
      GROUP: DSNDB0G                           STATISTICS REPORT - LONG                   REQUESTED FROM: 09/09/95 08:41:15.00
     MEMBER: DB1G                                                                                     TO: 09/09/95 08:51:17.00
  SUBSYSTEM: DB1G                                                                          INTERVAL FROM: 09/09/95 08:41:16.36
DB2 VERSION: V4                                      SCOPE: MEMBER                                    TO: 09/09/95 08:51:16.73

---- HIGHLIGHTS ----------------------------------------------------------------------------------------------------
INTERVAL START  : 09/09/95 08:41:16.36    SAMPLING START: 09/09/95 08:41:16.36    TOTAL THREADS      :     0.00
INTERVAL END    : 09/09/95 08:51:16.73    SAMPLING END  : 09/09/95 08:51:16.73    TOTAL COMMITS      : 18056.00
INTERVAL ELAPSED:         10:00.377445    OUTAGE ELAPSED:             0.000000    DATA SHARING MEMBER:      N/A

LOCKING ACTIVITY            QUANTITY  /MINUTE  /THREAD  /COMMIT   DATA SHARING LOCKING        QUANTITY  /MINUTE  /THREAD  /COMMIT
--------------------------- --------  -------  -------  -------   --------------------------- --------  -------  -------  -------
SUSPENSIONS (ALL)            3391.00   338.89      N/C     0.19   LOCK REQUESTS (P-LOCKS)     25890.00  2587.37      N/C     1.43
SUSPENSIONS (LOCK ONLY)       343.00    34.28      N/C     0.02   UNLOCK REQUESTS (P-LOCKS)   25855.00  2583.87      N/C     1.43
SUSPENSIONS (LATCH ONLY)     3048.00   304.61      N/C     0.17   CHANGE REQUESTS (P-LOCKS)      10.00     1.00      N/C     0.00
SUSPENSIONS (OTHER)             0.00     0.00      N/C     0.00
                                                                  SYNCH.XES - LOCK REQUESTS      128.9K    12.9K     N/C     7.14
TIMEOUTS                        0.00     0.00      N/C     0.00   SYNCH.XES - CHANGE REQUESTS 64952.00  6491.12      N/C     3.60
DEADLOCKS                       0.00     0.00      N/C     0.00   SYNCH.XES - UNLOCK REQUESTS    119.6K    12.0K     N/C     6.62
                                                                  ASYNCH.XES - RESOURCES          0.00     0.00      N/C     0.00
LOCK REQUESTS                 147.5K    14.7K      N/C     8.17
UNLOCK REQUESTS             54317.00  5428.29      N/C     3.01   SUSPENDS - IRLM GLOBAL CONT   621.00    62.06      N/C     0.03
QUERY REQUESTS                  0.00     0.00      N/C     0.00   SUSPENDS - XES GLOBAL CONT.     0.00     0.00      N/C     0.00
CHANGE REQUESTS             65204.00  6516.30      N/C     3.61   SUSPENDS - FALSE CONTENTION  1016.00   101.54      N/C     0.06
OTHER REQUESTS                  0.00     0.00      N/C     0.00   INCOMPATIBLE RETAINED LOCK      0.00     0.00      N/C     0.00

LOCK ESCALATION (SHARED)        0.00     0.00      N/C     0.00   NOTIFY MESSAGES SENT          692.00    69.16      N/C     0.04
LOCK ESCALATION (EXCLUSIVE)     0.00     0.00      N/C     0.00   NOTIFY MESSAGES RECEIVED        0.00     0.00      N/C     0.00
                                                                  P-LOCK/NOTIFY EXITS ENGINES    10.00      N/A      N/A      N/A
DRAIN REQUESTS                  0.00     0.00      N/C     0.00   P-LCK/NFY EX.ENGINE UNAVAIL     0.00     0.00      N/C     0.00
DRAIN REQUESTS FAILED           0.00     0.00      N/C     0.00
CLAIM REQUESTS                165.8K    16.6K      N/C     9.18   PSET/PART P-LCK NEGOTIATION     0.00     0.00      N/C     0.00
CLAIM REQUESTS FAILED           0.00     0.00      N/C     0.00   PAGE P-LOCK NEGOTIATION       388.00    38.78      N/C     0.02
                                                                  OTHER P-LOCK NEGOTIATION        0.00     0.00      N/C     0.00
                                                                  P-LOCK CHANGE DURING NEG.     388.00    38.78      N/C     0.02

GLOBAL DDF ACTIVITY         QUANTITY  /MINUTE  /THREAD  /COMMIT   QUERY PARALLELISM           QUANTITY  /MINUTE  /THREAD  /COMMIT
--------------------------- --------  -------  -------  -------   --------------------------- --------  -------  -------  -------
DBAT QUEUED-MAXIMUM ACTIVE       N/P      N/P      N/P      N/A   MAX.DEGREE OF PARALLELISM       0.00      N/A      N/A      N/A
CONV.DEALLOC-MAX.CONNECTED       N/P      N/P      N/P      N/A   PARALLEL GROUPS EXECUTED        0.00     0.00      N/C     0.00
INACTIVE DBATS - CURRENTLY       N/P      N/A      N/A      N/A   EXECUTED AS PLANNED             0.00     0.00      N/C     0.00
INACTIVE DBATS - HWM             N/P      N/A      N/A      N/A   REDUCED DEGREE - NO BUFFER      0.00     0.00      N/C     0.00
COLD START CONNECTIONS           N/P      N/P      N/P      N/P   FALL TO SEQUENTIAL-CURSOR       0.00     0.00      N/C     0.00
WARM START CONNECTIONS           N/P      N/P      N/P      N/P   FALL TO SEQUENTIAL-NO ESA       0.00     0.00      N/C     0.00
RESYNCHRONIZATION ATTEMPTED      N/P      N/P      N/P      N/P   FALL TO SEQUENTIAL-NO BUFF.     0.00     0.00      N/C     0.00
RESYNCHRONIZATION SUCCEEDED      N/P      N/P      N/P      N/P   FALL TO SEQUENTIAL-ENCL.SER     0.00     0.00      N/C     0.00
```

Figure 65. Run 2WP: 2-Way First CPC, Locking Activity

A.4 DB2 PM Statistics Report: 2-Way Second CPC, Locking Activity

```
   LOCATION: DSNDB0G                              DB2 PERFORMANCE MONITOR (V4)                                PAGE: 1-4
      GROUP: DSNDB0G                               STATISTICS REPORT - LONG                    REQUESTED FROM: 09/09/95 08:41:15.00
     MEMBER: DB2G                                                                                          TO: 09/09/95 08:51:17.00
  SUBSYSTEM: DB2G                                                                               INTERVAL FROM: 09/09/95 08:41:16.41
DB2 VERSION: V4                                          SCOPE: MEMBER                                     TO: 09/09/95 08:51:16.85

---- HIGHLIGHTS ----------------------------------------------------------------------------------------------------
INTERVAL START  : 09/09/95 08:41:16.41   SAMPLING START: 09/09/95 08:41:16.41   TOTAL THREADS      :    0.00
INTERVAL END    : 09/09/95 08:51:16.85   SAMPLING END  : 09/09/95 08:51:16.85   TOTAL COMMITS      : 18191.00
INTERVAL ELAPSED:         10:00.441800   OUTAGE ELAPSED:              0.000000   DATA SHARING MEMBER:     N/A

LOCKING ACTIVITY            QUANTITY  /MINUTE  /THREAD  /COMMIT   DATA SHARING LOCKING        QUANTITY  /MINUTE  /THREAD  /COMMIT
--------------------------- --------  -------  -------  -------   --------------------------- --------  -------  -------  -------
SUSPENSIONS (ALL)            3752.00   374.92      N/C     0.21   LOCK REQUESTS (P-LOCKS)     27573.00  2755.27      N/C     1.52
SUSPENSIONS (LOCK ONLY)       375.00    37.47      N/C     0.02   UNLOCK REQUESTS (P-LOCKS)   27553.00  2753.27      N/C     1.51
SUSPENSIONS (LATCH ONLY)     3377.00   337.45      N/C     0.19   CHANGE REQUESTS (P-LOCKS)       4.00     0.40      N/C     0.00
SUSPENSIONS (OTHER)             0.00     0.00      N/C     0.00
                                                                  SYNCH.XES - LOCK REQUESTS     143.9K    14.4K      N/C     7.91
TIMEOUTS                        0.00     0.00      N/C     0.00   SYNCH.XES - CHANGE REQUESTS 68893.00  6884.23      N/C     3.79
DEADLOCKS                       0.00     0.00      N/C     0.00   SYNCH.XES - UNLOCK REQUESTS   127.9K    12.8K      N/C     7.03
                                                                  ASYNCH.XES - RESOURCES          0.00     0.00      N/C     0.00
LOCK REQUESTS                 160.6K    16.0K      N/C     8.83
UNLOCK REQUESTS             59529.00  5948.52      N/C     3.27   SUSPENDS - IRLM GLOBAL CONT   643.00    64.25      N/C     0.04
QUERY REQUESTS                  0.00     0.00      N/C     0.00   SUSPENDS - XES GLOBAL CONT.     0.00     0.00      N/C     0.00
CHANGE REQUESTS             69149.00  6909.81      N/C     3.80   SUSPENDS - FALSE CONTENTION  1085.00   108.42      N/C     0.06
OTHER REQUESTS                  0.00     0.00      N/C     0.00   INCOMPATIBLE RETAINED LOCK      0.00     0.00      N/C     0.00

LOCK ESCALATION (SHARED)        0.00     0.00      N/C     0.00   NOTIFY MESSAGES SENT            0.00     0.00      N/C     0.00
LOCK ESCALATION (EXCLUSIVE)     0.00     0.00      N/C     0.00   NOTIFY MESSAGES RECEIVED      692.00    69.15      N/C     0.04
                                                                  P-LOCK/NOTIFY EXITS ENGINES    10.00      N/A      N/A      N/A
DRAIN REQUESTS                  0.00     0.00      N/C     0.00   P-LCK/NFY EX.ENGINE UNAVAIL     0.00     0.00      N/C     0.00
DRAIN REQUESTS FAILED           0.00     0.00      N/C     0.00
CLAIM REQUESTS                166.1K    16.6K      N/C     9.13   PSET/PART P-LCK NEGOTIATION     0.00     0.00      N/C     0.00
CLAIM REQUESTS FAILED           0.00     0.00      N/C     0.00   PAGE P-LOCK NEGOTIATION       394.00    39.37      N/C     0.02
                                                                  OTHER P-LOCK NEGOTIATION        0.00     0.00      N/C     0.00
                                                                  P-LOCK CHANGE DURING NEG.     394.00    39.37      N/C     0.02

GLOBAL DDF ACTIVITY         QUANTITY  /MINUTE  /THREAD  /COMMIT   QUERY PARALLELISM           QUANTITY  /MINUTE  /THREAD  /COMMIT
--------------------------- --------  -------  -------  -------   --------------------------- --------  -------  -------  -------
DBAT QUEUED-MAXIMUM ACTIVE       N/P      N/P      N/P      N/A   MAX.DEGREE OF PARALLELISM       0.00      N/A      N/A      N/A
CONV.DEALLOC-MAX.CONNECTED       N/P      N/P      N/P      N/A   PARALLEL GROUPS EXECUTED        0.00     0.00      N/C     0.00
INACTIVE DBATS - CURRENTLY       N/P      N/A      N/A      N/A   EXECUTED AS PLANNED             0.00     0.00      N/C     0.00
INACTIVE DBATS - HWM             N/P      N/A      N/A      N/A   REDUCED DEGREE - NO BUFFER      0.00     0.00      N/C     0.00
COLD START CONNECTIONS           N/P      N/P      N/P      N/P   FALL TO SEQUENTIAL-CURSOR       0.00     0.00      N/C     0.00
WARM START CONNECTIONS           N/P      N/P      N/P      N/P   FALL TO SEQUENTIAL-NO ESA       0.00     0.00      N/C     0.00
RESYNCHRONIZATION ATTEMPTED      N/P      N/P      N/P      N/P   FALL TO SEQUENTIAL-NO BUFF.     0.00     0.00      N/C     0.00
RESYNCHRONIZATION SUCCEEDED      N/P      N/P      N/P      N/P   FALL TO SEQUENTIAL-ENCL.SER     0.00     0.00      N/C     0.00
```

Figure 66. Run 2WP: 2-Way Second CPC, Locking Activity

A.5 DB2 PM Statistics Report: 2-Way First CPC, BP5 General

```
   LOCATION: DSNDB0G                              DB2 PERFORMANCE MONITOR (V4)                                 PAGE: 1-16
      GROUP: DSNDB0G                                STATISTICS REPORT - LONG                    REQUESTED FROM: 09/09/95 08:41:15.00
     MEMBER: DB1G                                                                                            TO: 09/09/95 08:51:17.00
  SUBSYSTEM: DB1G                                                                                INTERVAL FROM: 09/09/95 08:41:16.36
DB2 VERSION: V4                                           SCOPE: MEMBER                                      TO: 09/09/95 08:51:16.73

---- HIGHLIGHTS ----------------------------------------------------------------------------------------------------
INTERVAL START  : 09/09/95 08:41:16.36   SAMPLING START: 09/09/95 08:41:16.36    TOTAL THREADS      :     0.00
INTERVAL END    : 09/09/95 08:51:16.73   SAMPLING END  : 09/09/95 08:51:16.73    TOTAL COMMITS      : 18056.00
INTERVAL ELAPSED:         10:00.377445   OUTAGE ELAPSED:               0.000000   DATA SHARING MEMBER:      N/A

BP5    GENERAL                  QUANTITY  /MINUTE  /THREAD  /COMMIT   BP5    READ OPERATIONS          QUANTITY  /MINUTE  /THREAD  /COMMIT
---------------------------    --------  -------  -------  -------   ---------------------------    --------  -------  -------  -------
CURRENT ACTIVE BUFFERS            46.00      N/A      N/A      N/A   GETPAGE REQUEST                15157.00  1514.75      N/C     0.84
UNAVAIL.BUFFER-VPOOL FULL          0.00     0.00      N/C     0.00   GETPAGE REQUEST-SEQUENTIAL      4473.00   447.02      N/C     0.25
                                                                     GETPAGE REQUEST-RANDOM         10684.00  1067.73      N/C     0.59
NUMBER OF DATASET OPENS            0.00     0.00      N/C     0.00
                                                                     SYNCHRONOUS READS               4310.00   430.73      N/C     0.24
BUFFERS ALLOCATED - VPOOL       6250.00   624.61      N/C     0.35   SYNCHRON. READS-SEQUENTIAL      4178.00   417.54      N/C     0.23
BUFFERS ALLOCATED - HPOOL          0.00     0.00      N/C     0.00   SYNCHRON. READS-RANDOM           132.00    13.19      N/C     0.01
HPOOL BUFFERS BACKED               0.00     0.00      N/C     0.00
                                                                     GETPAGE PER SYN.READ-RANDOM       80.94
DFHSM MIGRATED DATASET             0.00     0.00      N/C     0.00
DFHSM RECALL TIMEOUTS              0.00     0.00      N/C     0.00   SEQUENTIAL PREFETCH REQUEST     8554.00   854.86      N/C     0.47
                                                                     SEQUENTIAL PREFETCH READS          0.00     0.00      N/C     0.00
HPOOL EXPANS. OR CONTRACT.         0.00     0.00      N/C     0.00   PAGES READ VIA SEQ.PREFETCH        0.00     0.00      N/C     0.00
VPOOL EXPANS. OR CONTRACT.         0.00     0.00      N/C     0.00   S.PRF.PAGES READ/S.PRF.READ         N/C
VPOOL OR HPOOL EXP.FAILURE         0.00     0.00      N/C     0.00
                                                                     LIST PREFETCH REQUESTS             0.00     0.00      N/C     0.00
CONCUR.PREF.I/O STREAMS-HWM        0.00      N/A      N/A      N/A   LIST PREFETCH READS                0.00     0.00      N/C     0.00
PREF.I/O STREAMS REDUCTION         0.00     0.00      N/C     0.00   PAGES READ VIA LIST PREFTCH        0.00     0.00      N/C     0.00
PARALLEL QUERY REQUESTS            0.00     0.00      N/C     0.00   L.PRF.PAGES READ/L.PRF.READ         N/C
PARALL.QUERY REQ.REDUCTION         0.00     0.00      N/C     0.00
PREF.QUANT.REDUCED TO 1/2          0.00     0.00      N/C     0.00   DYNAMIC PREFETCH REQUESTED         0.00     0.00      N/C     0.00
PREF.QUANT.REDUCED TO 1/4          0.00     0.00      N/C     0.00   DYNAMIC PREFETCH READS             0.00     0.00      N/C     0.00
                                                                     PAGES READ VIA DYN.PREFETCH        0.00     0.00      N/C     0.00
BUFFERPOOL EXPANSIONS               N/A      N/A      N/A      N/A   D.PRF.PAGES READ/D.PRF.READ         N/C

                                                                     PREF.DISABLED-NO BUFFER         8554.00   854.86      N/C     0.47
                                                                     PREF.DISABLED-NO READ ENG          0.00     0.00      N/C     0.00

                                                                     SYNC.HPOOL READ                    0.00     0.00      N/C     0.00
                                                                     ASYNC.HPOOL READ                   0.00     0.00      N/C     0.00
                                                                     HPOOL READ FAILED                  0.00     0.00      N/C     0.00
                                                                     ASYN.DA.MOVER HPOOL READ-S         0.00     0.00      N/C     0.00
                                                                     ASYN.DA.MOVER HPOOL READ-F         0.00     0.00      N/C     0.00

                                                                     PAGE-INS REQUIRED FOR READ         0.00     0.00      N/C     0.00
```

Figure 67. Run 2WP: 2-Way First CPC, BP5 General

A.6 DB2 PM Statistics Report: 2-Way Second CPC, BP5 General

```
  LOCATION: DSNDB0G                              DB2 PERFORMANCE MONITOR (V4)                                   PAGE: 1-16
     GROUP: DSNDB0G                               STATISTICS REPORT - LONG                      REQUESTED FROM: 09/09/95 08:41:15.00
    MEMBER: DB2G                                                                                           TO: 09/09/95 08:51:17.00
 SUBSYSTEM: DB2G                                                                                 INTERVAL FROM: 09/09/95 08:41:16.41
DB2 VERSION: V4                                        SCOPE: MEMBER                                          TO: 09/09/95 08:51:16.85

---- HIGHLIGHTS ----------------------------------------------------------------------------------------------------
INTERVAL START  : 09/09/95 08:41:16.41   SAMPLING START: 09/09/95 08:41:16.41   TOTAL THREADS     :     0.00
INTERVAL END    : 09/09/95 08:51:16.85   SAMPLING END  : 09/09/95 08:51:16.85   TOTAL COMMITS     : 18191.00
INTERVAL ELAPSED:         10:00.441800   OUTAGE ELAPSED:               0.000000   DATA SHARING MEMBER:     N/A

BP5    GENERAL                QUANTITY  /MINUTE  /THREAD  /COMMIT   BP5    READ OPERATIONS        QUANTITY  /MINUTE  /THREAD  /COMMIT
---------------------------   --------  -------  -------  -------   ---------------------------   --------  -------  -------  -------
CURRENT ACTIVE BUFFERS           29.00      N/A      N/A      N/A   GETPAGE REQUEST               16459.00  1644.69      N/C     0.90
UNAVAIL.BUFFER-VPOOL FULL         0.00     0.00      N/C     0.00   GETPAGE REQUEST-SEQUENTIAL     4460.00   445.67      N/C     0.25
                                                                    GETPAGE REQUEST-RANDOM        11999.00  1199.02      N/C     0.66
NUMBER OF DATASET OPENS           0.00     0.00      N/C     0.00
                                                                    SYNCHRONOUS READS              4342.00   433.88      N/C     0.24
BUFFERS ALLOCATED - VPOOL      6250.00   624.54      N/C     0.34   SYNCHRON. READS-SEQUENTIAL     4184.00   418.09      N/C     0.23
BUFFERS ALLOCATED - HPOOL         0.00     0.00      N/C     0.00   SYNCHRON. READS-RANDOM          158.00    15.79      N/C     0.01
HPOOL BUFFERS BACKED              0.00     0.00      N/C     0.00
                                                                    GETPAGE PER SYN.READ-RANDOM      75.94
DFHSM MIGRATED DATASET            0.00     0.00      N/C     0.00
DFHSM RECALL TIMEOUTS             0.00     0.00      N/C     0.00   SEQUENTIAL PREFETCH REQUEST    8460.00   845.38      N/C     0.47
                                                                    SEQUENTIAL PREFETCH READS         0.00     0.00      N/C     0.00
HPOOL EXPANS. OR CONTRACT.        0.00     0.00      N/C     0.00   PAGES READ VIA SEQ.PREFETCH       0.00     0.00      N/C     0.00
VPOOL EXPANS. OR CONTRACT.        0.00     0.00      N/C     0.00   S.PRF.PAGES READ/S.PRF.READ        N/C
VPOOL OR HPOOL EXP.FAILURE        0.00     0.00      N/C     0.00
                                                                    LIST PREFETCH REQUESTS            0.00     0.00      N/C     0.00
CONCUR.PREF.I/O STREAMS-HWM       0.00      N/A      N/A      N/A   LIST PREFETCH READS               0.00     0.00      N/C     0.00
PREF.I/O STREAMS REDUCTION        0.00     0.00      N/C     0.00   PAGES READ VIA LIST PREFTCH       0.00     0.00      N/C     0.00
PARALLEL QUERY REQUESTS           0.00     0.00      N/C     0.00   L.PRF.PAGES READ/L.PRF.READ        N/C
PARALL.QUERY REQ.REDUCTION        0.00     0.00      N/C     0.00
PREF.QUANT.REDUCED TO 1/2         0.00     0.00      N/C     0.00   DYNAMIC PREFETCH REQUESTED        0.00     0.00      N/C     0.00
PREF.QUANT.REDUCED TO 1/4         0.00     0.00      N/C     0.00   DYNAMIC PREFETCH READS            0.00     0.00      N/C     0.00
                                                                    PAGES READ VIA DYN.PREFETCH       0.00     0.00      N/C     0.00
BUFFERPOOL EXPANSIONS              N/A      N/A      N/A      N/A   D.PRF.PAGES READ/D.PRF.READ        N/C

                                                                    PREF.DISABLED-NO BUFFER        8460.00   845.38      N/C     0.47
                                                                    PREF.DISABLED-NO READ ENG         0.00     0.00      N/C     0.00

                                                                    SYNC.HPOOL READ                   0.00     0.00      N/C     0.00
                                                                    ASYNC.HPOOL READ                  0.00     0.00      N/C     0.00
                                                                    HPOOL READ FAILED                 0.00     0.00      N/C     0.00
                                                                    ASYN.DA.MOVER HPOOL READ-S        0.00     0.00      N/C     0.00
                                                                    ASYN.DA.MOVER HPOOL READ-F        0.00     0.00      N/C     0.00

                                                                    PAGE-INS REQUIRED FOR READ        0.00     0.00      N/C     0.00
```

Figure 68. Run 2WP: 2-Way Second CPC, BP5 General

A.7 DB2 PM Statistics Report: 2-Way First CPC, Group Buffer Pools GBP7

```
   LOCATION: DSNDB0G                                DB2 PERFORMANCE MONITOR (V4)                                PAGE: 1-29
      GROUP: DSNDB0G                                  STATISTICS REPORT - LONG                       REQUESTED FROM: 09/09/95 08:41:15.00
     MEMBER: DB1G                                                                                                TO: 09/09/95 08:51:17.00
  SUBSYSTEM: DB1G                                                                                     INTERVAL FROM: 09/09/95 08:41:16.36
DB2 VERSION: V4                                             SCOPE: MEMBER                                        TO: 09/09/95 08:51:16.73

---- HIGHLIGHTS ----------------------------------------------------------------------------------------------------
INTERVAL START  : 09/09/95 08:41:16.36   SAMPLING START: 09/09/95 08:41:16.36   TOTAL THREADS     :     0.00
INTERVAL END    : 09/09/95 08:51:16.73   SAMPLING END  : 09/09/95 08:51:16.73   TOTAL COMMITS     : 18056.00
INTERVAL ELAPSED:        10:00.377445    OUTAGE ELAPSED:              0.000000   DATA SHARING MEMBER:      N/A

GROUP BP7                     QUANTITY  /MINUTE  /THREAD  /COMMIT   GROUP BP8                     QUANTITY  /MINUTE  /THREAD  /COMMIT
---------------------------   --------  -------  -------  -------   ---------------------------   --------  -------  -------  -------
SYN.READS(XI)-DATA RETURNED    1388.00   138.71      N/C     0.08   SYN.READS(XI)-DATA RETURNED    2413.00   241.15      N/C     0.13
SYN.READS(XI)-R/W INTEREST        0.00     0.00      N/C     0.00   SYN.READS(XI)-R/W INTEREST        0.00     0.00      N/C     0.00
SYN.READS(XI)-NO R/W INTER.       0.00     0.00      N/C     0.00   SYN.READS(XI)-NO R/W INTER.       0.00     0.00      N/C     0.00
SYN.READS(NF)-DATA RETURNED     845.00    84.45      N/C     0.05   SYN.READS(NF)-DATA RETURNED    2998.00   299.61      N/C     0.17
SYN.READS(NF)-R/W INTEREST    14828.00  1481.87      N/C     0.82   SYN.READS(NF)-R/W INTEREST    35271.00  3524.88      N/C     1.95
SYN.READS(NF)-NO R/W INTER.       0.00     0.00      N/C     0.00   SYN.READS(NF)-NO R/W INTER.       0.00     0.00      N/C     0.00
ASYNC.READS-DATA RETURNED         0.00     0.00      N/C     0.00   ASYNC.READS-DATA RETURNED         0.00     0.00      N/C     0.00
ASYNC.READS-R/W INTEREST          0.00     0.00      N/C     0.00   ASYNC.READS-R/W INTEREST          0.00     0.00      N/C     0.00
ASYNC.READS-NO R/W INTEREST       0.00     0.00      N/C     0.00   ASYNC.READS-NO R/W INTEREST       0.00     0.00      N/C     0.00

CLEAN PAGES SYNC.WRITTEN          0.00     0.00      N/C     0.00   CLEAN PAGES SYNC.WRITTEN          0.00     0.00      N/C     0.00
CHANGED PAGES SYNC.WRITTEN    10330.00  1032.35      N/C     0.57   CHANGED PAGES SYNC.WRITTEN    45510.00  4548.14      N/C     2.52
CLEAN PAGES ASYNC.WRITTEN         0.00     0.00      N/C     0.00   CLEAN PAGES ASYNC.WRITTEN         0.00     0.00      N/C     0.00
CHANGED PAGES ASYNC.WRITTEN       0.00     0.00      N/C     0.00   CHANGED PAGES ASYNC.WRITTEN       1.00     0.10      N/C     0.00

PAGES CASTOUT                  7673.00   766.82      N/C     0.42   PAGES CASTOUT                 46414.00  4638.48      N/C     2.57
CASTOUT CLASS THRESHOLD           0.00     0.00      N/C     0.00   CASTOUT CLASS THRESHOLD           0.00     0.00      N/C     0.00
GROUP BP CASTOUT THRESHOLD       26.00     2.60      N/C     0.00   GROUP BP CASTOUT THRESHOLD       20.00     2.00      N/C     0.00
CASTOUT ENGINE NOT AVAIL.         0.00     0.00      N/C     0.00   CASTOUT ENGINE NOT AVAIL.         0.00     0.00      N/C     0.00

WRITE ENGINE NOT AVAILABLE        0.00     0.00      N/C     0.00   WRITE ENGINE NOT AVAILABLE        0.00     0.00      N/C     0.00
READ FAILED-NO STORAGE            0.00     0.00      N/C     0.00   READ FAILED-NO STORAGE            0.00     0.00      N/C     0.00
WRITE FAILED-NO STORAGE           0.00     0.00      N/C     0.00   WRITE FAILED-NO STORAGE           0.00     0.00      N/C     0.00

OTHER REQUESTS                 5308.00   530.47      N/C     0.29   OTHER REQUESTS                18396.00  1838.44      N/C     1.02
```

Figure 69. Run 2WP: 2-Way First CPC, Group Buffer Pools GBP7

A.8 DB2 PM Statistics Report: 2-Way Second CPC, Group Buffer Pools GBP7

```
   LOCATION: DSNDB0G                            DB2 PERFORMANCE MONITOR (V4)                            PAGE: 1-29
      GROUP: DSNDB0G                             STATISTICS REPORT - LONG                   REQUESTED FROM: 09/09/95 08:41:15.00
     MEMBER: DB2G                                                                                      TO: 09/09/95 08:51:17.00
  SUBSYSTEM: DB2G                                                                           INTERVAL FROM: 09/09/95 08:41:16.41
DB2 VERSION: V4                                       SCOPE: MEMBER                                    TO: 09/09/95 08:51:16.85

---- HIGHLIGHTS ----------------------------------------------------------------------------------------------------
INTERVAL START  : 09/09/95 08:41:16.41   SAMPLING START: 09/09/95 08:41:16.41   TOTAL THREADS     :     0.00
INTERVAL END    : 09/09/95 08:51:16.85   SAMPLING END  : 09/09/95 08:51:16.85   TOTAL COMMITS     : 18191.00
INTERVAL ELAPSED:         10:00.441800   OUTAGE ELAPSED:             0.000000   DATA SHARING MEMBER:     N/A

GROUP BP7                     QUANTITY  /MINUTE  /THREAD  /COMMIT    GROUP BP8                     QUANTITY  /MINUTE  /THREAD  /COMMIT
---------------------------   --------  -------  -------  -------    ---------------------------   --------  -------  -------  -------
SYN.READS(XI)-DATA RETURNED    1316.00   131.50      N/C     0.07    SYN.READS(XI)-DATA RETURNED    2424.00   242.22      N/C     0.13
SYN.READS(XI)-R/W INTEREST        0.00     0.00      N/C     0.00    SYN.READS(XI)-R/W INTEREST        0.00     0.00      N/C     0.00
SYN.READS(XI)-NO R/W INTER.       0.00     0.00      N/C     0.00    SYN.READS(XI)-NO R/W INTER.       0.00     0.00      N/C     0.00
SYN.READS(NF)-DATA RETURNED     971.00    97.03      N/C     0.05    SYN.READS(NF)-DATA RETURNED    2932.00   292.98      N/C     0.16
SYN.READS(NF)-R/W INTEREST    16121.00  1610.91      N/C     0.89    SYN.READS(NF)-R/W INTEREST    34837.00  3481.14      N/C     1.92
SYN.READS(NF)-NO R/W INTER.       0.00     0.00      N/C     0.00    SYN.READS(NF)-NO R/W INTER.       0.00     0.00      N/C     0.00
ASYNC.READS-DATA RETURNED         0.00     0.00      N/C     0.00    ASYNC.READS-DATA RETURNED         0.00     0.00      N/C     0.00
ASYNC.READS-R/W INTEREST          0.00     0.00      N/C     0.00    ASYNC.READS-R/W INTEREST          0.00     0.00      N/C     0.00
ASYNC.READS-NO R/W INTEREST       0.00     0.00      N/C     0.00    ASYNC.READS-NO R/W INTEREST       0.00     0.00      N/C     0.00

CLEAN PAGES SYNC.WRITTEN          0.00     0.00      N/C     0.00    CLEAN PAGES SYNC.WRITTEN          0.00     0.00      N/C     0.00
CHANGED PAGES SYNC.WRITTEN    11790.00  1178.13      N/C     0.65    CHANGED PAGES SYNC.WRITTEN    45246.00  4521.27      N/C     2.49
CLEAN PAGES ASYNC.WRITTEN         0.00     0.00      N/C     0.00    CLEAN PAGES ASYNC.WRITTEN         0.00     0.00      N/C     0.00
CHANGED PAGES ASYNC.WRITTEN       0.00     0.00      N/C     0.00    CHANGED PAGES ASYNC.WRITTEN       0.00     0.00      N/C     0.00

PAGES CASTOUT                  7193.00   718.77      N/C     0.40    PAGES CASTOUT                 31073.00  3105.01      N/C     1.71
CASTOUT CLASS THRESHOLD           0.00     0.00      N/C     0.00    CASTOUT CLASS THRESHOLD           0.00     0.00      N/C     0.00
GROUP BP CASTOUT THRESHOLD        0.00     0.00      N/C     0.00    GROUP BP CASTOUT THRESHOLD        0.00     0.00      N/C     0.00
CASTOUT ENGINE NOT AVAIL.         0.00     0.00      N/C     0.00    CASTOUT ENGINE NOT AVAIL.         0.00     0.00      N/C     0.00

WRITE ENGINE NOT AVAILABLE        0.00     0.00      N/C     0.00    WRITE ENGINE NOT AVAILABLE        0.00     0.00      N/C     0.00
READ FAILED-NO STORAGE            0.00     0.00      N/C     0.00    READ FAILED-NO STORAGE            0.00     0.00      N/C     0.00
WRITE FAILED-NO STORAGE           0.00     0.00      N/C     0.00    WRITE FAILED-NO STORAGE           0.00     0.00      N/C     0.00

OTHER REQUESTS                 4496.00   449.27      N/C     0.25    OTHER REQUESTS                11324.00  1131.57      N/C     0.62
```

Figure 70. Run 2WP: 2-Way Second CPC, Group Buffer Pools GBP7

A.9 DB2 PM Statistics Report: 2-Way First CPC, Group Buffer Pool Totals

```
   LOCATION: DSNDB0G                                DB2 PERFORMANCE MONITOR (V4)                             PAGE: 1-30
      GROUP: DSNDB0G                                  STATISTICS REPORT - LONG                      REQUESTED FROM: 09/09/95 08:41:15.00
     MEMBER: DB1G                                                                                               TO: 09/09/95 08:51:17.00
  SUBSYSTEM: DB1G                                                                                   INTERVAL FROM: 09/09/95 08:41:16.36
DB2 VERSION: V4                                            SCOPE: MEMBER                                        TO: 09/09/95 08:51:16.73

---- HIGHLIGHTS ----------------------------------------------------------------------------------------------------
INTERVAL START  : 09/09/95 08:41:16.36   SAMPLING START: 09/09/95 08:41:16.36   TOTAL THREADS      :     0.00
INTERVAL END    : 09/09/95 08:51:16.73   SAMPLING END  : 09/09/95 08:51:16.73   TOTAL COMMITS      : 18056.00
INTERVAL ELAPSED:         10:00.377445   OUTAGE ELAPSED:               0.000000   DATA SHARING MEMBER:      N/A

GROUP TOT4K                  QUANTITY  /MINUTE  /THREAD  /COMMIT
---------------------------  --------  -------  -------  -------
SYN.READS(XI)-DATA RETURNED  22279.00  2226.50      N/C     1.23
SYN.READS(XI)-R/W INTEREST      10.00     1.00      N/C     0.00
SYN.READS(XI)-NO R/W INTER.      0.00     0.00      N/C     0.00
SYN.READS(NF)-DATA RETURNED  27537.00  2751.97      N/C     1.53
SYN.READS(NF)-R/W INTEREST   81112.00  8106.10      N/C     4.49
SYN.READS(NF)-NO R/W INTER.      0.00     0.00      N/C     0.00
ASYNC.READS-DATA RETURNED        0.00     0.00      N/C     0.00
ASYNC.READS-R/W INTEREST         0.00     0.00      N/C     0.00
ASYNC.READS-NO R/W INTEREST      0.00     0.00      N/C     0.00

CLEAN PAGES SYNC.WRITTEN         0.00     0.00      N/C     0.00
CHANGED PAGES SYNC.WRITTEN     110.7K    11.1K      N/C     6.13
CLEAN PAGES ASYNC.WRITTEN        0.00     0.00      N/C     0.00
CHANGED PAGES ASYNC.WRITTEN     56.00     5.60      N/C     0.00

PAGES CASTOUT                61901.00  6186.21      N/C     3.43
CASTOUT CLASS THRESHOLD          0.00     0.00      N/C     0.00
GROUP BP CASTOUT THRESHOLD      64.00     6.40      N/C     0.00
CASTOUT ENGINE NOT AVAIL.        0.00     0.00      N/C     0.00

WRITE ENGINE NOT AVAILABLE       0.00     0.00      N/C     0.00
READ FAILED-NO STORAGE           0.00     0.00      N/C     0.00
WRITE FAILED-NO STORAGE          0.00     0.00      N/C     0.00

OTHER REQUESTS               28564.00  2854.60      N/C     1.58

STATISTICS REPORT COMPLETE
```

Figure 71. Run 2WP: 2-Way First CPC, Group Buffer Pool Totals

A.10 DB2 PM Statistics Report: 2-Way Second CPC, Group Buffer Pool Totals

```
   LOCATION: DSNDB0G                              DB2 PERFORMANCE MONITOR (V4)                              PAGE: 1-30
      GROUP: DSNDB0G                               STATISTICS REPORT - LONG                     REQUESTED FROM: 09/09/95 08:41:15.00
     MEMBER: DB2G                                                                                           TO: 09/09/95 08:51:17.00
  SUBSYSTEM: DB2G                                                                                INTERVAL FROM: 09/09/95 08:41:16.41
DB2 VERSION: V4                                          SCOPE: MEMBER                                      TO: 09/09/95 08:51:16.85

---- HIGHLIGHTS ----------------------------------------------------------------------------------------------------
INTERVAL START  : 09/09/95 08:41:16.41    SAMPLING START: 09/09/95 08:41:16.41    TOTAL THREADS      :     0.00
INTERVAL END    : 09/09/95 08:51:16.85    SAMPLING END  : 09/09/95 08:51:16.85    TOTAL COMMITS      : 18191.00
INTERVAL ELAPSED:         10:00.441800    OUTAGE ELAPSED:               0.000000  DATA SHARING MEMBER:      N/A

GROUP TOT4K                    QUANTITY  /MINUTE  /THREAD  /COMMIT
---------------------------    --------  -------  -------  -------
SYN.READS(XI)-DATA RETURNED    22291.00  2227.46      N/C     1.23
SYN.READS(XI)-R/W INTEREST        15.00     1.50      N/C     0.00
SYN.READS(XI)-NO R/W INTER.        0.00     0.00      N/C     0.00
SYN.READS(NF)-DATA RETURNED    29591.00  2956.92      N/C     1.63
SYN.READS(NF)-R/W INTEREST     88606.00  8854.08      N/C     4.87
SYN.READS(NF)-NO R/W INTER.        0.00     0.00      N/C     0.00
ASYNC.READS-DATA RETURNED          0.00     0.00      N/C     0.00
ASYNC.READS-R/W INTEREST           0.00     0.00      N/C     0.00
ASYNC.READS-NO R/W INTEREST        0.00     0.00      N/C     0.00

CLEAN PAGES SYNC.WRITTEN           0.00     0.00      N/C     0.00
CHANGED PAGES SYNC.WRITTEN       117.7K    11.8K      N/C     6.47
CLEAN PAGES ASYNC.WRITTEN          0.00     0.00      N/C     0.00
CHANGED PAGES ASYNC.WRITTEN       40.00     4.00      N/C     0.00

PAGES CASTOUT                  45299.00  4526.57      N/C     2.49
CASTOUT CLASS THRESHOLD            0.00     0.00      N/C     0.00
GROUP BP CASTOUT THRESHOLD         0.00     0.00      N/C     0.00
CASTOUT ENGINE NOT AVAIL.          0.00     0.00      N/C     0.00

WRITE ENGINE NOT AVAILABLE         0.00     0.00      N/C     0.00
READ FAILED-NO STORAGE             0.00     0.00      N/C     0.00
WRITE FAILED-NO STORAGE            0.00     0.00      N/C     0.00

OTHER REQUESTS                 19845.00  1983.04      N/C     1.09

STATISTICS REPORT COMPLETE
```

Figure 72. Run 2WP: 2-Way Second CPC, Group Buffer Pool Totals

Appendix B. DB2 PM 1-Way and 2-Way Statistics Reports (1WP and 2WP)

This appendix compares DB2 PM Statistic reports for the 1-way measurements and the first CPC for the 2-way measurements.

B.1 DB2 PM Statistics Report: 1-Way CPC, SQL Activity

LOCATION: DSNDB0G	DB2 PERFORMANCE MONITOR (V4)	PAGE: 1-1
GROUP: DSNDB0G	STATISTICS REPORT - LONG	REQUESTED FROM: 09/09/95 12:47:18.00
MEMBER: DB1G		TO: 09/09/95 12:57:20.00
SUBSYSTEM: DB1G		INTERVAL FROM: 09/09/95 12:47:19.27
DB2 VERSION: V4	SCOPE: MEMBER	TO: 09/09/95 12:57:19.33

---- HIGHLIGHTS --

INTERVAL START : 09/09/95 12:47:19.27	SAMPLING START: 09/09/95 12:47:19.27	TOTAL THREADS : 0.00
INTERVAL END : 09/09/95 12:57:19.33	SAMPLING END : 09/09/95 12:57:19.33	TOTAL COMMITS : 21493.00
INTERVAL ELAPSED: 10:00.063955	OUTAGE ELAPSED: 0.000000	DATA SHARING MEMBER: N/A

SQL DML	QUANTITY	/MINUTE	/THREAD	/COMMIT
SELECT	104.9K	10.5K	N/C	4.88
INSERT	62058.00	6205.14	N/C	2.89
UPDATE	80418.00	8040.94	N/C	3.74
DELETE	4503.00	450.25	N/C	0.21
PREPARE	0.00	0.00	N/C	0.00
DESCRIBE	0.00	0.00	N/C	0.00
DESCRIBE TABLE	0.00	0.00	N/C	0.00
OPEN	98754.00	9874.35	N/C	4.59
CLOSE	56799.00	5679.29	N/C	2.64
FETCH	199.5K	19.9K	N/C	9.28
TOTAL	606.9K	60.7K	N/C	28.24

SQL DCL	QUANTITY	/MINUTE	/THREAD	/COMMIT
LOCK TABLE	0.00	0.00	N/C	0.00
GRANT	0.00	0.00	N/C	0.00
REVOKE	0.00	0.00	N/C	0.00
SET HOST VARIABLE	0.00	0.00	N/C	0.00
SET CURRENT SQLID	0.00	0.00	N/C	0.00
SET CURRENT DEGREE	0.00	0.00	N/C	0.00
SET CURRENT RULES	0.00	0.00	N/C	0.00
CONNECT TYPE 1	0.00	0.00	N/C	0.00
CONNECT TYPE 2	0.00	0.00	N/C	0.00
RELEASE	0.00	0.00	N/C	0.00
SET CONNECTION	0.00	0.00	N/C	0.00
TOTAL	0.00	0.00	N/C	0.00

SQL DDL	QUANTITY	/MINUTE	/THREAD	/COMMIT
CREATE TABLE	0.00	0.00	N/C	0.00
CREATE INDEX	0.00	0.00	N/C	0.00
CREATE VIEW	0.00	0.00	N/C	0.00
CREATE SYNONYM	0.00	0.00	N/C	0.00
CREATE TABLESPACE	0.00	0.00	N/C	0.00
CREATE DATABASE	0.00	0.00	N/C	0.00
CREATE STOGROUP	0.00	0.00	N/C	0.00
CREATE ALIAS	0.00	0.00	N/C	0.00
ALTER TABLE	0.00	0.00	N/C	0.00
ALTER INDEX	0.00	0.00	N/C	0.00
ALTER TABLESPACE	0.00	0.00	N/C	0.00
ALTER DATABASE	0.00	0.00	N/C	0.00
ALTER STOGROUP	0.00	0.00	N/C	0.00
DROP TABLE	0.00	0.00	N/C	0.00
DROP INDEX	0.00	0.00	N/C	0.00
DROP VIEW	0.00	0.00	N/C	0.00
DROP SYNONYM	0.00	0.00	N/C	0.00
DROP TABLESPACE	0.00	0.00	N/C	0.00
DROP DATABASE	0.00	0.00	N/C	0.00
DROP STOGROUP	0.00	0.00	N/C	0.00
DROP ALIAS	0.00	0.00	N/C	0.00
DROP PACKAGE	0.00	0.00	N/C	0.00
COMMENT ON	0.00	0.00	N/C	0.00
LABEL ON	0.00	0.00	N/C	0.00
TOTAL	0.00	0.00	N/C	0.00

STORED PROCEDURES	QUANTITY	/MINUTE	/THREAD	/COMMIT
CALL STATEMENTS EXECUTED	0.00	0.00	N/C	0.00
PROCEDURE ABENDS	0.00	0.00	N/C	0.00
CALL STATEMENT TIMEOUTS	0.00	0.00	N/C	0.00
CALL STATEMENT REJECTED	0.00	0.00	N/C	0.00

Figure 73. Run 1WP: 1-Way CPC, SQL Activity

B.2 DB2 PM Statistics Report: 2-Way First CPC, SQL Activity

```
   LOCATION: DSNDB0G                               DB2 PERFORMANCE MONITOR (V4)                              PAGE: 1-1
      GROUP: DSNDB0G                                STATISTICS REPORT - LONG                    REQUESTED FROM: 09/09/95 08:41:15.00
     MEMBER: DB1G                                                                                           TO: 09/09/95 08:51:17.00
  SUBSYSTEM: DB1G                                                                                INTERVAL FROM: 09/09/95 08:41:16.36
DB2 VERSION: V4                                           SCOPE: MEMBER                                     TO: 09/09/95 08:51:16.73

---- HIGHLIGHTS ----------------------------------------------------------------------------------------------------
INTERVAL START  : 09/09/95 08:41:16.36    SAMPLING START: 09/09/95 08:41:16.36    TOTAL THREADS      :     0.00
INTERVAL END    : 09/09/95 08:51:16.73    SAMPLING END  : 09/09/95 08:51:16.73    TOTAL COMMITS      : 18056.00
INTERVAL ELAPSED:         10:00.377445    OUTAGE ELAPSED:             0.000000    DATA SHARING MEMBER:      N/A

SQL DML                      QUANTITY  /MINUTE  /THREAD  /COMMIT    SQL DCL                      QUANTITY  /MINUTE  /THREAD  /COMMIT
---------------------------  --------  -------  -------  -------    ---------------------------  --------  -------  -------  -------
SELECT                       86981.00  8692.63      N/C     4.82    LOCK TABLE                       0.00     0.00      N/C     0.00
INSERT                       53011.00  5297.77      N/C     2.94    GRANT                            0.00     0.00      N/C     0.00
UPDATE                       64530.00  6448.94      N/C     3.57    REVOKE                           0.00     0.00      N/C     0.00
DELETE                        2505.00   250.34      N/C     0.14    SET HOST VARIABLE                0.00     0.00      N/C     0.00
                                                                    SET CURRENT SQLID                0.00     0.00      N/C     0.00
PREPARE                          0.00     0.00      N/C     0.00    SET CURRENT DEGREE               0.00     0.00      N/C     0.00
DESCRIBE                         0.00     0.00      N/C     0.00    SET CURRENT RULES                0.00     0.00      N/C     0.00
DESCRIBE TABLE                   0.00     0.00      N/C     0.00
OPEN                         81417.00  8136.58      N/C     4.51    CONNECT TYPE 1                   0.00     0.00      N/C     0.00
CLOSE                        45883.00  4585.42      N/C     2.54    CONNECT TYPE 2                   0.00     0.00      N/C     0.00
FETCH                          166.8K    16.7K      N/C     9.24    RELEASE                          0.00     0.00      N/C     0.00
                                                                    SET CONNECTION                   0.00     0.00      N/C     0.00
TOTAL                          501.2K    50.1K      N/C    27.76
                                                                    TOTAL                            0.00     0.00      N/C     0.00

SQL DDL                      QUANTITY  /MINUTE  /THREAD  /COMMIT    STORED PROCEDURES            QUANTITY  /MINUTE  /THREAD  /COMMIT
---------------------------  --------  -------  -------  -------    ---------------------------  --------  -------  -------  -------
CREATE TABLE                     0.00     0.00      N/C     0.00    CALL STATEMENTS EXECUTED         0.00     0.00      N/C     0.00
CREATE INDEX                     0.00     0.00      N/C     0.00    PROCEDURE ABENDS                 0.00     0.00      N/C     0.00
CREATE VIEW                      0.00     0.00      N/C     0.00    CALL STATEMENT TIMEOUTS          0.00     0.00      N/C     0.00
CREATE SYNONYM                   0.00     0.00      N/C     0.00    CALL STATEMENT REJECTED          0.00     0.00      N/C     0.00
CREATE TABLESPACE                0.00     0.00      N/C     0.00
CREATE DATABASE                  0.00     0.00      N/C     0.00
CREATE STOGROUP                  0.00     0.00      N/C     0.00
CREATE ALIAS                     0.00     0.00      N/C     0.00

ALTER  TABLE                     0.00     0.00      N/C     0.00
ALTER  INDEX                     0.00     0.00      N/C     0.00
ALTER  TABLESPACE                0.00     0.00      N/C     0.00
ALTER  DATABASE                  0.00     0.00      N/C     0.00
ALTER  STOGROUP                  0.00     0.00      N/C     0.00

DROP   TABLE                     0.00     0.00      N/C     0.00
DROP   INDEX                     0.00     0.00      N/C     0.00
DROP   VIEW                      0.00     0.00      N/C     0.00
DROP   SYNONYM                   0.00     0.00      N/C     0.00
DROP   TABLESPACE                0.00     0.00      N/C     0.00
DROP   DATABASE                  0.00     0.00      N/C     0.00
DROP   STOGROUP                  0.00     0.00      N/C     0.00
DROP   ALIAS                     0.00     0.00      N/C     0.00
DROP   PACKAGE                   0.00     0.00      N/C     0.00

COMMENT ON                       0.00     0.00      N/C     0.00
LABEL ON                         0.00     0.00      N/C     0.00

TOTAL                            0.00     0.00      N/C     0.00
```

Figure 74. Run 2WP: 2-Way First CPC, SQL Activity

B.3 DB2 PM Statistics Report: 1-Way CPC, Locking Activity

```
   LOCATION: DSNDB0G                                  DB2 PERFORMANCE MONITOR (V4)                                PAGE: 1-4
      GROUP: DSNDB0G                                    STATISTICS REPORT - LONG                      REQUESTED FROM: 09/09/95 12:47:18.00
     MEMBER: DB1G                                                                                                 TO: 09/09/95 12:57:20.00
  SUBSYSTEM: DB1G                                                                                     INTERVAL FROM: 09/09/95 12:47:19.27
DB2 VERSION: V4                                              SCOPE: MEMBER                                        TO: 09/09/95 12:57:19.33

---- HIGHLIGHTS ----------------------------------------------------------------------------------------------------
INTERVAL START  : 09/09/95 12:47:19.27    SAMPLING START: 09/09/95 12:47:19.27    TOTAL THREADS      :     0.00
INTERVAL END    : 09/09/95 12:57:19.33    SAMPLING END  : 09/09/95 12:57:19.33    TOTAL COMMITS      : 21493.00
INTERVAL ELAPSED:         10:00.063955    OUTAGE ELAPSED:                0.000000  DATA SHARING MEMBER:      N/A

LOCKING ACTIVITY             QUANTITY  /MINUTE  /THREAD  /COMMIT    DATA SHARING LOCKING         QUANTITY  /MINUTE  /THREAD  /COMMIT
---------------------------  --------  -------  -------  -------    ---------------------------  --------  -------  -------  -------
SUSPENSIONS (ALL)             2399.00   239.87      N/C     0.11    LOCK REQUESTS (P-LOCKS)          0.00     0.00      N/C     0.00
SUSPENSIONS (LOCK ONLY)        100.00    10.00      N/C     0.00    UNLOCK REQUESTS (P-LOCKS)        0.00     0.00      N/C     0.00
SUSPENSIONS (LATCH ONLY)      2299.00   229.88      N/C     0.11    CHANGE REQUESTS (P-LOCKS)        0.00     0.00      N/C     0.00
SUSPENSIONS (OTHER)              0.00     0.00      N/C     0.00
                                                                    SYNCH.XES - LOCK REQUESTS      857.00    85.69      N/C     0.04
TIMEOUTS                         0.00     0.00      N/C     0.00    SYNCH.XES - CHANGE REQUESTS      0.00     0.00      N/C     0.00
DEADLOCKS                        0.00     0.00      N/C     0.00    SYNCH.XES - UNLOCK REQUESTS    856.00    85.59      N/C     0.04
                                                                    ASYNCH.XES - RESOURCES           0.00     0.00      N/C     0.00
LOCK REQUESTS                  142.5K    14.3K      N/C     6.63
UNLOCK REQUESTS              32810.00  3280.65      N/C     1.53    SUSPENDS - IRLM GLOBAL CONT      0.00     0.00      N/C     0.00
QUERY REQUESTS                   0.00     0.00      N/C     0.00    SUSPENDS - XES GLOBAL CONT.      0.00     0.00      N/C     0.00
CHANGE REQUESTS              81741.00  8173.23      N/C     3.80    SUSPENDS - FALSE CONTENTION      0.00     0.00      N/C     0.00
OTHER REQUESTS                   0.00     0.00      N/C     0.00    INCOMPATIBLE RETAINED LOCK       0.00     0.00      N/C     0.00

LOCK ESCALATION (SHARED)         0.00     0.00      N/C     0.00    NOTIFY MESSAGES SENT             0.00     0.00      N/C     0.00
LOCK ESCALATION (EXCLUSIVE)      0.00     0.00      N/C     0.00    NOTIFY MESSAGES RECEIVED         0.00     0.00      N/C     0.00
                                                                    P-LOCK/NOTIFY EXITS ENGINES     10.00      N/A      N/A      N/A
DRAIN REQUESTS                   0.00     0.00      N/C     0.00    P-LCK/NFY EX.ENGINE UNAVAIL      0.00     0.00      N/C     0.00
DRAIN REQUESTS FAILED            0.00     0.00      N/C     0.00
CLAIM REQUESTS                 195.6K    19.6K      N/C     9.10    PSET/PART P-LCK NEGOTIATION      0.00     0.00      N/C     0.00
CLAIM REQUESTS FAILED            0.00     0.00      N/C     0.00    PAGE P-LOCK NEGOTIATION          0.00     0.00      N/C     0.00
                                                                    OTHER P-LOCK NEGOTIATION         0.00     0.00      N/C     0.00
                                                                    P-LOCK CHANGE DURING NEG.        0.00     0.00      N/C     0.00

GLOBAL DDF ACTIVITY          QUANTITY  /MINUTE  /THREAD  /COMMIT    QUERY PARALLELISM            QUANTITY  /MINUTE  /THREAD  /COMMIT
---------------------------  --------  -------  -------  -------    ---------------------------  --------  -------  -------  -------
DBAT QUEUED-MAXIMUM ACTIVE        N/P      N/P      N/P      N/A    MAX.DEGREE OF PARALLELISM        0.00      N/A      N/A      N/A
CONV.DEALLOC-MAX.CONNECTED        N/P      N/P      N/P      N/A    PARALLEL GROUPS EXECUTED         0.00     0.00      N/C     0.00
INACTIVE DBATS - CURRENTLY        N/P      N/A      N/A      N/A    EXECUTED AS PLANNED              0.00     0.00      N/C     0.00
INACTIVE DBATS - HWM              N/P      N/A      N/A      N/A    REDUCED DEGREE - NO BUFFER       0.00     0.00      N/C     0.00
COLD START CONNECTIONS            N/P      N/P      N/P      N/P    FALL TO SEQUENTIAL-CURSOR        0.00     0.00      N/C     0.00
WARM START CONNECTIONS            N/P      N/P      N/P      N/P    FALL TO SEQUENTIAL-NO ESA        0.00     0.00      N/C     0.00
RESYNCHRONIZATION ATTEMPTED       N/P      N/P      N/P      N/P    FALL TO SEQUENTIAL-NO BUFF.      0.00     0.00      N/C     0.00
RESYNCHRONIZATION SUCCEEDED       N/P      N/P      N/P      N/P    FALL TO SEQUENTIAL-ENCL.SER      0.00     0.00      N/C     0.00
```

Figure 75. Run 1WP: 1-Way CPC, Locking Activity

B.4 DB2 PM Statistics Report: 2-Way First CPC, Locking Activity

```
  LOCATION: DSNDB0G                    DB2 PERFORMANCE MONITOR (V4)                         PAGE: 1-4
     GROUP: DSNDB0G                      STATISTICS REPORT - LONG               REQUESTED FROM: 09/09/95 08:41:15.00
    MEMBER: DB1G                                                                            TO: 09/09/95 08:51:17.00
 SUBSYSTEM: DB1G                                                                 INTERVAL FROM: 09/09/95 08:41:16.36
DB2 VERSION: V4                               SCOPE: MEMBER                                 TO: 09/09/95 08:51:16.73

---- HIGHLIGHTS ------------------------------------------------------------------------------------------------------
INTERVAL START  : 09/09/95 08:41:16.36   SAMPLING START: 09/09/95 08:41:16.36   TOTAL THREADS      :     0.00
INTERVAL END    : 09/09/95 08:51:16.73   SAMPLING END  : 09/09/95 08:51:16.73   TOTAL COMMITS      : 18056.00
INTERVAL ELAPSED:         10:00.377445   OUTAGE ELAPSED:              0.000000   DATA SHARING MEMBER:      N/A
```

LOCKING ACTIVITY	QUANTITY	/MINUTE	/THREAD	/COMMIT
SUSPENSIONS (ALL)	3391.00	338.89	N/C	0.19
SUSPENSIONS (LOCK ONLY)	343.00	34.28	N/C	0.02
SUSPENSIONS (LATCH ONLY)	3048.00	304.61	N/C	0.17
SUSPENSIONS (OTHER)	0.00	0.00	N/C	0.00
TIMEOUTS	0.00	0.00	N/C	0.00
DEADLOCKS	0.00	0.00	N/C	0.00
LOCK REQUESTS	147.5K	14.7K	N/C	8.17
UNLOCK REQUESTS	54317.00	5428.29	N/C	3.01
QUERY REQUESTS	0.00	0.00	N/C	0.00
CHANGE REQUESTS	65204.00	6516.30	N/C	3.61
OTHER REQUESTS	0.00	0.00	N/C	0.00
LOCK ESCALATION (SHARED)	0.00	0.00	N/C	0.00
LOCK ESCALATION (EXCLUSIVE)	0.00	0.00	N/C	0.00
DRAIN REQUESTS	0.00	0.00	N/C	0.00
DRAIN REQUESTS FAILED	0.00	0.00	N/C	0.00
CLAIM REQUESTS	165.8K	16.6K	N/C	9.18
CLAIM REQUESTS FAILED	0.00	0.00	N/C	0.00

DATA SHARING LOCKING	QUANTITY	/MINUTE	/THREAD	/COMMIT
LOCK REQUESTS (P-LOCKS)	25890.00	2587.37	N/C	1.43
UNLOCK REQUESTS (P-LOCKS)	25855.00	2583.87	N/C	1.43
CHANGE REQUESTS (P-LOCKS)	10.00	1.00	N/C	0.00
SYNCH.XES - LOCK REQUESTS	128.9K	12.9K	N/C	7.14
SYNCH.XES - CHANGE REQUESTS	64952.00	6491.12	N/C	3.60
SYNCH.XES - UNLOCK REQUESTS	119.6K	12.0K	N/C	6.62
ASYNCH.XES - RESOURCES	0.00	0.00	N/C	0.00
SUSPENDS - IRLM GLOBAL CONT	621.00	62.06	N/C	0.03
SUSPENDS - XES GLOBAL CONT.	0.00	0.00	N/C	0.00
SUSPENDS - FALSE CONTENTION	1016.00	101.54	N/C	0.06
INCOMPATIBLE RETAINED LOCK	0.00	0.00	N/C	0.00
NOTIFY MESSAGES SENT	692.00	69.16	N/C	0.04
NOTIFY MESSAGES RECEIVED	0.00	0.00	N/C	0.00
P-LOCK/NOTIFY EXITS ENGINES	10.00	N/A	N/A	N/A
P-LCK/NFY EX.ENGINE UNAVAIL	0.00	0.00	N/C	0.00
PSET/PART P-LCK NEGOTIATION	0.00	0.00	N/C	0.00
PAGE P-LOCK NEGOTIATION	388.00	38.78	N/C	0.02
OTHER P-LOCK NEGOTIATION	0.00	0.00	N/C	0.00
P-LOCK CHANGE DURING NEG.	388.00	38.78	N/C	0.02

GLOBAL DDF ACTIVITY	QUANTITY	/MINUTE	/THREAD	/COMMIT
DBAT QUEUED-MAXIMUM ACTIVE	N/P	N/P	N/P	N/A
CONV.DEALLOC-MAX.CONNECTED	N/P	N/P	N/P	N/A
INACTIVE DBATS - CURRENTLY	N/P	N/A	N/A	N/A
INACTIVE DBATS - HWM	N/P	N/A	N/A	N/A
COLD START CONNECTIONS	N/P	N/P	N/P	N/P
WARM START CONNECTIONS	N/P	N/P	N/P	N/P
RESYNCHRONIZATION ATTEMPTED	N/P	N/P	N/P	N/P
RESYNCHRONIZATION SUCCEEDED	N/P	N/P	N/P	N/P

QUERY PARALLELISM	QUANTITY	/MINUTE	/THREAD	/COMMIT
MAX.DEGREE OF PARALLELISM	0.00	N/A	N/A	N/A
PARALLEL GROUPS EXECUTED	0.00	0.00	N/C	0.00
EXECUTED AS PLANNED	0.00	0.00	N/C	0.00
REDUCED DEGREE - NO BUFFER	0.00	0.00	N/C	0.00
FALL TO SEQUENTIAL-CURSOR	0.00	0.00	N/C	0.00
FALL TO SEQUENTIAL-NO ESA	0.00	0.00	N/C	0.00
FALL TO SEQUENTIAL-NO BUFF.	0.00	0.00	N/C	0.00
FALL TO SEQUENTIAL-ENCL.SER	0.00	0.00	N/C	0.00

Figure 76. Run 2WP: 2-Way First CPC, Locking Activity

B.5 DB2 PM Statistics Report: 1-Way CPC, BP5 General

```
   LOCATION: DSNDB0G                           DB2 PERFORMANCE MONITOR (V4)                              PAGE: 1-16
      GROUP: DSNDB0G                            STATISTICS REPORT - LONG                    REQUESTED FROM: 09/09/95 12:47:18.00
     MEMBER: DB1G                                                                                      TO: 09/09/95 12:57:20.00
  SUBSYSTEM: DB1G                                                                            INTERVAL FROM: 09/09/95 12:47:19.27
DB2 VERSION: V4                                      SCOPE: MEMBER                                     TO: 09/09/95 12:57:19.33

---- HIGHLIGHTS ----------------------------------------------------------------------------------------------------
INTERVAL START  : 09/09/95 12:47:19.27   SAMPLING START: 09/09/95 12:47:19.27   TOTAL THREADS     :     0.00
INTERVAL END    : 09/09/95 12:57:19.33   SAMPLING END  : 09/09/95 12:57:19.33   TOTAL COMMITS     : 21493.00
INTERVAL ELAPSED:         10:00.063955   OUTAGE ELAPSED:             0.000000   DATA SHARING MEMBER:     N/A
```

BP5 GENERAL	QUANTITY	/MINUTE	/THREAD	/COMMIT
CURRENT ACTIVE BUFFERS	2563.00	N/A	N/A	N/A
UNAVAIL.BUFFER-VPOOL FULL	0.00	0.00	N/C	0.00
NUMBER OF DATASET OPENS	0.00	0.00	N/C	0.00
BUFFERS ALLOCATED - VPOOL	6250.00	624.93	N/C	0.29
BUFFERS ALLOCATED - HPOOL	0.00	0.00	N/C	0.00
HPOOL BUFFERS BACKED	0.00	0.00	N/C	0.00
DFHSM MIGRATED DATASET	0.00	0.00	N/C	0.00
DFHSM RECALL TIMEOUTS	0.00	0.00	N/C	0.00
HPOOL EXPANS. OR CONTRACT.	0.00	0.00	N/C	0.00
VPOOL EXPANS. OR CONTRACT.	0.00	0.00	N/C	0.00
VPOOL OR HPOOL EXP.FAILURE	0.00	0.00	N/C	0.00
CONCUR.PREF.I/O STREAMS-HWM	0.00	N/A	N/A	N/A
PREF.I/O STREAMS REDUCTION	0.00	0.00	N/C	0.00
PARALLEL QUERY REQUESTS	0.00	0.00	N/C	0.00
PARALL.QUERY REQ.REDUCTION	0.00	0.00	N/C	0.00
PREF.QUANT.REDUCED TO 1/2	0.00	0.00	N/C	0.00
PREF.QUANT.REDUCED TO 1/4	0.00	0.00	N/C	0.00
BUFFERPOOL EXPANSIONS	N/A	N/A	N/A	N/A

BP5 READ OPERATIONS	QUANTITY	/MINUTE	/THREAD	/COMMIT
GETPAGE REQUEST	19249.00	1924.69	N/C	0.90
GETPAGE REQUEST-SEQUENTIAL	5232.00	523.14	N/C	0.24
GETPAGE REQUEST-RANDOM	14017.00	1401.55	N/C	0.65
SYNCHRONOUS READS	4958.00	495.75	N/C	0.23
SYNCHRON. READS-SEQUENTIAL	4956.00	495.55	N/C	0.23
SYNCHRON. READS-RANDOM	2.00	0.20	N/C	0.00
GETPAGE PER SYN.READ-RANDOM	7008.50			
SEQUENTIAL PREFETCH REQUEST	10012.00	1001.09	N/C	0.47
SEQUENTIAL PREFETCH READS	0.00	0.00	N/C	0.00
PAGES READ VIA SEQ.PREFETCH	0.00	0.00	N/C	0.00
S.PRF.PAGES READ/S.PRF.READ	N/C			
LIST PREFETCH REQUESTS	0.00	0.00	N/C	0.00
LIST PREFETCH READS	0.00	0.00	N/C	0.00
PAGES READ VIA LIST PREFTCH	0.00	0.00	N/C	0.00
L.PRF.PAGES READ/L.PRF.READ	N/C			
DYNAMIC PREFETCH REQUESTED	0.00	0.00	N/C	0.00
DYNAMIC PREFETCH READS	0.00	0.00	N/C	0.00
PAGES READ VIA DYN.PREFETCH	0.00	0.00	N/C	0.00
D.PRF.PAGES READ/D.PRF.READ	N/C			
PREF.DISABLED-NO BUFFER	10012.00	1001.09	N/C	0.47
PREF.DISABLED-NO READ ENG	0.00	0.00	N/C	0.00
SYNC.HPOOL READ	0.00	0.00	N/C	0.00
ASYNC.HPOOL READ	0.00	0.00	N/C	0.00
HPOOL READ FAILED	0.00	0.00	N/C	0.00
ASYN.DA.MOVER HPOOL READ-S	0.00	0.00	N/C	0.00
ASYN.DA.MOVER HPOOL READ-F	0.00	0.00	N/C	0.00
PAGE-INS REQUIRED FOR READ	3.00	0.30	N/C	0.00

Figure 77. Run 1WP: 1-Way CPC, BP5 General

B.6 DB2 PM Statistics Report: 2-Way First CPC, BP5 General

```
   LOCATION: DSNDB0G                              DB2 PERFORMANCE MONITOR (V4)                                     PAGE: 1-16
      GROUP: DSNDB0G                               STATISTICS REPORT - LONG                        REQUESTED FROM: 09/09/95 08:41:15.00
     MEMBER: DB1G                                                                                              TO: 09/09/95 08:51:17.00
  SUBSYSTEM: DB1G                                                                                   INTERVAL FROM: 09/09/95 08:41:16.36
DB2 VERSION: V4                                        SCOPE: MEMBER                                           TO: 09/09/95 08:51:16.73

---- HIGHLIGHTS ----------------------------------------------------------------------------------------------------
INTERVAL START   : 09/09/95 08:41:16.36    SAMPLING START: 09/09/95 08:41:16.36    TOTAL THREADS      :      0.00
INTERVAL END     : 09/09/95 08:51:16.73    SAMPLING END  : 09/09/95 08:51:16.73    TOTAL COMMITS      : 18056.00
INTERVAL ELAPSED:        10:00.377445      OUTAGE ELAPSED:             0.000000    DATA SHARING MEMBER:       N/A

BP5    GENERAL                QUANTITY  /MINUTE  /THREAD  /COMMIT    BP5    READ OPERATIONS        QUANTITY  /MINUTE  /THREAD  /COMMIT
---------------------------   --------  -------  -------  -------    ---------------------------   --------  -------  -------  -------
CURRENT ACTIVE BUFFERS           46.00      N/A      N/A      N/A    GETPAGE REQUEST               15157.00  1514.75      N/C     0.84
UNAVAIL.BUFFER-VPOOL FULL         0.00     0.00      N/C     0.00    GETPAGE REQUEST-SEQUENTIAL     4473.00   447.02      N/C     0.25
                                                                     GETPAGE REQUEST-RANDOM        10684.00  1067.73      N/C     0.59
NUMBER OF DATASET OPENS           0.00     0.00      N/C     0.00
                                                                     SYNCHRONOUS READS              4310.00   430.73      N/C     0.24
BUFFERS ALLOCATED - VPOOL      6250.00   624.61      N/C     0.35    SYNCHRON. READS-SEQUENTIAL     4178.00   417.54      N/C     0.23
BUFFERS ALLOCATED - HPOOL         0.00     0.00      N/C     0.00    SYNCHRON. READS-RANDOM          132.00    13.19      N/C     0.01
HPOOL BUFFERS BACKED              0.00     0.00      N/C     0.00
                                                                     GETPAGE PER SYN.READ-RANDOM      80.94
DFHSM MIGRATED DATASET            0.00     0.00      N/C     0.00
DFHSM RECALL TIMEOUTS             0.00     0.00      N/C     0.00    SEQUENTIAL PREFETCH REQUEST    8554.00   854.86      N/C     0.47
                                                                     SEQUENTIAL PREFETCH READS         0.00     0.00      N/C     0.00
HPOOL EXPANS. OR CONTRACT.        0.00     0.00      N/C     0.00    PAGES READ VIA SEQ.PREFETCH       0.00     0.00      N/C     0.00
VPOOL EXPANS. OR CONTRACT.        0.00     0.00      N/C     0.00    S.PRF.PAGES READ/S.PRF.READ        N/C
VPOOL OR HPOOL EXP.FAILURE        0.00     0.00      N/C     0.00
                                                                     LIST PREFETCH REQUESTS            0.00     0.00      N/C     0.00
CONCUR.PREF.I/O STREAMS-HWM       0.00      N/A      N/A      N/A    LIST PREFETCH READS               0.00     0.00      N/C     0.00
PREF.I/O STREAMS REDUCTION        0.00     0.00      N/C     0.00    PAGES READ VIA LIST PREFTCH       0.00     0.00      N/C     0.00
PARALLEL QUERY REQUESTS           0.00     0.00      N/C     0.00    L.PRF.PAGES READ/L.PRF.READ        N/C
PARALL.QUERY REQ.REDUCTION        0.00     0.00      N/C     0.00
PREF.QUANT.REDUCED TO 1/2         0.00     0.00      N/C     0.00    DYNAMIC PREFETCH REQUESTED        0.00     0.00      N/C     0.00
PREF.QUANT.REDUCED TO 1/4         0.00     0.00      N/C     0.00    DYNAMIC PREFETCH READS            0.00     0.00      N/C     0.00
                                                                     PAGES READ VIA DYN.PREFETCH       0.00     0.00      N/C     0.00
BUFFERPOOL EXPANSIONS              N/A      N/A      N/A      N/A    D.PRF.PAGES READ/D.PRF.READ        N/C

                                                                     PREF.DISABLED-NO BUFFER        8554.00   854.86      N/C     0.47
                                                                     PREF.DISABLED-NO READ ENG         0.00     0.00      N/C     0.00

                                                                     SYNC.HPOOL READ                   0.00     0.00      N/C     0.00
                                                                     ASYNC.HPOOL READ                  0.00     0.00      N/C     0.00
                                                                     HPOOL READ FAILED                 0.00     0.00      N/C     0.00
                                                                     ASYN.DA.MOVER HPOOL READ-S        0.00     0.00      N/C     0.00
                                                                     ASYN.DA.MOVER HPOOL READ-F        0.00     0.00      N/C     0.00

                                                                     PAGE-INS REQUIRED FOR READ        0.00     0.00      N/C     0.00
```

Figure 78. Run 2WP: 2-Way First CPC, BP5 General

Appendix C. DB2 PM 2-Way Accounting Report (2WP)

This appendix compares DB2 PM Accounting reports for the 2-way measurements.

C.1 DB2 PM Accounting Report: 2-Way First CPC, Highlights Total

```
  LOCATION: DSNDB0G                             DB2 PERFORMANCE MONITOR (V4)                                PAGE: 1-66
     GROUP: DSNDB0G                               ACCOUNTING REPORT - LONG                    REQUESTED FROM: 09/09/95 08:41:15.00
    MEMBER: DB1G                                                                                          TO: 09/09/95 08:51:17.00
 SUBSYSTEM: DB1G                                     ORDER: PLANNAME                           INTERVAL FROM: 09/09/95 08:41:15.04
DB2 VERSION: V4                                                                                           TO: 09/09/95 08:51:16.45

***  GRAND TOTAL  ***
AVERAGE       APPL (CLASS 1)  DB2 (CLASS 2)  IFI (CLASS 5)    CLASS 3 SUSP.   AVERAGE TIME  AV.EVENT    HIGHLIGHTS
------------  --------------  --------------  --------------  --------------  ------------  --------    --------------------------
ELAPSED TIME        2.101547        0.451835             N/P  LOCK/LATCH          0.006334      1.90    #OCCURRENCES    :    18096
CPU TIME            0.083094        0.064449             N/P  SYNCHRON. I/O       0.223908      9.04    #ALLIEDS        :    18096
 TCB                0.083094        0.064449             N/P  OTHER READ I/O      0.166700      1.17    #ALLIEDS DISTRIB:        0
 TCB-STPROC         0.000000        0.000000             N/A  OTHER WRTE I/O      0.000006      0.00    #DBATS          :        0
 CPU-PARALL         0.000000        0.000000             N/A  SER.TASK SWTCH      0.031217      0.54    #DBATS DISTRIB. :        0
NOT ACCOUNT.             N/A             N/C             N/A  ARC.LOG(QUIES)      0.000000      0.00    #NO PROGRAM DATA:    18096
DB2 ENT/EXIT             N/A           61.44             N/A  ARC.LOG READ        0.000000      0.00    #NORMAL TERMINAT:    18096
EN/EX-STPROC             N/A            0.00             N/A  DRAIN LOCK          0.000000      0.00    #ABNORMAL TERMIN:        0
                                                              CLAIM RELEASE       0.000000      0.00    #CPU PARALLELISM:        0
DCAPT.DESCR.             N/A             N/A             N/P  PAGE LATCH          0.000398      0.00    #IO PARALLELISM :        0
LOG EXTRACT.             N/A             N/A             N/P  STORED PROC.        0.000000      0.00    #INCREMENT. BIND:        0
                                                              NOTIFY MSGS.        0.000000      0.00    #COMMITS        :    18096
NOT NULL               18096           18096               0  GLOBAL CONT.        0.023198      0.11    #ROLLBACKS      :        0
                                                              TOTAL CLASS 3       0.451760     12.76    UPDATE/COMMIT   :     6.65

SQL DML   AVERAGE     TOTAL   SQL DCL            TOTAL   SQL DDL     CREATE   DROP  ALTER   LOCKING          AVERAGE     TOTAL
--------  --------  --------  --------------  --------   ----------  ------  ------ ------  --------------  --------  --------
SELECT        4.81     87112  LOCK TABLE             0   TABLE            0       0      0  TIMEOUTS            0.00         0
INSERT        2.94     53138  GRANT                  0   INDEX            0       0      0  DEADLOCKS           0.00         0
UPDATE        3.57     64672  REVOKE                 0   TABLESPACE       0       0      0  ESCAL.(SHARED)      0.00         0
DELETE        0.14      2510  SET CURR.SQLID         0   DATABASE         0       0      0  ESCAL.(EXCLUS)      0.00         0
                              SET HOST VAR.          0   STOGROUP         0       0      0  MAX LOCKS HELD      4.97        46
DESCRIBE      0.00         0  SET CUR.DEGREE         0   SYNONYM          0       0    N/A  LOCK REQUEST        6.74    121919
DESC.TBL      0.00         0  SET RULES              0   VIEW             0       0    N/A  UNLOCK REQUEST      1.57     28408
PREPARE       0.00         0  CONNECT TYPE 1         0   ALIAS            0       0    N/A  QUERY REQUEST       0.00         0
OPEN          4.51     81592  CONNECT TYPE 2         0   PACKAGE        N/A       0    N/A  CHANGE REQUEST      3.61     65344
FETCH         9.24    167247  SET CONNECTION         0                                      OTHER REQUEST       0.00         0
CLOSE         2.54     45998  RELEASE                0   TOTAL            0       0      0  LOCK SUSPENS.       0.02       345
                                                                                            LATCH SUSPENS.      0.17      3051
DML-ALL      27.76    502269  DCL-ALL                0   COMMENT ON       0                 OTHER SUSPENS.      0.00         0
                                                         LABEL ON         0                 TOTAL SUSPENS.      0.19      3396

NORMAL TERM.       AVERAGE     TOTAL    ABNORMAL TERM.        TOTAL    IN DOUBT           TOTAL   DRAIN/CLAIM       AVERAGE     TOTAL
---------------  --------  --------    -----------------  --------    --------------  --------   --------------  --------  --------
NEW USER             1.00     18096    APPL.PROGR. ABEND         0    APPL.PGM ABEND         0   DRAIN REQUESTS      0.00         0
DEALLOCATION         0.00         0    END OF MEMORY             0    END OF MEMORY          0   DRAIN FAILED        0.00         0
APPL.PROGR. END      0.00         0    RESOL.IN DOUBT            0    END OF TASK            0   CLAIM REQUESTS      9.10    164644
RESIGNON             0.00         0    CANCEL FORCE              0    CANCEL FORCE           0   CLAIM FAILED        0.00         0
DBAT INACTIVE        0.00         0
```

Figure 79. Run 2WP: 2-Way First CPC, Highlights Total

C.2 DB2 PM Accounting Report: 2-Way Second CPC, Highlights Total

```
   LOCATION: DSNDB0G                                 DB2 PERFORMANCE MONITOR (V4)                                     PAGE: 1-70
      GROUP: DSNDB0G                                   ACCOUNTING REPORT - LONG                          REQUESTED FROM: 09/09/95 08:41:15.00
     MEMBER: DB2G                                                                                                   TO: 09/09/95 08:51:17.00
  SUBSYSTEM: DB2G                                           ORDER: PLANNAME                               INTERVAL FROM: 09/09/95 08:41:15.00
DB2 VERSION: V4                                                                                                     TO: 09/09/95 08:51:16.57

***  GRAND TOTAL  ***
AVERAGE       APPL (CLASS 1)  DB2 (CLASS 2)  IFI (CLASS 5)    CLASS 3 SUSP.   AVERAGE TIME  AV.EVENT    HIGHLIGHTS
------------  --------------  --------------  --------------    --------------  ------------  --------    --------------------------
ELAPSED TIME        2.084017        0.459894             N/P    LOCK/LATCH          0.007494      2.04    #OCCURRENCES    :    18233
CPU TIME            0.084746        0.065662             N/P    SYNCHRON. I/O       0.232183      9.35    #ALLIEDS        :    18233
 TCB                0.084746        0.065662             N/P    OTHER READ I/O      0.162587      1.14    #ALLIEDS DISTRIB:        0
 TCB-STPROC         0.000000        0.000000             N/A    OTHER WRTE I/O      0.000013      0.00    #DBATS          :        0
 CPU-PARALL         0.000000        0.000000             N/A    SER.TASK SWTCH      0.031082      0.54    #DBATS DISTRIB. :        0
NOT ACCOUNT.             N/A             N/C             N/A    ARC.LOG(QUIES)      0.000000      0.00    #NO PROGRAM DATA:    18233
DB2 ENT/EXIT             N/A           62.42             N/A    ARC.LOG READ        0.000000      0.00    #NORMAL TERMINAT:    18233
EN/EX-STPROC             N/A            0.00             N/A    DRAIN LOCK          0.000000      0.00    #ABNORMAL TERMIN:        0
                                                                CLAIM RELEASE       0.000000      0.00    #CPU PARALLELISM:        0
DCAPT.DESCR.             N/A             N/A             N/P    PAGE LATCH          0.000228      0.00    #IO PARALLELISM :        0
LOG EXTRACT.             N/A             N/A             N/P    STORED PROC.        0.000000      0.00    #INCREMENT. BIND:        0
                                                                NOTIFY MSGS.        0.000000      0.00    #COMMITS        :    18233
NOT NULL               18233           18233               0    GLOBAL CONT.        0.022649      0.11    #ROLLBACKS      :        0
                                                                TOTAL CLASS 3       0.456235     13.18    UPDATE/COMMIT   :     6.85

SQL DML     AVERAGE     TOTAL    SQL DCL           TOTAL    SQL DDL     CREATE    DROP   ALTER    LOCKING            AVERAGE     TOTAL
--------   --------  --------    --------------  --------    ----------  ------  ------  ------    --------------   --------  --------
SELECT         4.85     88421    LOCK TABLE             0    TABLE            0       0       0    TIMEOUTS             0.00         0
INSERT         2.90     52803    GRANT                  0    INDEX            0       0       0    DEADLOCKS            0.00         0
UPDATE         3.74     68226    REVOKE                 0    TABLESPACE       0       0       0    ESCAL.(SHARED)       0.00         0
DELETE         0.21      3810    SET CURR.SQLID         0    DATABASE         0       0       0    ESCAL.(EXCLUS)       0.00         0
                                 SET HOST VAR.          0    STOGROUP         0       0       0    MAX LOCKS HELD       5.23        47
DESCRIBE       0.00         0    SET CUR.DEGREE         0    SYNONYM          0       0     N/A    LOCK REQUEST         7.30    133016
DESC.TBL       0.00         0    SET RULES              0    VIEW             0       0     N/A    UNLOCK REQUEST       1.71     31255
PREPARE        0.00         0    CONNECT TYPE 1         0    ALIAS            0       0     N/A    QUERY REQUEST        0.00         0
OPEN           4.60     83786    CONNECT TYPE 2         0    PACKAGE        N/A       0     N/A    CHANGE REQUEST       3.80     69323
FETCH          9.30    169636    SET CONNECTION         0                                          OTHER REQUEST        0.00         0
CLOSE          2.65     48261    RELEASE                0    TOTAL            0       0       0    LOCK SUSPENS.        0.02       376
                                                                                                   LATCH SUSPENS.       0.19      3377
DML-ALL       28.24    514943    DCL-ALL                0    COMMENT ON       0                    OTHER SUSPENS.       0.00         0
                                                             LABEL ON         0                    TOTAL SUSPENS.       0.21      3753

NORMAL TERM.       AVERAGE     TOTAL    ABNORMAL TERM.        TOTAL    IN DOUBT           TOTAL    DRAIN/CLAIM        AVERAGE     TOTAL
---------------   --------  --------    -----------------  --------    --------------  --------    --------------   --------  --------
NEW USER              1.00     18233    APPL.PROGR. ABEND         0    APPL.PGM ABEND         0    DRAIN REQUESTS       0.00         0
DEALLOCATION          0.00         0    END OF MEMORY             0    END OF MEMORY          0    DRAIN FAILED         0.00         0
APPL.PROGR. END       0.00         0    RESOL.IN DOUBT            0    END OF TASK            0    CLAIM REQUESTS       9.13    166530
RESIGNON              0.00         0    CANCEL FORCE              0    CANCEL FORCE           0    CLAIM FAILED         0.00         0
DBAT INACTIVE         0.00         0
```

Figure 80. Run 2WP: 2-Way Second CPC, Highlights Total

C.3 DB2 PM Accounting Report: 2-Way First CPC, Data Sharing Total

```
   LOCATION: DSNDB0G                            DB2 PERFORMANCE MONITOR (V4)                               PAGE: 1-67
      GROUP: DSNDB0G                              ACCOUNTING REPORT - LONG                    REQUESTED FROM: 09/09/95 08:41:15.00
     MEMBER: DB1G                                                                                        TO: 09/09/95 08:51:17.00
  SUBSYSTEM: DB1G                                      ORDER: PLANNAME                         INTERVAL FROM: 09/09/95 08:41:15.04
DB2 VERSION: V4                                                                                          TO: 09/09/95 08:51:16.45

***  GRAND TOTAL  ***
DATA CAPTURE        AVERAGE     TOTAL   DATA SHARING        AVERAGE     TOTAL   QUERY PARALLELISM               TOTAL   AVERAGE
-----------------  --------  --------   -------------------  --------  --------   ----------------------------  --------  --------
IFI CALLS MADE          N/P       N/P   LOCK REQ   - PLOCKS     1.43     25943   MAXIMUM DEGREE                      0       N/A
RECORDS CAPTURED        N/P       N/P   UNLOCK REQ - PLOCKS     1.41     25490   GROUPS EXECUTED                     0      0.00
LOG RECORDS READ        N/P       N/P   CHANGE REQ - PLOCKS     0.00        10   PLANNED DEGREE                      0      0.00
ROWS RETURNED           N/P       N/P   LOCK REQ   - XES        6.95    125767   REDUCED - NO BUFFER                 0      0.00
RECORDS RETURNED        N/P       N/P   UNLOCK REQ - XES        2.31     41772   SEQUENTIAL- CURSOR                  0      0.00
DATA DESC. RETURN       N/P       N/P   CHANGE REQ - XES        3.60     65195   SEQUENTIAL- NO ESA SORT             0      0.00
TABLES RETURNED         N/P       N/P   SUSPENDS   - IRLM       0.03       622   SEQUENTIAL- NO BUFFER               0      0.00
DESCRIBES               N/P       N/P   SUSPENDS   - XES        0.00         0   SEQUENTIAL- ENCLAVE SERVICES        0      0.00
                                        SUSPENDS   - FALSE      0.06      1018
                                        INCOMPATIBLE LOCKS      0.00         0
                                        NOTIFY MSGS SENT        0.00         0

STORED PROCEDURES   AVERAGE     TOTAL   RID LIST            AVERAGE     TOTAL
-----------------  --------  --------   -------------------  --------  --------
CALL STATEMENTS        0.00         0   USED                    8.12    146864
PROCEDURE ABENDS       0.00         0   FAIL-NO STORAGE         0.00         0
CALL TIMEOUT           0.00         0   FAIL-LIMIT EXCEEDED     0.00         0
CALL REJECT            0.00         0

BP0                   AVERAGE     TOTAL   BP1                   AVERAGE     TOTAL   BP2                   AVERAGE     TOTAL
-------------------  --------  --------   -------------------  --------  --------   -------------------  --------  --------
EXPANSIONS                N/A       N/A   EXPANSIONS                N/A       N/A   EXPANSIONS                N/A       N/A
GETPAGES                 4.48     40966   GETPAGES                27.60    371078   GETPAGES                28.99    390453
BUFFER UPDATES           3.24     29585   BUFFER UPDATES           0.91     12204   BUFFER UPDATES           0.00         0
SYNCHRONOUS WRITE        0.00         0   SYNCHRONOUS WRITE        0.00         0   SYNCHRONOUS WRITE        0.00         0
SYNCHRONOUS READ         0.00         0   SYNCHRONOUS READ         0.00         0   SYNCHRONOUS READ         3.75     50556
SEQUENTIAL PREFETCH      0.17      1562   SEQUENTIAL PREFETCH      0.00         0   SEQUENTIAL PREFETCH      0.00         0
LIST PREFETCH            0.00         0   LIST PREFETCH            0.00         0   LIST PREFETCH            0.00         0
DYNAMIC PREFETCH         0.00        22   DYNAMIC PREFETCH         0.16      2141   DYNAMIC PREFETCH         0.88     11842
PAGES READ ASYNCHR.      0.00        38   PAGES READ ASYNCHR.      0.00         0   PAGES READ ASYNCHR.     19.90    267988
HPOOL WRITES             0.00         0   HPOOL WRITES             0.00         0   HPOOL WRITES             0.00         0
HPOOL WRITES-FAILED      0.00         0   HPOOL WRITES-FAILED      0.00         0   HPOOL WRITES-FAILED      0.00         0
PAGES READ-HPOOL         0.00         0   PAGES READ-HPOOL         0.00         0   PAGES READ-HPOOL         0.00         0
HPOOL READS              0.00         0   HPOOL READS              0.00         0   HPOOL READS              0.00         0
HPOOL READS FAILED       0.00         0   HPOOL READS FAILED       0.00         0   HPOOL READS FAILED       0.00         0
```

Figure 81. Run 2WP: 2-Way First CPC, Data Sharing Total

C.4 DB2 PM Accounting Report: 2-Way Second CPC, Data Sharing Total

```
   LOCATION: DSNDB0G                        DB2 PERFORMANCE MONITOR (V4)                          PAGE: 1-71
      GROUP: DSNDB0G                          ACCOUNTING REPORT - LONG                  REQUESTED FROM: 09/09/95 08:41:15.00
     MEMBER: DB2G                                                                                   TO: 09/09/95 08:51:17.00
  SUBSYSTEM: DB2G                                 ORDER: PLANNAME                        INTERVAL FROM: 09/09/95 08:41:15.00
DB2 VERSION: V4                                                                                     TO: 09/09/95 08:51:16.57

***  GRAND TOTAL  ***
DATA CAPTURE       AVERAGE    TOTAL    DATA SHARING         AVERAGE    TOTAL    QUERY PARALLELISM              TOTAL   AVERAGE
-----------------  --------  --------  -------------------  --------  --------  ---------------------------  --------  --------
IFI CALLS MADE          N/P       N/P  LOCK REQ   - PLOCKS      1.52     27637  MAXIMUM DEGREE                      0       N/A
RECORDS CAPTURED        N/P       N/P  UNLOCK REQ - PLOCKS      1.49     27202  GROUPS EXECUTED                     0      0.00
LOG RECORDS READ        N/P       N/P  CHANGE REQ - PLOCKS      0.00         4  PLANNED DEGREE                      0      0.00
ROWS RETURNED           N/P       N/P  LOCK REQ   - XES         7.52    137060  REDUCED - NO BUFFER                 0      0.00
RECORDS RETURNED        N/P       N/P  UNLOCK REQ - XES         2.47     45049  SEQUENTIAL- CURSOR                  0      0.00
DATA DESC. RETURN       N/P       N/P  CHANGE REQ - XES         3.79     69175  SEQUENTIAL- NO ESA SORT             0      0.00
TABLES RETURNED         N/P       N/P  SUSPENDS   - IRLM        0.04       640  SEQUENTIAL- NO BUFFER               0      0.00
DESCRIBES               N/P       N/P  SUSPENDS   - XES         0.00         0  SEQUENTIAL- ENCLAVE SERVICES        0      0.00
                                       SUSPENDS   - FALSE       0.06      1089
                                       INCOMPATIBLE LOCKS       0.00         0
                                       NOTIFY MSGS SENT         0.00         1

STORED PROCEDURES  AVERAGE    TOTAL    RID LIST             AVERAGE    TOTAL
-----------------  --------  --------  -------------------  --------  --------
CALL STATEMENTS        0.00         0  USED                     8.00    145905
PROCEDURE ABENDS       0.00         0  FAIL-NO STORAGE          0.00         0
CALL TIMEOUT           0.00         0  FAIL-LIMIT EXCEEDED      0.00         0
CALL REJECT            0.00         0

BP0                  AVERAGE    TOTAL    BP1                  AVERAGE    TOTAL    BP2                  AVERAGE    TOTAL
-------------------  --------  --------  -------------------  --------  --------  -------------------  --------  --------
EXPANSIONS                N/A       N/A  EXPANSIONS                N/A       N/A  EXPANSIONS                N/A       N/A
GETPAGES                 4.47     40854  GETPAGES                27.54    370395  GETPAGES                28.77    391302
BUFFER UPDATES           3.23     29505  BUFFER UPDATES           0.91     12196  BUFFER UPDATES           0.00         0
SYNCHRONOUS WRITE        0.00         0  SYNCHRONOUS WRITE        0.00         0  SYNCHRONOUS WRITE        0.00         0
SYNCHRONOUS READ         0.00         0  SYNCHRONOUS READ         0.00         0  SYNCHRONOUS READ         3.69     50190
SEQUENTIAL PREFETCH      0.17      1535  SEQUENTIAL PREFETCH      0.00         0  SEQUENTIAL PREFETCH      0.00         0
LIST PREFETCH            0.00         0  LIST PREFETCH            0.00         0  LIST PREFETCH            0.00         0
DYNAMIC PREFETCH         0.00        19  DYNAMIC PREFETCH         0.17      2222  DYNAMIC PREFETCH         0.91     12335
PAGES READ ASYNCHR.      0.00        31  PAGES READ ASYNCHR.      0.00         0  PAGES READ ASYNCHR.     20.34    276648
HPOOL WRITES             0.00         0  HPOOL WRITES             0.00         0  HPOOL WRITES             0.00         0
HPOOL WRITES-FAILED      0.00         0  HPOOL WRITES-FAILED      0.00         0  HPOOL WRITES-FAILED      0.00         0
PAGES READ-HPOOL         0.00         0  PAGES READ-HPOOL         0.00         0  PAGES READ-HPOOL         0.00         0
HPOOL READS              0.00         0  HPOOL READS              0.00         0  HPOOL READS              0.00         0
HPOOL READS FAILED       0.00         0  HPOOL READS FAILED       0.00         0  HPOOL READS FAILED       0.00         0
```

Figure 82. Run 2WP: 2-Way Second CPC, Data Sharing Total

Appendix D. DB2 PM 1-Way and 2-Way Accounting Reports (1WP and 2WP)

This appendix compares DB2 PM Accounting reports for the 1-way measurements and the first CPC for the 2-way measurements.

D.1 DB2 PM Accounting Report: 1-Way CPC, Highlights Total

```
   LOCATION: DSNDB0G                                 DB2 PERFORMANCE MONITOR (V4)                                PAGE: 1-58
      GROUP: DSNDB0G                                   ACCOUNTING REPORT - LONG                       REQUESTED FROM: 09/09/95 12:47:18.00
     MEMBER: DB1G                                                                                                TO: 09/09/95 12:57:20.00
  SUBSYSTEM: DB1G                                          ORDER: PLANNAME                             INTERVAL FROM: 09/09/95 12:47:18.07
DB2 VERSION: V4                                                                                                  TO: 09/09/95 12:57:19.99

***  GRAND TOTAL  ***
AVERAGE       APPL (CLASS 1)  DB2 (CLASS 2)  IFI (CLASS 5)    CLASS 3 SUSP.   AVERAGE TIME  AV.EVENT    HIGHLIGHTS
------------  --------------  --------------  --------------   --------------  ------------  --------    --------------------------
ELAPSED TIME        1.763538        0.482750             N/P    LOCK/LATCH          0.003570      2.39    #OCCURRENCES    :     21553
CPU TIME            0.075018        0.055852             N/P    SYNCHRON. I/O       0.285165     11.20    #ALLIEDS        :     21553
 TCB                0.075018        0.055852             N/P    OTHER READ I/O      0.180544      1.14    #ALLIEDS DISTRIB:         0
 TCB-STPROC         0.000000        0.000000             N/A    OTHER WRTE I/O      0.000948      0.01    #DBATS          :         0
 CPU-PARALL         0.000000        0.000000             N/A    SER.TASK SWTCH      0.022624      0.54    #DBATS DISTRIB. :         0
NOT ACCOUNT.             N/A             N/C             N/A    ARC.LOG(QUIES)      0.000000      0.00    #NO PROGRAM DATA:     21553
DB2 ENT/EXIT             N/A           62.38             N/A    ARC.LOG READ        0.000000      0.00    #NORMAL TERMINAT:     21553
EN/EX-STPROC             N/A            0.00             N/A    DRAIN LOCK          0.000000      0.00    #ABNORMAL TERMIN:         0
                                                                CLAIM RELEASE       0.000000      0.00    #CPU PARALLELISM:         0
DCAPT.DESCR.             N/A             N/A             N/P    PAGE LATCH          0.000213      0.00    #IO PARALLELISM :         0
LOG EXTRACT.             N/A             N/A             N/P    STORED PROC.        0.000000      0.00    #INCREMENT. BIND:         0
                                                                NOTIFY MSGS.        0.000000      0.00    #COMMITS        :     21553
NOT NULL               21553           21553               0    GLOBAL CONT.        0.000000      0.00    #ROLLBACKS      :         0
                                                                TOTAL CLASS 3       0.493063     15.28    UPDATE/COMMIT   :      6.84

SQL DML   AVERAGE     TOTAL   SQL DCL            TOTAL    SQL DDL     CREATE   DROP  ALTER    LOCKING          AVERAGE     TOTAL
--------  --------  --------  --------------  --------   ----------  ------  ------  ------   --------------  --------  --------
SELECT      4.87    105049   LOCK TABLE            0    TABLE            0       0      0    TIMEOUTS            0.00         0
INSERT      2.89     62190   GRANT                 0    INDEX            0       0      0    DEADLOCKS           0.00         0
UPDATE      3.74     80611   REVOKE                0    TABLESPACE       0       0      0    ESCAL.(SHARED)      0.00         0
DELETE      0.21      4520   SET CURR.SQLID        0    DATABASE         0       0      0    ESCAL.(EXCLUS)      0.00         0
                             SET HOST VAR.         0    STOGROUP         0       0      0    MAX LOCKS HELD      5.20        48
DESCRIBE    0.00         0   SET CUR.DEGREE        0    SYNONYM          0       0    N/A    LOCK REQUEST        6.63    142820
DESC.TBL    0.00         0   SET RULES             0    VIEW             0       0    N/A    UNLOCK REQUEST      1.52     32835
PREPARE     0.00         0   CONNECT TYPE 1        0    ALIAS            0       0    N/A    QUERY REQUEST       0.00         0
OPEN        4.59     98979   CONNECT TYPE 2        0    PACKAGE        N/A       0    N/A    CHANGE REQUEST      3.80     81941
FETCH       9.28    199942   SET CONNECTION        0                                         OTHER REQUEST       0.00         0
CLOSE       2.64     56936   RELEASE               0    TOTAL            0       0      0    LOCK SUSPENS.       0.00       100
                                                                                             LATCH SUSPENS.      0.11      2300
DML-ALL    28.22    608227   DCL-ALL               0    COMMENT ON       0                   OTHER SUSPENS.      0.00         0
                                                        LABEL ON         0                   TOTAL SUSPENS.      0.11      2400

NORMAL TERM.      AVERAGE     TOTAL   ABNORMAL TERM.        TOTAL    IN DOUBT           TOTAL   DRAIN/CLAIM       AVERAGE     TOTAL
---------------  --------  --------   -----------------  --------   --------------  --------   --------------  --------  --------
NEW USER             1.00     21553   APPL.PROGR. ABEND         0   APPL.PGM ABEND         0   DRAIN REQUESTS      0.00         0
DEALLOCATION         0.00         0   END OF MEMORY             0   END OF MEMORY          0   DRAIN FAILED        0.00         0
APPL.PROGR. END      0.00         0   RESOL.IN DOUBT            0   END OF TASK            0   CLAIM REQUESTS      9.10    196070
RESIGNON             0.00         0   CANCEL FORCE              0   CANCEL FORCE           0   CLAIM FAILED        0.00         0
DBAT INACTIVE        0.00         0
```

Figure 83. Run 1WP: 1-Way CPC, Highlights Total

D.2 DB2 PM Accounting Report: 2-Way First CPC, Highlights Total

```
   LOCATION: DSNDB0G                                   DB2 PERFORMANCE MONITOR (V4)                                   PAGE: 1-66
      GROUP: DSNDB0G                                     ACCOUNTING REPORT - LONG                        REQUESTED FROM: 09/09/95 08:41:15.00
     MEMBER: DB1G                                                                                                   TO: 09/09/95 08:51:17.00
  SUBSYSTEM: DB1G                                           ORDER: PLANNAME                               INTERVAL FROM: 09/09/95 08:41:15.04
DB2 VERSION: V4                                                                                                     TO: 09/09/95 08:51:16.45

***  GRAND TOTAL  ***
AVERAGE       APPL (CLASS 1)  DB2 (CLASS 2)  IFI (CLASS 5)    CLASS 3 SUSP.   AVERAGE TIME  AV.EVENT     HIGHLIGHTS
------------  --------------  --------------  --------------    --------------  ------------  --------     --------------------------
ELAPSED TIME        2.101547        0.451835             N/P    LOCK/LATCH          0.006334      1.90     #OCCURRENCES    :    18096
CPU TIME            0.083094        0.064449             N/P    SYNCHRON. I/O       0.223908      9.04     #ALLIEDS        :    18096
 TCB                0.083094        0.064449             N/P    OTHER READ I/O      0.166700      1.17     #ALLIEDS DISTRIB:        0
 TCB-STPROC         0.000000        0.000000             N/A    OTHER WRTE I/O      0.000006      0.00     #DBATS          :        0
 CPU-PARALL         0.000000        0.000000             N/A    SER.TASK SWTCH      0.031217      0.54     #DBATS DISTRIB. :        0
NOT ACCOUNT.             N/A             N/C             N/A    ARC.LOG(QUIES)      0.000000      0.00     #NO PROGRAM DATA:    18096
DB2 ENT/EXIT             N/A           61.44             N/A    ARC.LOG READ        0.000000      0.00     #NORMAL TERMINAT:    18096
EN/EX-STPROC             N/A            0.00             N/A    DRAIN LOCK          0.000000      0.00     #ABNORMAL TERMIN:        0
                                                                CLAIM RELEASE       0.000000      0.00     #CPU PARALLELISM:        0
DCAPT.DESCR.             N/A             N/A             N/P    PAGE LATCH          0.000398      0.00     #IO PARALLELISM :        0
LOG EXTRACT.             N/A             N/A             N/P    STORED PROC.        0.000000      0.00     #INCREMENT. BIND:        0
                                                                NOTIFY MSGS.        0.000000      0.00     #COMMITS        :    18096
NOT NULL               18096           18096               0    GLOBAL CONT.        0.023198      0.11     #ROLLBACKS      :        0
                                                                TOTAL CLASS 3       0.451760     12.76     UPDATE/COMMIT   :     6.65

SQL DML    AVERAGE     TOTAL   SQL DCL           TOTAL    SQL DDL     CREATE   DROP  ALTER     LOCKING          AVERAGE     TOTAL
--------  --------  --------   --------------  --------   ----------  ------  ------  ------    --------------  --------  --------
SELECT        4.81     87112   LOCK TABLE             0   TABLE            0       0       0    TIMEOUTS            0.00         0
INSERT        2.94     53138   GRANT                  0   INDEX            0       0       0    DEADLOCKS           0.00         0
UPDATE        3.57     64672   REVOKE                 0   TABLESPACE       0       0       0    ESCAL.(SHARED)      0.00         0
DELETE        0.14      2510   SET CURR.SQLID         0   DATABASE         0       0       0    ESCAL.(EXCLUS)      0.00         0
                               SET HOST VAR.          0   STOGROUP         0       0       0    MAX LOCKS HELD      4.97        46
DESCRIBE      0.00         0   SET CUR.DEGREE         0   SYNONYM          0       0     N/A    LOCK REQUEST        6.74    121919
DESC.TBL      0.00         0   SET RULES              0   VIEW             0       0     N/A    UNLOCK REQUEST      1.57     28408
PREPARE       0.00         0   CONNECT TYPE 1         0   ALIAS            0       0     N/A    QUERY REQUEST       0.00         0
OPEN          4.51     81592   CONNECT TYPE 2         0   PACKAGE        N/A       0     N/A    CHANGE REQUEST      3.61     65344
FETCH         9.24    167247   SET CONNECTION         0                                         OTHER REQUEST       0.00         0
CLOSE         2.54     45998   RELEASE                0   TOTAL            0       0       0    LOCK SUSPENS.       0.02       345
                                                                                                LATCH SUSPENS.      0.17      3051
DML-ALL      27.76    502269   DCL-ALL                0   COMMENT ON       0                    OTHER SUSPENS.      0.00         0
                                                          LABEL ON         0                    TOTAL SUSPENS.      0.19      3396

NORMAL TERM.      AVERAGE     TOTAL   ABNORMAL TERM.        TOTAL    IN DOUBT           TOTAL    DRAIN/CLAIM       AVERAGE     TOTAL
---------------  --------  --------   -----------------  --------    --------------  --------    --------------  --------  --------
NEW USER             1.00     18096   APPL.PROGR. ABEND         0    APPL.PGM ABEND         0    DRAIN REQUESTS      0.00         0
DEALLOCATION         0.00         0   END OF MEMORY             0    END OF MEMORY          0    DRAIN FAILED        0.00         0
APPL.PROGR. END      0.00         0   RESOL.IN DOUBT            0    END OF TASK            0    CLAIM REQUESTS      9.10    164644
RESIGNON             0.00         0   CANCEL FORCE              0    CANCEL FORCE           0    CLAIM FAILED        0.00         0
DBAT INACTIVE        0.00         0
```

Figure 84. Run 2WP: 2-Way First CPC, Highlights Total

D.3 DB2 PM Accounting Report: 1-Way CPC, Data Sharing Total

LOCATION: DSNDB0G
GROUP: DSNDB0G
MEMBER: DB1G
SUBSYSTEM: DB1G
DB2 VERSION: V4

DB2 PERFORMANCE MONITOR (V4)
ACCOUNTING REPORT - LONG

ORDER: PLANNAME

PAGE: 1-59
REQUESTED FROM: 09/09/95 12:47:18.00
TO: 09/09/95 12:57:20.00
INTERVAL FROM: 09/09/95 12:47:18.07
TO: 09/09/95 12:57:19.99

*** GRAND TOTAL ***

DATA CAPTURE	AVERAGE	TOTAL
IFI CALLS MADE	N/P	N/P
RECORDS CAPTURED	N/P	N/P
LOG RECORDS READ	N/P	N/P
ROWS RETURNED	N/P	N/P
RECORDS RETURNED	N/P	N/P
DATA DESC. RETURN	N/P	N/P
TABLES RETURNED	N/P	N/P
DESCRIBES	N/P	N/P

DATA SHARING	AVERAGE	TOTAL
LOCK REQ - PLOCKS	0.00	0
UNLOCK REQ - PLOCKS	0.00	0
CHANGE REQ - PLOCKS	0.00	0
LOCK REQ - XES	0.04	801
UNLOCK REQ - XES	0.04	780
CHANGE REQ - XES	0.00	0
SUSPENDS - IRLM	0.00	0
SUSPENDS - XES	0.00	0
SUSPENDS - FALSE	0.00	0
INCOMPATIBLE LOCKS	0.00	0
NOTIFY MSGS SENT	0.00	0

QUERY PARALLELISM	TOTAL	AVERAGE
MAXIMUM DEGREE	0	N/A
GROUPS EXECUTED	0	0.00
PLANNED DEGREE	0	0.00
REDUCED - NO BUFFER	0	0.00
SEQUENTIAL- CURSOR	0	0.00
SEQUENTIAL- NO ESA SORT	0	0.00
SEQUENTIAL- NO BUFFER	0	0.00
SEQUENTIAL- ENCLAVE SERVICES	0	0.00

STORED PROCEDURES	AVERAGE	TOTAL
CALL STATEMENTS	0.00	0
PROCEDURE ABENDS	0.00	0
CALL TIMEOUT	0.00	0
CALL REJECT	0.00	0

RID LIST	AVERAGE	TOTAL
USED	7.95	171444
FAIL-NO STORAGE	0.00	0
FAIL-LIMIT EXCEEDED	0.00	0

BP0	AVERAGE	TOTAL
EXPANSIONS	N/A	N/A
GETPAGES	4.44	48082
BUFFER UPDATES	3.21	34729
SYNCHRONOUS WRITE	0.00	0
SYNCHRONOUS READ	0.00	0
SEQUENTIAL PREFETCH	0.16	1757
LIST PREFETCH	0.00	0
DYNAMIC PREFETCH	0.00	13
PAGES READ ASYNCHR.	0.00	0
HPOOL WRITES	0.00	0
HPOOL WRITES-FAILED	0.00	0
PAGES READ-HPOOL	0.00	0
HPOOL READS	0.00	0
HPOOL READS FAILED	0.00	0

BP1	AVERAGE	TOTAL
EXPANSIONS	N/A	N/A
GETPAGES	27.62	439130
BUFFER UPDATES	0.91	14467
SYNCHRONOUS WRITE	0.00	0
SYNCHRONOUS READ	0.00	0
SEQUENTIAL PREFETCH	0.00	0
LIST PREFETCH	0.00	0
DYNAMIC PREFETCH	0.16	2470
PAGES READ ASYNCHR.	0.00	0
HPOOL WRITES	0.00	0
HPOOL WRITES-FAILED	0.00	0
PAGES READ-HPOOL	0.00	0
HPOOL READS	0.00	0
HPOOL READS FAILED	0.00	0

BP2	AVERAGE	TOTAL
EXPANSIONS	N/A	N/A
GETPAGES	30.24	485082
BUFFER UPDATES	0.00	0
SYNCHRONOUS WRITE	0.00	0
SYNCHRONOUS READ	3.69	59258
SEQUENTIAL PREFETCH	0.00	0
LIST PREFETCH	0.00	0
DYNAMIC PREFETCH	0.91	14557
PAGES READ ASYNCHR.	20.54	329399
HPOOL WRITES	0.00	0
HPOOL WRITES-FAILED	0.00	0
PAGES READ-HPOOL	0.00	0
HPOOL READS	0.00	0
HPOOL READS FAILED	0.00	0

Figure 85. Run 1WP: 1-Way CPC, Data Sharing Total

D.4 DB2 PM Accounting Report: 2-Way First CPC, Data Sharing Total

```
   LOCATION: DSNDB0G                                   DB2 PERFORMANCE MONITOR (V4)                                    PAGE: 1-67
      GROUP: DSNDB0G                                    ACCOUNTING REPORT - LONG                       REQUESTED FROM: 09/09/95 08:41:15.00
     MEMBER: DB1G                                                                                                  TO: 09/09/95 08:51:17.00
  SUBSYSTEM: DB1G                                           ORDER: PLANNAME                             INTERVAL FROM: 09/09/95 08:41:15.04
DB2 VERSION: V4                                                                                                    TO: 09/09/95 08:51:16.45

***  GRAND TOTAL  ***
DATA CAPTURE       AVERAGE     TOTAL   DATA SHARING         AVERAGE     TOTAL   QUERY PARALLELISM               TOTAL   AVERAGE
-----------------  --------  --------  -------------------  --------  --------  ----------------------------  --------  --------
IFI CALLS MADE          N/P       N/P  LOCK REQ   - PLOCKS      1.43     25943  MAXIMUM DEGREE                       0       N/A
RECORDS CAPTURED        N/P       N/P  UNLOCK REQ - PLOCKS      1.41     25490  GROUPS EXECUTED                      0      0.00
LOG RECORDS READ        N/P       N/P  CHANGE REQ - PLOCKS      0.00        10  PLANNED DEGREE                       0      0.00
ROWS RETURNED           N/P       N/P  LOCK REQ   - XES         6.95    125767  REDUCED - NO BUFFER                  0      0.00
RECORDS RETURNED        N/P       N/P  UNLOCK REQ - XES         2.31     41772  SEQUENTIAL- CURSOR                   0      0.00
DATA DESC. RETURN       N/P       N/P  CHANGE REQ - XES         3.60     65195  SEQUENTIAL- NO ESA SORT              0      0.00
TABLES RETURNED         N/P       N/P  SUSPENDS   - IRLM        0.03       622  SEQUENTIAL- NO BUFFER                0      0.00
DESCRIBES               N/P       N/P  SUSPENDS   - XES         0.00         0  SEQUENTIAL- ENCLAVE SERVICES         0      0.00
                                       SUSPENDS   - FALSE       0.06      1018
                                       INCOMPATIBLE LOCKS       0.00         0
                                       NOTIFY MSGS SENT         0.00         0

STORED PROCEDURES  AVERAGE     TOTAL   RID LIST             AVERAGE     TOTAL
-----------------  --------  --------  -------------------  --------  --------
CALL STATEMENTS        0.00         0  USED                     8.12    146864
PROCEDURE ABENDS       0.00         0  FAIL-NO STORAGE          0.00         0
CALL TIMEOUT           0.00         0  FAIL-LIMIT EXCEEDED      0.00         0
CALL REJECT            0.00         0

BP0                  AVERAGE     TOTAL   BP1                  AVERAGE     TOTAL   BP2                  AVERAGE     TOTAL
-------------------  --------  --------  -------------------  --------  --------  -------------------  --------  --------
EXPANSIONS                N/A       N/A  EXPANSIONS                N/A       N/A  EXPANSIONS                N/A       N/A
GETPAGES                 4.48     40966  GETPAGES                27.60    371078  GETPAGES                28.99    390453
BUFFER UPDATES           3.24     29585  BUFFER UPDATES           0.91     12204  BUFFER UPDATES           0.00         0
SYNCHRONOUS WRITE        0.00         0  SYNCHRONOUS WRITE        0.00         0  SYNCHRONOUS WRITE        0.00         0
SYNCHRONOUS READ         0.00         0  SYNCHRONOUS READ         0.00         0  SYNCHRONOUS READ         3.75     50556
SEQUENTIAL PREFETCH      0.17      1562  SEQUENTIAL PREFETCH      0.00         0  SEQUENTIAL PREFETCH      0.00         0
LIST PREFETCH            0.00         0  LIST PREFETCH            0.00         0  LIST PREFETCH            0.00         0
DYNAMIC PREFETCH         0.00        22  DYNAMIC PREFETCH         0.16      2141  DYNAMIC PREFETCH         0.88     11842
PAGES READ ASYNCHR.      0.00        38  PAGES READ ASYNCHR.      0.00         0  PAGES READ ASYNCHR.     19.90    267988
HPOOL WRITES             0.00         0  HPOOL WRITES             0.00         0  HPOOL WRITES             0.00         0
HPOOL WRITES-FAILED      0.00         0  HPOOL WRITES-FAILED      0.00         0  HPOOL WRITES-FAILED      0.00         0
PAGES READ-HPOOL         0.00         0  PAGES READ-HPOOL         0.00         0  PAGES READ-HPOOL         0.00         0
HPOOL READS              0.00         0  HPOOL READS              0.00         0  HPOOL READS              0.00         0
HPOOL READS FAILED       0.00         0  HPOOL READS FAILED       0.00         0  HPOOL READS FAILED       0.00         0
```

Figure 86. Run 2WP: 2-Way First CPC, Data Sharing Total

Appendix E. MVS XCF DISPLAY Lock Structure

```
IXC360I  10.22.33  DISPLAY XCF 168
STRNAME: DSNDB0G_LOCK1
 STATUS: ALLOCATED
  POLICY SIZE    : 64000 K
  POLICY INITSIZE: N/A
  REBUILD PERCENT: N/A
  PREFERENCE LIST: LF01
  EXCLUSION LIST IS EMPTY
1
 ACTIVE STRUCTURE
 ----------------
  ALLOCATION TIME: 09/09/95 09:23:14
  CFNAME         : LF01
  COUPLING FACILITY: 009674.IBM.00.000000040016
                   PARTITION: 1   CPCID: 00
  ACTUAL SIZE    : 64000 K
  STORAGE INCREMENT SIZE: 256 K
  VERSION        : ABA5BA83 A3667001
  DISPOSITION    : KEEP
  ACCESS TIME    : 0
  MAX CONNECTIONS: 7
  # CONNECTIONS  : 3

 CONNECTION NAME  ID VERSION  SYSNAME  JOBNAME  ASID STATE
 ---------------- -- -------- -------- -------- ---- ----------------
 DXRDB0G$$DJ1G001 02 0002018F STLABC1  DB1GIRLM 0192 ACTIVE
 DXRDB0G$$DJ2G002 01 00010234 STLABC2  DB2GIRLM 0199 ACTIVE
 DXRDB0G$$DJ3G003 03 000300C0 STLABC2  DB3GIRLM 00C4 QUIESCED
```

Figure 87. XCF Display Lock Structure Output. Taken From XCF DISPAY Lock Structure

Appendix F. RMF Reports

F.1 RMF – Lock Structure Activities

```
                                   C O U P L I N G   F A C I L I T Y   A C T I V I T Y

                                                                                                                    PAGE   1

    MVS/ESA                SYSPLEX XESDEV            DATE 09/09/1995            INTERVAL 020.17.000
    SP5.2.0                RPT VERSION 5.2.0         TIME 08.31.00              CYCLE 01.000 SECONDS
--------------------------------------------------------------------------------------------------------------------------------
COUPLING FACILITY NAME = LF01
TOTAL SAMPLES(AVG) =  1217  (MAX) =  1217  (MIN) =  1217
--------------------------------------------------------------------------------------------------------------------------------
                                          COUPLING  FACILITY  USAGE  SUMMARY
--------------------------------------------------------------------------------------------------------------------------------
STRUCTURE SUMMARY
--------------------------------------------------------------------------------------------------------------------------------
                                                      % OF              % OF      AVG     LIST      LIST       LOCK     CACHE TOT
       STRUCTURE                            ALLOC     CF          #     ALL       REQ/    ENTRIES   ELEMENTS   ENTRIES  DIR ENTRY/
TYPE   NAME             STATUS CHG          SIZE      STORAGE    REQ    REQ       SEC     TOT/CUR   TOT/CUR    TOT/CUR  DATA ELEM
LOCK   DSNDB0G_LOCK1    ACTIVE               63M        3.1%  634018    100%    520.97     228K         0         17M      N/A

                                          -------      -----  -------  ------  -------
           STRUCTURE TOTALS                  63M        3.1%  634018    100%    520.97

STORAGE SUMMARY
--------------------------------------------------------------------------------------------------------------------------------
                                              ALLOC     % OF CF     ------- DUMP SPACE -------
                                              SIZE      STORAGE     % IN USE    MAX % REQUESTED

TOTAL CF STORAGE USED BY STRUCTURES            63M       3.1%
TOTAL CF DUMP STORAGE                           1M       0.1%          0.0%               0.0%
TOTAL CF STORAGE AVAILABLE                      2G      96.8%
                                             -------
TOTAL CF STORAGE SIZE                           2G

                                              ALLOC     % ALLOCATED
                                              SIZE
TOTAL CONTROL STORAGE DEFINED                   2G       3.2%
TOTAL DATA STORAGE DEFINED                      0K       0.0%

PROCESSOR SUMMARY
--------------------------------------------------------------------------------------------------------------------------------

                                              % BUSY
AVG. CF UTILIZATION                            2.1%
```

Figure 88 (Part 1 of 2). RMF Coupling Facility Activity Report – Coupling Facility Model 9674/R61

```
                              C O U P L I N G   F A C I L I T Y   A C T I V I T Y
                                                                                                              PAGE   2
     MVS/ESA                 SYSPLEX XESDEV            DATE 09/09/1995           INTERVAL 020.17.000
     SP5.2.0                 RPT VERSION 5.2.0         TIME 08.31.00             CYCLE 01.000 SECONDS

------------------------------------------------------------------------------------------------------------------------
 COUPLING FACILITY NAME = LF01
------------------------------------------------------------------------------------------------------------------------
                                                  COUPLING  FACILITY  STRUCTURE  ACTIVITY
------------------------------------------------------------------------------------------------------------------------
  1
 STRUCTURE NAME = DSNDB0G_LOCK1      TYPE = LOCK
           # REQ     -------------- REQUESTS -------------   -------------- QUEUED REQUESTS --------------
 SYSTEM    TOTAL              #     % OF  -SERV TIME(MIC)-          (ARRIVAL  RATE) --- QUEUE TIME(MIC)--       REQUEST
 NAME      AVG/SEC           REQ    ALL     AVG   STD_DEV           MIN   AVG   MAX          AVG  STD_DEV     CONTENTIONS

 STLABC1     634K   SYNC    634K   100%   223.4     51.4    HPRIO     0   0.0     0  NO SCH   0.0      0.0  # REQ         625K
           521.0    ASYNC     0    0.0%     0.0      0.0    LPRIO     0   0.0     0                         # REQ DELAYED 4538
                    CHNGD     0    0.0%  INCLUDED IN ASYNC                                                    -CONT       4535
                                                                                                              -FALSE CONT 1825

------------------------------------------------------------------------------------------------------------------------
 TOTAL       634K   SYNC    634K   100%   223.4     51.4                                 NO SCH   0.0      0.0  # REQ         625K
           521.0    ASYNC     0    0.0%     0.0      0.0                                                    # REQ DELAYED 4538
                    CHNGD     0    0.0%                                                                       -CONT       4535
                                                                                                              -FALSE CONT 1825


                              C O U P L I N G   F A C I L I T Y   A C T I V I T Y
                                                                                                              PAGE   3
     MVS/ESA                 SYSPLEX XESDEV            DATE 09/09/1995           INTERVAL 020.17.000
     SP5.2.0                 RPT VERSION 5.2.0         TIME 08.31.00             CYCLE 01.000 SECONDS

------------------------------------------------------------------------------------------------------------------------
 COUPLING FACILITY NAME = LF01
------------------------------------------------------------------------------------------------------------------------
                                                         SUBCHANNEL  ACTIVITY
------------------------------------------------------------------------------------------------------------------------
          # REQ                                 ----------- REQUESTS -----------   -------------- QUEUED REQUESTS --------------
 SYSTEM   TOTAL                   --BUSY--            #   -SERVICE TIME(MIC)-          (ARRIVAL  RATE) -- QUEUE TIME(MIC) --
 NAME     AVG/SEC -- CONFIG --  -COUNTS-             REQ     AVG    STD_DEV            MIN    AVG  MAX          AVG  STD_DEV

 STLABC1  636249  SCH GEN    4  PTH    0    SYNC   634018   223.4      51.4  HPRIO     0    0.0    0  NO SCH   0.0     0.0
           522.8  SCH USE    2  SCH 4902    ASYNC       0     0.0       0.0  LPRIO     0    0.0    0
                  SCH MAX    2              CHANGED     0  INCLUDED IN ASYNC DUMP      0    0.0    0  DUMP     0.0     0.0
                  PTH        1              UNSUCC      0     0.0       0.0
```

Figure 88 (Part 2 of 2). RMF Coupling Facility Activity Report – Coupling Facility Model 9674/R61

F.2 RMF – XCF Activities

```
                                              X C F  A C T I V I T Y
                                                                                                              PAGE 1
          MVS/ESA                SYSTEM ID 961              DATE 09/09/95            INTERVAL 01.00.000
          SP5.2.0                RPT VERSION 5.2.0          TIME 08.31.00            CYCLE 1.000 SECONDS

                                                  XCF USAGE BY SYSTEM
---------------------------------------------------------------------------------------------------------------------

                                        REMOTE SYSTEMS                                                        LOCAL
-----------------------------------------------------------------------------------------------    ----------------

 2                           OUTBOUND FROM STLABC1                             INBOUND TO STLABC1            STLABC1
---------------------------------------------------------------------------   --------------------------   ----------------

                                           ----- BUFFER -----      ALL
TO       TRANSPORT  BUFFER       REQ    %    %    %    %      PATHS      REQ   FROM           REQ      REQ    TRANSPORT      REQ
SYSTEM   CLASS      LENGTH       OUT  SML  FIT  BIG  OVR    UNAVAIL   REJECT   SYSTEM          IN   REJECT   CLASS       REJECT

STLABC2  DEFAULT       956     3,835    0  >99   <1  100          0        0   STLABC2      3,795        0   DEFAULT          0
STLABC3  DEFAULT       956     3,026    0  >99   <1  100          0        0   STLABC3      3,024        0
                            ----------                                                 ----------
TOTAL                          6,861                                           TOTAL        6,819
```

Figure 89 (Part 1 of 4). XCF Activity Report

```
                                                  X C F   A C T I V I T Y
                                                                                                     PAGE 2
         MVS/ESA                  SYSTEM ID 961              DATE 09/09/95              INTERVAL 01.00.000
         SP5.2.0                  RPT VERSION 5.2.0          TIME 08.31.00              CYCLE 1.000 SECONDS


                                                  XCF USAGE BY MEMBER
-------------------------------------------------------------------------------------------------------------------
                                                  XCF USAGE BY MEMBER
           MEMBERS COMMUNICATING WITH STLABC1                                                   MEMBERS ON STLABC1
------------------------------------------------------------------       -----------------------------------------------------
                                                FROM          TO                                             REQ           REQ
                                                 REQ         REQ
                                                FROM          TO                                             REQ           REQ
GROUP       MEMBER              SYSTEM        STLABC1     STLABC1    GROUP       MEMBER                      OUT            IN
                                                FROM          TO                                             REQ           REQ
COFVLFNO    STLABC2             STLABC2            0           0     COFVLFNO    STLABC1                       0             0
            STLABC3             STLABC3            0           0                                      ----------    ----------
                                          ----------  ----------     TOTAL                                     0             0
TOTAL                                              0           0
                                                FROM          TO                                             REQ           REQ
DSNDB0G     DB2INMVSDB2G        STLABC2            0           0     DSNDB0G     DB2INMVSDB1G                  0             0
                                          ----------  ----------                                      ----------    ----------
TOTAL                                              0           0     TOTAL                                     0             0
                                                FROM          TO                                             REQ           REQ
DXRDB0G     DXRDB0G$$DJ2G002    STLABC2          240         240     DXRDB0G     DXRDB0G$$DJ1G001            240           240
                                          ----------  ----------                                      ----------    ----------
TOTAL                                            240         240     TOTAL                                   240           240
                                                FROM          TO                                             REQ           REQ
IXCLO000    M238                STLABC2          569         529     IXCLO000    M237                      1,103         1,063
                                          ----------  ----------                                      ----------    ----------
TOTAL                                            569         529     TOTAL                                 1,103         1,063
                                                FROM          TO                                             REQ           REQ
                                                                     STLABC1     STLABC1$961                   0             0
                                                                                                      ----------    ----------
                                                                     TOTAL                                     0             0

                                                FROM          TO                                             REQ           REQ
STLABC2     STLABC2$962         STLABC2            0           0
                                          ----------  ----------
TOTAL                                              0           0
                                                FROM          TO                                             REQ           REQ
STLABC3     STLABC3$963         STLABC3            0           0
                                          ----------  ----------
TOTAL                                              0           0
                                                FROM          TO                                             REQ           REQ
SYSDAE      STLABC2             STLABC2            0           0     SYSDAE      STLABC1                       0             0
            STLABC3             STLABC3            0           0                                      ----------    ----------
                                          ----------  ----------     TOTAL                                     0             0
TOTAL                                              0           0
                                                FROM          TO                                             REQ           REQ
SYSGRS      STLABC2             STLABC2        3,023       3,023     SYSGRS      STLABC1                   6,046         6,046
            STLABC3             STLABC3        3,023       3,023                                      ----------    ----------
                                          ----------  ----------     TOTAL                                 6,046         6,046
TOTAL                                          6,046       6,046
```

Figure 89 (Part 2 of 4). XCF Activity Report

```
                                        X C F   A C T I V I T Y                                         PAGE 3

         MVS/ESA              SYSTEM ID 961          DATE 09/09/95          INTERVAL 01.00.000
         SP5.2.0              RPT VERSION 5.2.0      TIME 08.31.00          CYCLE 1.000 SECONDS

                                              XCF USAGE BY MEMBER
-----------------------------------------------------------------------------------------------------------------

           MEMBERS COMMUNICATING WITH STLABC1                                        MEMBERS ON STLABC1
-----------------------------------------------------------------    ------------------------------------------------------

                                               REQ         REQ
                                              FROM          TO                                                 REQ        REQ
GROUP      MEMBER           SYSTEM          STLABC1     STLABC1      GROUP      MEMBER                         OUT         IN

SYSIGW00   IGWCLM01STLABC2  STLABC2               0           0      SYSIGW00   IGWCLM01STLABC1                  0          0
           IGWCLM01STLABC3  STLABC3               0           0                                        ----------  ----------
                                         ----------  ----------      TOTAL                                       0          0
TOTAL                                             0           0

SYSIGW01   IGWCLM01STLABC2  STLABC2               0           0      SYSIGW01   IGWCLM01STLABC1                  0          0
           IGWCLM01STLABC3  STLABC3               0           0                                        ----------  ----------
                                         ----------  ----------      TOTAL                                       0          0
TOTAL                                             0           0

SYSIKJBC   STLABC2          STLABC2               0           0      SYSIKJBC   STLABC1                          0          0
           STLABC3          STLABC3               0           0                                        ----------  ----------
                                         ----------  ----------      TOTAL                                       0          0
TOTAL                                             0           0

SYSMCS     STLABC2          STLABC2               2           2      SYSMCS     STLABC1                          4          2
           STLABC3          STLABC3               2           0                                        ----------  ----------
                                         ----------  ----------      TOTAL                                       4          2
TOTAL                                             4           2

SYSMCS2    STLABC2          STLABC2               0           0      SYSMCS2    STLABC1                          0          0
           STLABC3          STLABC3               0           0                                        ----------  ----------
                                         ----------  ----------      TOTAL                                       0          0
TOTAL                                             0           0

SYSRMF     SYSRMF@@STLABC2  STLABC2               0           0      SYSRMF     SYSRMF@@STLABC1                  0          0
                                         ----------  ----------                                        ----------  ----------
TOTAL                                             0           0      TOTAL                                       0          0

SYSWLM     STLABC2          STLABC2               1           1      SYSWLM     STLABC1                          2          2
           STLABC3          STLABC3               1           1                                        ----------  ----------
                                         ----------  ----------      TOTAL                                       2          2
TOTAL                                             2           2
```

Figure 89 (Part 3 of 4). XCF Activity Report

```
                                             X C F  A C T I V I T Y
                                                                                                          PAGE 4
        MVS/ESA                  SYSTEM ID 961           DATE 09/09/95          INTERVAL 01.00.000
        SP5.2.0                  RPT VERSION 5.2.0       TIME 08.31.00          CYCLE 1.000 SECONDS


TOTAL SAMPLES = 60                                  XCF PATH STATISTICS
-------------------------------------------------------------------------------------------------------------------------------------

 2                          OUTBOUND FROM STLABC1                                                   INBOUND TO STLABC1
--------------------------------------------------------------------------------          -----------------------------------------
        T FROM/TO                                                                                  T FROM/TO
 TO     Y DEVICE, OR     TRANSPORT      REQ     AVG Q                                      FROM    Y DEVICE, OR          REQ BUFFERS
 SYSTEM P STRUCTURE      CLASS          OUT     LNGTH    AVAIL   BUSY   RETRY              SYSTEM  P STRUCTURE            IN UNAVAIL

 STLABC2 C 0B14 TO 0B14  DEFAULT      2,456      0.00    2,447      9       0              STLABC2 C 0B16 TO 0B16      1,387       0
         C 0B15 TO 0B15  DEFAULT      1,387      0.00    1,384      3       0                      C 0B17 TO 0B17      2,427       0
 STLABC3 C 0B24 TO 0B24  DEFAULT      1,448      0.00    1,448      0       0              STLABC3 C 0B26 TO 0B26        771       0
         C 0B25 TO 0B25  DEFAULT      1,610      0.02    1,610      0       0                      C 0B27 TO 0B27      2,319       0
*UNKNOWN C 0B34 TO       DEFAULT          0      0.00        0      0       0             *UNKNOWN C      TO 0B36          0       0
         C 0B35 TO       DEFAULT          0      0.00        0      0       0                      C      TO 0B37          0       0
         C 0B44 TO       DEFAULT          0      0.00        0      0       0                      C      TO 0B46          0       0
         C 0B45 TO       DEFAULT          0      0.00        0      0       0                      C      TO 0B47          0       0
         C 0B54 TO       DEFAULT          0      0.00        0      0       0                      C      TO 0B56          0       0
         C 0B55 TO       DEFAULT          0      0.00        0      0       0                      C      TO 0B57          0       0
         C 0B64 TO       DEFAULT          0      0.00        0      0       0                      C      TO 0B66          0       0
         C 0B65 TO       DEFAULT          0      0.00        0      0       0                      C      TO 0B67          0       0
                                     ----------                                                                      ----------
 TOTAL                                6,901                                                TOTAL                       6,904
```

Figure 89 (Part 4 of 4). XCF Activity Report

F.3 RMF – I/O Queuing Activity

```
                    3              I/O  Q U E U I N G   A C T I V I T Y
                                                                                                PAGE  1
        MVS/ESA                   SYSTEM ID 961          DATE 09/09/95          INERVAL 09.37.205
        SP5.2.0                  RPT VERSION 5.2.0       TIME 08.41.39         CYCLE 1.000 SECONDS

TOTAL SAMPLES =   577     IOP   ACTIVITY RATE  AVG Q LNGTH     IODF = C1   CR-DATE: 07/04/95   CR-TIME: 06.55.18      ACT: POR


                          00      611.928       0.19


                   DELAY    % ALL
LCU  CONTENTION      Q     CH PATH  CONTROL UNITS    CHAN     CHPID    % DP   % CU
        RATE       LNGTH    BUSY                     PATHS    TAKEN    BUSY   BUSY


0001     0.106      0.02     0.00  0120              20       3.063    0.00  15.41
                                   0121              2C       3.021    0.00  16.03
0004     0.000      0.00     0.00  0160              20       7.940    0.00   0.04
                                   0161              2C       8.025    0.00   0.02
0005     0.000      0.00     0.00  0180              20       0.404    0.85   0.00
                                                     21       0.269    0.00   0.00
                                                     2C       0.381    0.00   0.00
                                                     2D       0.270    0.00   0.00
0008     0.000      0.00     0.00  01C0              20       5.286    0.00   0.03
                                   01C1              2C       5.036    0.00   0.00
0009     0.000      0.00     0.00  0200              21       4.333    0.00   0.04
                                   0201              2D       4.302    0.00   0.00
000A     0.000      0.00     0.00  0240              24       4.522    0.00   0.00
                                   0241              30       4.369    0.00   0.00
000C     0.042      0.00     0.00  02C0              21       3.118    0.00   7.50
                                   02C1              2D       3.067    0.00   9.65
000D     0.055      0.03     0.00  02D0              21       3.042    0.00   9.62
                                   02D1              2D       3.072    0.00   9.77
000E     0.050      0.03     0.00  02E0              21       3.118    0.00   8.68
                                   02E1              2D       3.257    0.00   9.96
000F     0.000      0.00     0.00  0300              21      12.779    0.00   0.12
                                   0301              2D      13.411    0.00   0.03
0010     0.000      0.00     0.00  0340              24      12.429    0.00   0.14
                                   0341              30      12.114    0.00   0.03
0012     0.000      0.00     0.00  0380              22       1.311    0.00   1.69
                                   0381              2E       1.188    0.00   1.58
0014     0.000      0.00     0.00  0400              22       3.358    0.00   0.00
                                   0401              2E       3.096    0.00   0.06
0015     0.000      0.00     0.00  0440              25       3.290    0.00   0.05
                                   0441              31       3.274    0.00   0.00
0016     0.021      0.08     0.00  0480              22       2.235    0.00   6.11
                                   0481              2E       2.036    0.00   8.77
0017     0.017      0.00     0.00  04A0              22       2.497    0.00   4.44
                                   04A1              2E       2.176    0.00   7.31
0018     0.000      0.00     0.00  04C0              22      12.614    0.00   0.00
                                   04C1              2E      11.916    0.00   0.00
```

Figure 90. RMF – I/O Queuing Activity Report. I/O Queuing Activity

F.4 RMF – Coupling Facility Activity

```
                           C O U P L I N G   F A C I L I T Y   A C T I V I T Y
                                                                                                   PAGE   1
   MVS/ESA               SYSPLEX XESDEV            DATE 06/03/1995          INTERVAL 016.56.000
   SP5.1.0               RPT VERSION 5.1.0         TIME 09.24.00            CYCLE 01.000 SECONDS

-----------------------------------------------------------------------------------------------------------
COUPLING FACILITY NAME = LF02
TOTAL SAMPLES(AVG) =   977  (MAX) =   977  (MIN) =   977
-----------------------------------------------------------------------------------------------------------
       1                            COUPLING  FACILITY  USAGE  SUMMARY
-----------------------------------------------------------------------------------------------------------
STRUCTURE SUMMARY
-----------------------------------------------------------------------------------------------------------

                                                % OF             % OF     AVG     LIST     LIST      LOCK     CACHE TOT
      STRUCTURE                        ALLOC    CF         #     ALL      REQ/    ENTRIES  ELEMENTS  ENTRIES  DIR ENTRY/
TYPE  NAME            STATUS CHG       SIZE     STORAGE    REQ   REQ      SEC     TOT/CUR  TOT/CUR   TOT/CUR  DATA ELEM

LIST  DSNDB0G_SCA     ACTIVE            10M       0.5%     221   0.0%     0.22      16K      32K      N/A       N/A
                                                                                    232      421      N/A       N/A

CACHE DSNDB0G_GBP0    ACTIVE             3M       0.1%   10294   0.8%    10.13     N/A      N/A      N/A      1259
                                                                                   N/A      N/A      N/A       625
      DSNDB0G_GBP1    ACTIVE            14M       0.7%   30215   2.5%    29.74     N/A      N/A      N/A      6189
                                                                                   N/A      N/A      N/A      3249
      DSNDB0G_GBP3    ACTIVE            86M       4.3%  516624  42.1%   508.49     N/A      N/A      N/A        65K
                                                                                   N/A      N/A      N/A        19K
      DSNDB0G_GBP4    ACTIVE            20M       1.0%   20834   1.7%    20.51     N/A      N/A      N/A      7025
                                                                                   N/A      N/A      N/A      4674
      DSNDB0G_GBP5    ACTIVE            73M       3.7%   26185   2.1%    25.77     N/A      N/A      N/A        49K
                                                                                   N/A      N/A      N/A        16K
      DSNDB0G_GBP6    ACTIVE            39M       1.9%  118817   9.7%   116.95     N/A      N/A      N/A        33K
                                                                                   N/A      N/A      N/A      8230
      DSNDB0G_GBP7    ACTIVE            22M       1.1%   69811   5.7%    68.71     N/A      N/A      N/A        16K
                                                                                   N/A      N/A      N/A      4658
      DSNDB0G_GBP8    ACTIVE           130M       6.5%  433518  35.3%   426.69     N/A      N/A      N/A        89K
                                                                                   N/A      N/A      N/A        29K

                                      -------    -----  -------  ------  -------
          STRUCTURE TOTALS              396M     19.9%    1227K    100%   1207.2
```

Figure 91. RMF – Coupling Facility Structure Usage Summary Report. For Coupling Facility Model 9674/R61

F.5 RMF – Coupling Facility Structure Activity

```
                              C O U P L I N G   F A C I L I T Y   A C T I V I T Y
                                                                                                      PAGE   3
   MVS/ESA                  SYSPLEX XESDEV              DATE 06/03/1995           INTERVAL 016.56.000
   SP5.1.0                  RPT VERSION 5.1.0           TIME 09.24.00             CYCLE 01.000 SECONDS

------------------------------------------------------------------------------------------------------------------
COUPLING FACILITY NAME = LF02
------------------------------------------------------------------------------------------------------------------
                                        COUPLING  FACILITY  STRUCTURE  ACTIVITY
------------------------------------------------------------------------------------------------------------------
 1
STRUCTURE NAME = DSNDB0G_SCA       TYPE = LIST
         # REQ    -------------- REQUESTS -------------   -------------- QUEUED REQUESTS --------------
SYSTEM   TOTAL             #     % OF  -SERV TIME(MIC)-           (ARRIVAL  RATE) --- QUEUE TIME(MIC)--
NAME     AVG/SEC          REQ    ALL     AVG   STD_DEV           MIN   AVG   MAX          AVG  STD_DEV

STLABC1     221   SYNC    219   99.1%   500.7     59.9   HPRIO     0   0.0     2  NO SCH 685.5   546.6
           0.22   ASYNC     0    0.0%  3415.0    572.8   LPRIO     0   0.0     0
                  CHNGD     2    0.9%  INCLUDED IN ASYNC  DUMP      0   0.0     0  DUMP     0.0     0.0


------------------------------------------------------------------------------------------------------------------
TOTAL       221   SYNC    219   99.1%   500.7     59.9                             NO SCH 685.5   546.6
           0.22   ASYNC     0    0.0%  3415.0    572.8                             DUMP     0.0     0.0
                  CHNGD     2    0.9%

STRUCTURE NAME = DSNDB0G_GBP0      TYPE = CACHE
         # REQ    -------------- REQUESTS -------------   -------------- QUEUED REQUESTS --------------
SYSTEM   TOTAL             #     % OF  -SERV TIME(MIC)-           (ARRIVAL  RATE) --- QUEUE TIME(MIC)--
NAME     AVG/SEC          REQ    ALL     AVG   STD_DEV           MIN   AVG   MAX          AVG  STD_DEV

STLABC1   10294   SYNC    10K   99.0%   373.7    107.6   HPRIO     0   1.6     5  NO SCH  4880    4128
          10.13   ASYNC     4    0.0%  3032.2   1245.3   LPRIO     0   0.0     0
                  CHNGD    96    0.9%  INCLUDED IN ASYNC  DUMP      0   0.0     0  DUMP     0.0     0.0


------------------------------------------------------------------------------------------------------------------
TOTAL     10294   SYNC    10K   99.0%   373.7    107.6                             NO SCH  4880    4128
          10.13   ASYNC     4    0.0%  3032.2   1245.3                             DUMP     0.0     0.0
                  CHNGD    96    0.9%

STRUCTURE NAME = DSNDB0G_GBP1      TYPE = CACHE
         # REQ    -------------- REQUESTS -------------   -------------- QUEUED REQUESTS --------------
SYSTEM   TOTAL             #     % OF  -SERV TIME(MIC)-           (ARRIVAL  RATE) --- QUEUE TIME(MIC)--
NAME     AVG/SEC          REQ    ALL     AVG   STD_DEV           MIN   AVG   MAX          AVG  STD_DEV

STLABC1   30215   SYNC    30K   98.9%   351.8    101.0   HPRIO     0   5.5    16  NO SCH  5262    3540
          29.74   ASYNC     3    0.0%  3117.7   1361.5   LPRIO     0   0.0     0
                  CHNGD   325    1.1%  INCLUDED IN ASYNC  DUMP      0   0.0     0  DUMP     0.0     0.0


------------------------------------------------------------------------------------------------------------------
TOTAL     30215   SYNC    30K   98.9%   351.8    101.0                             NO SCH  5262    3540
          29.74   ASYNC     3    0.0%  3117.7   1361.5                             DUMP     0.0     0.0
                  CHNGD   325    1.1%
```

Figure 92. RMF – Coupling Facility Structure Usage Detail Report. For Coupling Facility Model 9674/R61

```
                              C O U P L I N G   F A C I L I T Y   A C T I V I T Y
                                                                                                      PAGE   6
   MVS/ESA               SYSPLEX XESDEV          DATE 06/03/1995          INTERVAL 016.56.000
   SP5.1.0               RPT VERSION 5.1.0       TIME 09.24.00            CYCLE 01.000 SECONDS

-------------------------------------------------------------------------------------------------------------------------------
COUPLING FACILITY NAME = LF02
-------------------------------------------------------------------------------------------------------------------------------
                                                   SUBCHANNEL  ACTIVITY
-------------------------------------------------------------------------------------------------------------------------------
        # REQ                               ----------- REQUESTS -----------   -------------- QUEUED REQUESTS --------------
SYSTEM  TOTAL              --BUSY--              #   -SERVICE TIME(MIC)-          (ARRIVAL  RATE) -- QUEUE TIME(MIC) --
NAME    AVG/SEC -- CONFIG --  -COUNTS-           REQ     AVG    STD_DEV            MIN    AVG  MAX           AVG  STD_DEV
 2
STLABC1   1228K SCH GEN    4  PTH    0     SYNC     1200K  265.8      94.6  HPRIO    0   25.4  982  NO SCH  5252   4249
        1208.3  SCH USE    4  SCH   18     ASYNC      202  3045.0     2708  LPRIO    0    0.0    0
                SCH MAX    4               CHANGED  26439  INCLUDED IN ASYNC  DUMP     0    0.0    0  DUMP     0.0    0.0
                PTH        2               UNSUCC       0     0.0      0.0
```

Figure 93. RMF – Coupling Facility Subchannel Activity Summary Report. For Coupling Facility Model 9674/R61

Index

D

E

F

G

H

I

L

M

ITSO Technical Bulletin Evaluation RED000

International Technical Support Organization
DB2 for MVS/ESA Version 4
Data Sharing Performance Topics
December 1995

Publication No. SG24-4611-00

Your feedback is very important to help us maintain the quality of ITSO Bulletins. **Please fill out this questionnaire and return it using one of the following methods:**

- Mail it to the address on the back (postage paid in U.S. only)
- Give it to an IBM marketing representative for mailing
- Fax it to: Your International Access Code + 1 914 432 8246
- Send a note to REDBOOK@VNET.IBM.COM

Please rate on a scale of 1 to 5 the subjects below.
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction ____

Organization of the book	____	Grammar/punctuation/spelling	____
Accuracy of the information	____	Ease of reading and understanding	____
Relevance of the information	____	Ease of finding information	____
Completeness of the information	____	Level of technical detail	____
Value of illustrations	____	Print quality	____

Please answer the following questions:

a) If you are an employee of IBM or its subsidiaries:

Do you provide billable services for 20% or more of your time? Yes____ No____

Are you in a Services Organization? Yes____ No____

b) Are you working in the USA? Yes____ No____

c) Was the Bulletin published in time for your needs? Yes____ No____

d) Did this Bulletin meet your needs? Yes____ No____

If no, please explain:

____

____

What other topics would you like to see in this Bulletin?

____

____

What other Technical Bulletins would you like to see published?

____

Comments/Suggestions: **(THANK YOU FOR YOUR FEEDBACK!)**

______________________________ ______________________________
Name Address

______________________________ ______________________________
Company or Organization

______________________________ ______________________________
Phone No.

Fold and Tape **Please do not staple** Fold and Tape

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM International Technical Support Organization
Department 471/E2
650 Harry Road
San Jose, CA
USA 95120-6099

Fold and Tape **Please do not staple** Fold and Tape

Printed in U.S.A.

SG24-4611-00

Artwork Definitions

id	File	Page	References
ITSLOGO	4611SU	i	i

Table Definitions

id	File	Page	References
TL001C	4611CH04	22	22
TMA36A	4611CH05	39	39
TMA36X	4611CH05	39	39
TMA37A	4611CH05	40	40, 41, 41, 41
TMA37B	4611CH05	40	41, 41, 41
TMA37C	4611CH05	40	41
TMA37X	4611CH05	40	41, 41, 41
TMA38A	4611CH05	42	42
TMA38X	4611CH05	42	42
TMA40A	4611CH05	43	43
TMA40X	4611CH05	43	43
TMA41A	4611CH05	43	43
TMA41X	4611CH05	43	43
TMA39A	4611CH05	43	44, 44
TMA39B	4611CH05	43	44, 44
TMA39X	4611CH05	43	43
TMA45A	4611CH05	44	44, 44
TMA45B	4611CH05	44	44, 44
TMA45X	4611CH05	44	44
TMA46A	4611CH05	45	45, 45
TMA46B	4611CH05	45	45, 45
TMA46X	4611CH05	45	45
TMA42A	4611CH05	52	52, 52, 52
TMA42X	4611CH05	52	52, 52
TMA33A	4611CH05	53	53
TMA33B	4611CH05	53	53
TMA34A	4611CH05	53	53
TMA34B	4611CH05	53	53
TTULT11	4611CH08	92	92, 92, 92, 92
TLCKSZ	4611CH08	110	110

Example Definitions

id	File	Page	References
4611XMP	4611VARS	i	138, 139, 140, 141, 142, 143, 144, 145, 146, 147, 150, 151, 152, 153, 154, 155, 158, 159, 160, 161, 164, 165, 166, 167, 169, 171, 172, 172, 173, 174, 175, 176, 177, 178, 179, 180

Figures

id	File	Page	References
FDI001	4611CH02	8	1 7
FOGROUP	4611CH03	12	2 11
FLLS	4611CH04	23	3 23, 23, 23
FLLHTE2	4611CH04	24	4 24, 24
FLLHTE4	4611CH04	25	5 24, 24
FDG001	4611CH04	29	6 28, 28, 28
FDG002	4611CH04	30	7 28, 28, 28
FLAV1	4611CH04	31	8 31, 32, 32
FLAV2	4611CH04	32	9 32, 32
FDO001	4611CH04	34	10 33, 33, 35
FMA43	4611CH05	46	11 45
FMA35	4611CH05	49	12 48, 50
FMA11	4611CH05	58	13 57
FMA01	4611CH05	58	14 58
FMA02	4611CH05	59	15 58
FMA03	4611CH05	59	16 59
FMA04	4611CH05	60	17 59
FMA12	4611CH05	60	18 60
FMA05	4611CH05	60	19 60
FMA06	4611CH05	61	20 61
FMA07	4611CH05	61	21 61
FMA08	4611CH05	61	22 61
FMA09	4611CH05	63	23 28, 62

FMA10	4611CH05	63	24	63
FMA13	4611CH05	64	25	64
FMA14	4611CH05	65	26	64
FMA15	4611CH05	66	27	65
FMA16	4611CH05	66	28	66
FDM003	4611CH06	74	29	73, 73
FDM004	4611CH06	74	30	73, 73
FDM03A	4611CH06	75	31	75
FDM001	4611CH06	76	32	77, 77, 77
FDM002	4611CH06	77	33	77, 77, 77
FDG006	4611CH07	80	34	79, 79, 79, 83
FDG003	4611CH07	83	35	83, 83
FDG004	4611CH07	85	36	85
FDG005	4611CH07	87	37	86
DISPXCF	4611CH08	93	38	93
FXCFSTR	4611CH08	93	39	93, 93, 111
FXFDL1	4611CH08	94	40	94
DSNDB0G	4611CH08	95	41	95, 109, 109
FTUSTS	4611CH08	96	42	96
TAB1	4611CH08	100	43	99
TAB2	4611CH08	101	44	100, 104, 105
TAB3	4611CH08	102	45	102, 105
TAB5	4611CH08	104	46	103
TAB6	4611CH08	106	47	106
FTUCF1	4611CH08	107	48	106
FTUCF2	4611CH08	108	49	108
FTUIO	4611CH08	108	50	108
FDGLOCK	4611CH08	109	51	109, 110, 110, 110
FTURMCF	4611CH08	110	52	110

FXDGBP	4611CH08	112	53	112
FTUDBH	4611CH08	113	54	113, 115, 116, 117, 125, 127, 128, 128
FTUDGD	4611CH08	115	55	115, 129
FTUDMD	4611CH08	119	56	
FTUGBP	4611CH08	123	57	122
F1CBPX	4611CH08	126	58	125
FSCAXD	4611CH08	132	59	131
FSCAR1	4611CH08	133	60	132
FSCAR3	4611CH08	134	61	133
FTUBL	4611CH08	134	62	134, 135, 135, 135
F1111	4611AP04	138	63	61
F1211	4611AP04	139	64	61
F1114	4611AP04	140	65	63, 66
F1214	4611AP04	141	66	63, 66
F11116	4611AP04	142	67	
F12116	4611AP04	143	68	
F11129	4611AP04	144	69	
F12129	4611AP04	145	70	
F11130	4611AP04	146	71	64
F12130	4611AP04	147	72	65
F0111	4611AP04	150	73	59
F2211	4611AP04	151	74	
F0114	4611AP04	152	75	59, 60
F2214	4611AP04	153	76	
F01116	4611AP04	154	77	
F22116	4611AP04	155	78	
F1BHI	4611AP05	158	79	
F1CHI	4611AP05	159	80	
F1BDS	4611AP05	160	81	
F1CDS	4611AP05	161	82	
F1DHI	4611AP05	164	83	
F1EHI	4611AP05	165	84	
F1DDS	4611AP05	166	85	
F1EDS	4611AP05	167	86	
FAXFDL	4611AP06	169	87	

FRMFCF1	4611AP06	171	88	
FXCFACT	4611AP06	173	89	
FAIOQ	4611AP06	177	90	
FSCARM1	4611AP06	178	91	
FSCARM3	4611AP06	179	92	134, 134
FSCARM4	4611AP06	180	93	134, 134

Headings

id	File	Page	References	
NOTICES	4611FM	xiii	Special Notices	ii
BIBL	4611PREF	xvi	Related Publications	xiii
INTRO	4611CH01	1	Chapter 1, Introduction	xv
HSO	4611CH02	7	Chapter 2, Sysplex Overview	xv
HSIPLX	4611CH02	7	2.1, Sysplex	
HSIPCF	4611CH02	8	2.1.1, Coupling Facility	
HSIPDB	4611CH02	9	2.1.2, DB2 in a Sysplex	
HDSO	4611CH03	11	Chapter 3, Data Sharing Overview	xv
HDSOA	4611CH03	11	3.1, Architecture	
HDSAG	4611CH03	11	3.1.1, Group	
HDSAM	4611CH03	12	3.1.2, Member	
HDSOGBP	4611CH03	14	3.2, Group Buffer Pools	13, 47
HDSOGB0	4611CH03	15	3.2.1, Inter-DB2 Read-Write Interest	
HDSOGB1	4611CH03	15	3.2.2, Group Buffer Pool Dependence	
HDSOFC	4611CH03	15	3.2.3, Force at Commit	15
HDSOGB3	4611CH03	16	3.2.4, Buffer Pool Coherency	14
HDSOGB4	4611CH03	16	3.2.5, Group Buffer Pool Castout	
HDSOL	4611CH03	17	3.3, Global Locking	47
HDSLGM	4611CH03	17	3.3.1, Global Lock Manager	
HDSLCO	4611CH03	17	3.3.2, Contentions	
HLCK	4611CH04	19	Chapter 4, Locking in a Data Sharing Group	xv, 17
HLGBL	4611CH04	19	4.1, Global Locking	
HLGLL	4611CH04	19	Logical Locks	
HLGPL	4611CH04	19	Physical Locks	
HLLL	4611CH04	20	4.1.1, Table Space Logical Locks	
HLHL	4611CH04	20	Explicit Hierarchical Locking	20
HLDRAIN	4611CH04	20	Drains and Claims	

47

HMACFA1	4611CH05	44	Coupling Facility Activity for Type 1 Index and Page Locking	47
HMACFAR	4611CH05	44	Coupling Facility Activity for Type 2 Index and Row Locking	
HMADSEN	4611CH05	45	5.2, Data Sharing Environments.	
HMADSCH	4611CH05	45	Some Migration Paths in Data Sharing Environments	
HMADSE1	4611CH05	45	5.2.1, Migrating within 1-way Environments	
HMADSE2	4611CH05	47	5.2.2, Migrating to 2-way Environments	
HMADSE3	4611CH05	48	5.2.3, Migrating from 2-way to 3-way Environment	
HMA2IP	4611CH05	48	5.3, IBM Relational Warehouse Workload Performance Tests	
HTE	4611CH05	48	5.3.1, IRWW Test Environment	
HMACF	4611CH05	50	5.3.2, Configuration and Definitions for the Performance Test	
HMACFPO	4611CH05	50	Coupling Facility Policies	
HMAGBPS	4611CH05	51	Virtual Buffer Pool and Global Buffer Pool Settings	
HMADIRW	4611CH05	52	5.3.3, Description of the IRWW Workload	
HMA3TD	4611CH05	53	5.3.4, Transactions defined for IRWW	
HMATXT	4611CH05	53	Transactions vs. Tables Matrix	
HMA2PC	4611CH05	54	5.4, Performance Calculations for Data Sharing Measurements	
HMAFABR	4611CH05	54	5.4.1, Abbreviations Used in Calculations	
HMAF3ME	4611CH05	55	5.4.2, Calculation formulas for Each Member	
HMAFCBU	4611CH05	55	CPU %Busy	
HMAFETR	4611CH05	55	External Throughput Rate	
HMAFITR	4611CH05	55	Formula to Calculate Internal Throughput Rate	
HMAFXLR	4611CH05	55	Formula to Calculate Total XES Lock Requests (XLR)	
HMACONT	4611CH05	55	Formula to Calculate Total Global Lock Suspensions	
HMAFGBR	4611CH05	56	Formula to Calculate Total GBP Accesses (GBPA)	
HMA3FRU	4611CH05	56	5.4.3, Calculation formulas for Each Run	
HMAWGRO	4611CH05	56	Formula to Calculate Data Sharing Capacity Delta	72
HMAWCEF	4611CH05	56	Formula to Calculate Coupling Efficiency	
HMAWPLI	4611CH05	56	Formula to Calculate Increased CPU Time	
HMAWTCP	4611CH05	56	Formula to Calculate Total Contention Percentage	90
HMAWTFS	4611CH05	57	Formula to Calculate Total False Contention Percentage	
HMAWXLC	4611CH05	57	Formula to Calculate Number of XES Lock Requests per Commit	
HMAWGBC	4611CH05	57	Formula to Calculate Number of GBP Accesses per Commit	
HMA31WP	4611CH05	57	5.4.4, Performance Indicators for 1WP (1-way, Type 2 Index and Page Locking)	54
HMA1WPB	4611CH05	58	Obtaining CPU %Busy for 1WP	
HMA1WPE	4611CH05	58	Obtaining ETR for 1WP	
HMA1WPI	4611CH05	59	Calculating ITR for 1WP	62, 62
HMA1WPX	4611CH05	59	Obtaining Total Number of XES Lock Requests for 1WP	
HMA1WPT	4611CH05	59	Obtaining Total Number of GBP Accesses for 1WP	
HMA1WPS	4611CH05	59	Obtaining Total Number of Global Lock Suspensions for 1WP	

				117, 118
HTUPREF	4611CH08	125	8.2.3.1, Using DB2 PM Report for Buffer Pool Statistics	121
HTUDINC	4611CH08	126	8.3, Sizing the Group Buffer Pool	126
S07212	4611CH08	126	8.3.1, DB2 Member Virtual Buffer Pool	
S07211	4611CH08	127	8.3.2, Coupling Facility Structure Size	
HTU203	4611CH08	127	8.3.3, Group Buffer Pool Directory Entries	129
HTUGDP	4611CH08	128	8.3.4, Group Buffer Pool Data Pages	84, 129
S07213	4611CH08	130	8.4, DB2 Castout Process	
HGBPCCT	4611CH08	130	8.4.1, Group Buffer Pool Class Castout Threshold	51
HTUGBCI	4611CH08	131	8.4.3, Group Buffer Pool Castout Interval	122
SCA	4611CH08	131	8.5, Analyzing the Use of the Shared Communication Area	
HSCAXD	4611CH08	131	8.5.1, Using the XCF DISPLAY Command	
HSCARM	4611CH08	132	8.5.2, Using RMF Coupling Facility Activity Report	
HXCF	4611CH08	134	8.6, Analyzing the MVS XCF Communication	
APMGBP	4611AP04	137	Appendix A, DB2 PM 2-Way Statistics Report (2WP)	96, 122, 126
H2111	4611AP04	138	A.1, DB2 PM Statistics Report: 2-Way First CPC, SQL Activity	
H2211	4611AP04	139	A.2, DB2 PM Statistics Report: 2-Way Second CPC, SQL Activity	
H2114	4611AP04	140	A.3, DB2 PM Statistics Report: 2-Way First CPC, Locking Activity	
H2214	4611AP04	141	A.4, DB2 PM Statistics Report: 2-Way Second CPC, Locking Activity	
H21116	4611AP04	142	A.5, DB2 PM Statistics Report: 2-Way First CPC, BP5 General	
H22116	4611AP04	143	A.6, DB2 PM Statistics Report: 2-Way Second CPC, BP5 General	
H21129	4611AP04	144	A.7, DB2 PM Statistics Report: 2-Way First CPC, Group Buffer Pools GBP7	
H22129	4611AP04	145	A.8, DB2 PM Statistics Report: 2-Way Second CPC, Group Buffer Pools GBP7	
H21130	4611AP04	146	A.9, DB2 PM Statistics Report: 2-Way First CPC, Group Buffer Pool Totals	
H22130	4611AP04	147	A.10, DB2 PM Statistics Report: 2-Way Second CPC, Group Buffer Pool Totals	
APM1BP	4611AP04	149	Appendix B, DB2 PM 1-Way and 2-Way Statistics Reports (1WP and 2WP)	
H1111	4611AP04	150	B.1, DB2 PM Statistics Report: 1-Way CPC, SQL Activity	
H1211	4611AP04	151	B.2, DB2 PM Statistics Report: 2-Way First CPC, SQL Activity	
H1114	4611AP04	152	B.3, DB2 PM Statistics Report: 1-Way CPC, Locking Activity	
H1214	4611AP04	153	B.4, DB2 PM Statistics Report: 2-Way First CPC, Locking Activity	
H11116	4611AP04	154	B.5, DB2 PM Statistics Report: 1-Way CPC, BP5 General	
H12116	4611AP04	155	B.6, DB2 PM Statistics Report: 2-Way First CPC, BP5 General	
APMACB	4611AP05	157	Appendix C, DB2 PM 2-Way Accounting Report (2WP)	
HABHI	4611AP05	158	C.1, DB2 PM Accounting Report: 2-Way First CPC, Highlights Total	
HACHI	4611AP05			

id	File	Page	References
		159	C.2, DB2 PM Accounting Report: 2-Way Second CPC, Highlights Total
HABDS	4611AP05	160	C.3, DB2 PM Accounting Report: 2-Way First CPC, Data Sharing Total
HACDS	4611AP05	161	C.4, DB2 PM Accounting Report: 2-Way Second CPC, Data Sharing Total
APMACC	4611AP05	163	Appendix D, DB2 PM 1-Way and 2-Way Accounting Reports (1WP and 2WP)
HADHI	4611AP05	164	D.1, DB2 PM Accounting Report: 1-Way CPC, Highlights Total
HAEHI	4611AP05	165	D.2, DB2 PM Accounting Report: 2-Way First CPC, Highlights Total
HADDS	4611AP05	166	D.3, DB2 PM Accounting Report: 1-Way CPC, Data Sharing Total
HAEDS	4611AP05	167	D.4, DB2 PM Accounting Report: 2-Way First CPC, Data Sharing Total
HAXFDL	4611AP06	169	Appendix E, MVS XCF DISPLAY Lock Structure
ARMFCF	4611AP06	171	F.1, RMF – Lock Structure Activities 106, 107, 110
AXCFRMF	4611AP06	172	F.2, RMF – XCF Activities 108, 109, 134
HRMFIOQ	4611AP06	176	F.3, RMF – I/O Queuing Activity 108
XSCARMF	4611AP06	177	F.4, RMF – Coupling Facility Activity 133, 133
ASCARM3	4611AP06	178	F.5, RMF – Coupling Facility Structure Activity

Spots

id	File	Page	References
SMFCL	4611CH06	73	(no text) 75, 76
STUMEM	4611CH08	94	Size of a lock table hash entry

Tables

id	File	Page	References
TL001	4611CH04	22	1 15, 22, 22
TMA36	4611CH05	39	2 39
TMA37	4611CH05	41	3 40, 41
TMA38	4611CH05	42	4 42, 72, 87, 88
TMA40	4611CH05	43	5 43, 87
TMA41	4611CH05	43	6 43, 88
TMA39	4611CH05	44	7 43, 44, 44, 67, 81
TMA45	4611CH05	44	8 44
TMA46	4611CH05	45	9 44
TMA44	4611CH05		

		51	10	50
TTUBP	4611CH05	51	11	51, 128, 128
TMA42	4611CH05	52	12	52
TMA33	4611CH05	53	13	53
TMA34	4611CH05	53	14	53
TTULT1	4611CH08	92	15	92
TTULT2	4611CH08	92	16	
TLCKSZ1	4611CH08	110	17	110, 110, 111, 112

Processing Options

Runtime values:

Document fileid	SG244611 SCRIPT
Document type	USERDOC
Document style	IBMXAGD
Profile	EDFPRF30
Service Level	0029
SCRIPT/VS Release	4.0.0
Date	95.12.21
Time	16:33:17
Device	3820A
Number of Passes	4
Index	YES
SYSVAR D	YES
SYSVAR G	INLINE
SYSVAR S	OFF
SYSVAR V	ITSCEVAL

Formatting values used:

Annotation	NO
Cross reference listing	YES
Cross reference head prefix only	NO
Dialog	LABEL
Duplex	YES
DVCF conditions file	(none)
DVCF value 1	(none)
DVCF value 2	(none)
DVCF value 3	(none)
DVCF value 4	(none)
DVCF value 5	(none)
DVCF value 6	(none)
DVCF value 7	(none)
DVCF value 8	(none)
DVCF value 9	(none)
Explode	NO
Figure list on new page	YES
Figure/table number separation	YES
Folio-by-chapter	NO
Head 0 body text	Part
Head 1 body text	Chapter
Head 1 appendix text	Appendix
Hyphenation	NO
Justification	NO
Language	ENGL
Layout	OFF
Leader dots	YES
Master index	(none)
Partial TOC (maximum level)	4
Partial TOC (new page after)	INLINE
Print example id′s	NO
Print cross reference page numbers	YES
Process value	(none)
Punctuation move characters	.,
Read cross-reference file	(none)
Running heading/footing rule	NONE
Show index entries	NO
Table of Contents (maximum level)	3
Table list on new page	YES

Title page (draft) alignment .... RIGHT
Write cross-reference file .... (none)

Imbed Trace

Page 0	4611SU
Page 0	4611VARS
Page 0	4611FM
Page i	4611EDNO
Page ii	4611ABST
Page xiii	4611SPEC
Page xiii	4611TMKS
Page xiv	4611PREF
Page xvii	4611ACKS
Page xviii	4611CH01
Page 5	4611CH02
Page 9	4611CH03
Page 17	4611CH04
Page 37	4611CH05
Page 67	4611CH06
Page 78	4611CH07
Page 90	4611CH08
Page 136	4611AP04
Page 138	4611111
Page 139	4611211
Page 140	4611114
Page 141	4611214
Page 142	46111116
Page 143	46112116
Page 144	46111129
Page 145	46112129
Page 146	46111130
Page 147	46112130
Page 150	4611011
Page 151	4611111
Page 152	4611014
Page 153	4611114
Page 154	46110116
Page 155	46111116
Page 155	4611AP05
Page 158	4611BHI
Page 159	4611CHI
Page 160	4611BDS
Page 161	4611CDS
Page 164	4611AHI
Page 165	4611BHI
Page 166	4611ADS
Page 167	4611BDS
Page 167	4611AP06
Page 185	4611EVAL
Page 185	RCFADDR
Page 185	ITSCADDR FILE
Page 187	RCFADDR
Page 187	ITSCADDR FILE