

IBM DB2 Universal Database



# System Monitor Guide and Reference

*Version 5*



IBM DB2 Universal Database



# System Monitor Guide and Reference

*Version 5*

Before using this information and the product it supports, be sure to read the general information under Appendix E, "Notices" on page 283.

This document contains proprietary information of IBM. It is provided under a license agreement and is protected by copyright law. The information contained in this publication does not include any product warranties and any statements provided in this manual should not be interpreted as such.

Order publications through your IBM representative or the IBM branch office serving your locality or by calling 1-800-879-2755 in U.S. or 1-800-IBM-4YOU in Canada.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© **Copyright International Business Machines Corporation 1993, 1997. All rights reserved.**

Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

---

# Contents

<b>About This Book</b> . . . . .	vii
Who Should Use This Book . . . . .	vii
How This Book is Structured . . . . .	vii
Conventions . . . . .	viii
<b>Chapter 1. Introducing the Database System Monitor</b> . . . . .	1
Database System Monitor Capabilities . . . . .	1
<b>Chapter 2. Using the Database System Monitor</b> . . . . .	3
Database Manager Maintains Operation and Performance Data . . . . .	3
Monitor Switches Control Data Collected by the Database Manager . . . . .	3
Accessing Monitor Data . . . . .	4
Snapshot Monitoring . . . . .	4
Authority Required for Snapshot Monitoring . . . . .	6
Snapshot Monitor Interface . . . . .	7
Information Available by Taking Snapshots . . . . .	7
Snapshot Uses an Instance Connection . . . . .	9
Availability of Snapshot Monitor Data . . . . .	10
Event Monitors . . . . .	10
Authority Required for Event Monitoring . . . . .	15
Using Event Monitors . . . . .	15
Querying the State of an Event Monitor . . . . .	17
Information Available from Event Monitors . . . . .	17
Using Pipe Event Monitors . . . . .	18
When Counters are Initialized . . . . .	20
Resetting Monitor Data . . . . .	21
System Monitor Memory Requirements - (mon_heap_sz) . . . . .	23
Partitioned Database Considerations . . . . .	23
Taking a Snapshot on Multi-node Systems . . . . .	23
Using Event Monitors on Multi-node Systems . . . . .	25
Monitoring Subsections . . . . .	26
DB2 Productivity Tools . . . . .	27
System Monitor Definitions . . . . .	28
<b>Chapter 3. Database System Monitor Data Elements</b> . . . . .	31
How to Read the Data Element Tables . . . . .	32
Element Types . . . . .	33
Server Identification and Status . . . . .	33
Start Database Manager Timestamp . . . . .	34
Configuration NNAME at Monitoring (Server) Node . . . . .	34
Server Instance Name . . . . .	35
Database Manager Type at Monitored (Server) Node . . . . .	35
Server Product/Version ID . . . . .	36
Server Version . . . . .	36
Service Level . . . . .	37

Server Operating System . . . . .	37
Product Name . . . . .	38
Product Identification . . . . .	38
Status of DB2 Instance . . . . .	39
Database Identification and Status . . . . .	39
Database Name . . . . .	39
Database Path . . . . .	40
Database Activation Timestamp . . . . .	41
Database Deactivation Timestamp . . . . .	41
Status of Database . . . . .	42
Catalog Node Network Name . . . . .	42
Database Location . . . . .	43
Catalog Node Number . . . . .	43
Last Backup Timestamp . . . . .	44
Application Identification and Status . . . . .	44
Application Handle (agent ID) . . . . .	45
Application Status . . . . .	46
ID of Code Page Used by Application . . . . .	48
Application Status Change Time . . . . .	48
Application Name . . . . .	49
Application ID . . . . .	50
Sequence Number . . . . .	52
Authorization ID . . . . .	52
Configuration NNAME of Client . . . . .	53
Client Product/Version ID . . . . .	53
Database Alias Used by Application . . . . .	54
Host Product/Version ID . . . . .	55
Outbound Application ID . . . . .	55
Outbound Sequence Number . . . . .	56
User Login ID . . . . .	56
DRDA Correlation Token . . . . .	57
Client Process ID . . . . .	57
Client Operating Platform . . . . .	58
Client Communication Protocol . . . . .	58
Database Country Code . . . . .	59
Application Agent Priority . . . . .	59
Application Priority Type . . . . .	60
User Authorization Level . . . . .	60
Coordinating Node . . . . .	61
Connection Request Start Timestamp . . . . .	62
Connection Request Completion Timestamp . . . . .	62
Previous Unit of Work Completion Timestamp . . . . .	62
Unit of Work Start Timestamp . . . . .	63
Unit of Work Stop Timestamp . . . . .	64
Unit of Work Completion Status . . . . .	65
Previous Transaction Stop Time . . . . .	65
Application Idle Time . . . . .	66
DB2 Agent Information . . . . .	66

Database Manager Configuration . . . . .	66
Agents and Connections . . . . .	67
Sort . . . . .	78
Fast Communication Manager . . . . .	84
Database Configuration . . . . .	90
Buffer Pool Activity . . . . .	90
Non-buffered I/O Activity . . . . .	114
Catalog Cache . . . . .	119
Package Cache . . . . .	122
Database Heap . . . . .	126
Logging . . . . .	127
Database and Application Activity . . . . .	130
Locks and Deadlocks . . . . .	130
Lock Wait Information . . . . .	140
Rollforward Monitoring . . . . .	147
Table Activity . . . . .	149
SQL Cursors . . . . .	160
SQL Statement Activity . . . . .	164
SQL Statement Details . . . . .	174
Subsection Details . . . . .	183
Intra-query Parallelism . . . . .	189
CPU Usage . . . . .	191
Snapshot Monitoring Elements . . . . .	192
<b>Chapter 4. Event Monitor Output . . . . .</b>	<b>195</b>
Output Stream Format . . . . .	195
Matching Event Records with Their Application . . . . .	200
File Event Monitor Buffering . . . . .	201
Blocked Event Monitors . . . . .	202
Non-Blocked Event Monitors . . . . .	202
File Event Monitor Target . . . . .	202
Programming to Read an Event Monitor Trace . . . . .	204
Reading the Data Stream . . . . .	205
Swapping Bytes in Numerical Values . . . . .	206
Reading the Event Records . . . . .	206
Reading the Log Header . . . . .	208
Printing Event Records . . . . .	209
Reading Events from a FILE Trace . . . . .	211
<b>Appendix A. Database System Monitor Interfaces . . . . .</b>	<b>213</b>
CREATE EVENT MONITOR Command and SQL . . . . .	214
db2eva - Event Analyzer Command . . . . .	222
db2evmon - Event Monitor Trace Formatter Command . . . . .	224
DROP EVENT MONITOR Command and SQL . . . . .	226
EVENT_MON_STATE SQL Function . . . . .	227
GET DATABASE MANAGER MONITOR SWITCHES Command . . . . .	228
GET MONITOR SWITCHES Command . . . . .	230
GET SNAPSHOT Command . . . . .	232

LIST ACTIVE DATABASES Command	235
LIST APPLICATIONS - Command	237
LIST DCS APPLICATIONS - Command	239
RESET MONITOR Command	241
SET EVENT MONITOR STATE Command and SQL	242
sqlmon - Get/Update Monitor Switches API	244
sqlmonss - Get Snapshot API	248
sqlmonsz - Estimate Size Required for sqlmonss() Output Buffer API	260
sqlmrset - Reset Monitor API	263
UPDATE MONITOR SWITCHES Command	266
<b>Appendix B. Parallel Edition Version 1.2 Users</b>	269
API Changes	270
Obsolete Commands	270
<b>Appendix C. DB2 Version 1 sqlestat Users</b>	271
<b>Appendix D. How the DB2 Library Is Structured</b>	273
SmartGuides	273
Online Help	274
DB2 Books	276
About the Information Center	280
<b>Appendix E. Notices</b>	283
Trademarks	283
Trademarks of Other Companies	284
<b>Index</b>	285
<b>Contacting IBM</b>	287



---

## About This Book

Your DB2 Database Manager is instrumented to gather data on its operation and performance. You can use this data to:

- Monitor database activities
- Assist in problem determination
- Analyze performance
- Help configure the system.

The DB2 DBMS function that collects this data is called the database system monitor. This book describes how to use the database system monitor.

Various tools allow users to exploit the strengths of the database system monitor with minimal explicit knowledge of its associated commands, APIs, or data formats. Some of these tools, for example the Control Center, are described briefly, but for detailed information you should refer to the *Administration Getting Started*.

---

## Who Should Use This Book

This book is for any users who require an understanding of the operation of the DB2 database system monitor, including how to program to its interface.

It is intended for database administrators, system administrators, security administrators and system operators who are maintaining a database accessed by local or remote clients. It is also for software developers who are interested in building software tools that use the DB2 database system monitor to assist in these administrative functions.

---

## How This Book is Structured

This book starts with a description of the database system monitor and then details the data that you can collect with it.

Chapter 1, *Introducing the Database System Monitor*, introduces the database system monitor and describes its capabilities.

Chapter 2, *Using the Database System Monitor*, describes the information that is available from the database system monitor: how to collect it and how to work with it.

Chapter 3, *Database System Monitor Data Elements*, provides details of the information elements that you can collect with the database system monitor.

Chapter 4, *Event Monitor Output*, is for programmers who want to write applications that read records from an event monitor trace.

Appendix A, *Database System Monitor Interfaces*, contains detailed descriptions of the commands, SQL statements, APIs, and tools that you may need to use with the database system monitor. Detailed information is provided for each API routine.

Appendix B, Parallel Edition Version 1.2 Users, is intended for DB2 Parallel Edition Version 1.2. users of database system monitor who are upgrading their system to DB2 Version 5.

Appendix C, DB2 Version 1 sqlestat Users, is intended for DB2 Version 1 sqlestat users.

Appendix D, How the DB2 Library Is Structured describes the DB2 library; including books and online help.

Appendix E, Notices contains notice and trademark information.

---

## Conventions

You will find this book easier to use if you look for these conventions:

- **Boldface type** indicates an important item or concept
- *Italics type* indicates new terms, data elements, configuration parameters, or book titles.
- Monospace type indicates an example of text that is displayed on the screen or contained in a file.
- UPPERCASE TYPE indicates a file name, command name, or acronym.

Text in examples can be black or a lighter type.

db2 commands and output associated with the database system monitor are in black type  
other db2 commands used are in lighter type

---

## Chapter 1. Introducing the Database System Monitor

This chapter gives you a brief overview of the database system monitor's capabilities. It also discusses the integral role that the database system monitor plays in monitoring database activity and performance.

If you want to get started quickly, read this chapter and Chapter 2, "Using the Database System Monitor" on page 3. The information in these two chapters, combined with the reference material in Appendix A, "Database System Monitor Interfaces" on page 213, provides the information required to use the database system monitor.

Chapter 3, "Database System Monitor Data Elements" on page 31 provides complete details on all the data available with the database system monitor.

---

### Database System Monitor Capabilities

The capabilities of the database system monitor opens several possibilities:

- **Activity monitoring**

For example, using the database system monitor you can obtain:

- The list of database connections:
  - The status of each connection.
  - The SQL that each is executing.
  - The locks that each holds.
- The tables being accessed and the number of rows read and written for each.

You can also track the progression of a query or application using information, such as:

- The cursors that are currently open for this application.
- The number of rows read or CPU consumed (if available from the operating system) by this application.
- How long each query has been running.
- How long an application has been idle.

- **Problem determination**

You can collect data to help diagnose the cause of poor system and application performance. For example:

- By tracing deadlocks you can determine conflicts between applications that lead to poor overall system performance.
- By looking at the amount of time applications spent waiting for locks and which application is holding these locks you can identify applications that fail to commit their transactions, a common cause of poor system performance.

- **Performance analysis**

You can use the information available to analyze the performance of individual applications or SQL queries. For example, you can monitor for:

- The CPU consumed by each individual statement or application.
- The time it takes to run a statement.
- The number of rows read and returned.
- The use of database resources, such as buffer pool, prefetchers, and SQL cache.

These run-time metrics are useful in tuning queries for optimal utilization of your database resources. Modifying a query or certain system parameters can result in dramatic performance improvements. The impact of your modifications can be measured with the database system monitor.

You can also track the usage of indexes and tables, and in a partitioned database, the progression of a query on each partition. Adding indices or repartitioning the data often results in significant performance improvements.

Carrying out some these performance analysis tasks may also require input that is obtained from the operating system, such as system load or the amount of free storage, or from other DB2 tools such as the **db2 explain facility**. For example, the db2expln application lets you analyze the *access plan* generated by the SQL compiler, which can then be compared with the run-time information available from the database system monitor.

- **System configuration**

You can assemble the information necessary to evaluate and tune the effectiveness of your database manager and database configuration.

You can use the database system monitor to help monitor, tune, and manage your databases whether they are local or remote.

---

## Chapter 2. Using the Database System Monitor

This chapter describes the data that is available from the DB2 Version 5 database system monitor. It explains how you can either take a **snapshot** of this data, or request the database manager to log information when certain **events** take place.

It describes the types of snapshots that you can take, and how they can be taken using CLP (command line processor) commands or APIs (application programming interfaces). It details the types of event monitors that can be used for data collection, and how to collect that information using commands or tools that come with DB2.

---

### Database Manager Maintains Operation and Performance Data

Built into the database manager is the ability to collect data about its operation and performance, and that of the applications using it. The database manager maintains information at the following levels:

- Database manager
- Database
- Application (database connection)
- Table
- Table space
- Buffer pool
- Transaction
- Statement
- Subsection

Collecting some of this data introduces some processing overhead. For example, in order to calculate the execution time of an SQL statement, the database manager must make a call to the operating system to obtain timestamps before and after statement execution. These types of system calls are generally expensive. In order to minimize the overhead involved in maintaining monitoring information, **monitor switches** control the collection of potentially expensive data by the database manager.

### Monitor Switches Control Data Collected by the Database Manager

The database system monitor will always collect some basic information, but you can use the switches to govern the amount of expensive data collected. Monitor switches can be set:

- **Explicitly**, this is usually done using the UPDATE MONITOR SWITCHES command.

You can also set these switches in the database manager configuration file if you want data collection to start from the moment the server is started. You should note that setting switches in this way means that they cannot be turned off without stopping the database management system. Switches are explained in "Resetting Monitor Data" on page 21. For more information on configuration see the *dft\_monswitches* configuration parameters in the *Administration Guide*.

- **Implicitly**, when an event monitor is activated. Event monitors are explained in “Event Monitors” on page 10.

To see if your database manager is currently collecting any monitor data issue the command:

```
db2 get database manager monitor switches
```

The resulting output indicates the database manager switch settings and the time that they were turned on.

```

DBM System Monitor Information Collected

Buffer Pool Activity Information (BUFFERPOOL) = OFF
Lock Information (LOCK) = OFF
Sorting Information (SORT) = ON 04-18-1997 10:11:01.738400
SQL Statement Information (STATEMENT) = OFF
Table Activity Information (TABLE) = OFF
Unit of Work Information (UOW) = OFF

```

In this example, in addition to collecting basic-level information, the database manager is collecting all information under control of the sort switch.

## Accessing Monitor Data

There are two ways to access the monitor data collected by the database manager:

- **Snapshot monitoring**

Taking a snapshot gives you information for a specific point in time. A snapshot is a picture of the current state of activity in the database manager for a particular object or group of objects.

- **Event monitors**

You can request the database manager to automatically log monitor data to files or a named pipe when specific events occur. This allows you to collect information about transient events that are difficult to monitor through snapshots, such as deadlocks and transaction completions.

---

## Snapshot Monitoring

The snapshot monitor provides two categories of information for each level being monitored:

- **State**

This includes information such as:

- the current status of the database
- information on the current or most recent unit of work
- the list of locks being held by an application
- the status of an application
- the current number of connections to a database

- the most recent SQL statement performed by an application
- run-time values for configurable system parameters.
- Counters

These accumulate counts for activities from the time monitoring started until the time a snapshot is taken. Such as:

- the number of deadlocks that have occurred
- the number of transactions performed on a database
- the amount of time an application has waited on locks.

For example, you can obtain a list of the locks held by applications connected to a database by taking a database lock snapshot. First, turn on the LOCK switch (UPDATE MONITOR SWITCHES), so that the time spent waiting for locks is collected.

```
db2 connect to sample
db2 update monitor switches using LOCK on
db2 +c list tables for all          # this command will require locks
                                   # on the database catalogs
db2 get snapshot for locks on sample
```

**Note:** You can create and populate the sample database by running `sql11ib/misc/db2saml`.

Issuing the GET SNAPSHOT command returns the following.

Database Lock Snapshot					
Database name	= SAMPLE				
Database path	= /home/bourbon/bourbon/NODE0000/SQL00005/				
Input database alias	= SAMPLE				
Locks held	= 6				
Applications currently connected	= 1				
Applications currently waiting on locks	= 0				
Snapshot timestamp	= 04-11-1997 10:40:29.976539				
Application handle	= 1				
Application ID	= LOCAL.bourbon.970411143813				
Sequence number	= 0001				
Application name	= db2bp_32				
Authorization ID	= BOURBON				
Application status	= UOW Waiting				
Status change time	= Not Collected				
Application code page	= 850				
Locks held	= 6				
Total wait time (ms)	= 0				
Object Type	Tablespace Name	Table Schema	Table Name	Mode	Status
Row	SYSCATSPACE	SYSIBM	SYSTABLES	NS	Granted
Table	SYSCATSPACE	SYSIBM	SYSTABLES	IS	Granted
Table	SYSCATSPACE	SYSIBM	SYSTABLESPACES	S	Granted
Row	SYSCATSPACE	SYSIBM	SYSPLAN	S	Granted
Table	SYSCATSPACE	SYSIBM	SYSPLAN	IS	Granted
Internal				S	Granted

Figure 1. Results of GET SNAPSHOT FOR LOCKS Command

From this snapshot, you can see that there is currently one application connected to the SAMPLE database, and it is holding six locks.

```
Locks held = 6
Applications currently connected = 1
```

Note that the time (Status change time) when the Application status became UOW Waiting is returned as Not Collected, because the UOW switch is OFF.

The lock snapshot also returns the total time spent waiting for locks (so far), by applications connected to this database.

```
Total wait time (ms) = 0
```

This is an example of an accumulating counter. “Resetting Monitor Data” on page 21 explains how counters can be reset to zero.

## Authority Required for Snapshot Monitoring

To perform any of the snapshot monitor tasks, you must have SYSMAINT, SYSCTRL, or SYSADM authority for the database manager instance that you wish to monitor.



## Snapshot Monitor Interface

Snapshot monitoring is invoked using the following application programming interfaces (APIs):

<b>sqlmon()</b>	set or query monitor switch settings
<b>sqlmonrset()</b>	reset system monitor counters
<b>sqlmonss()</b>	take a snapshot
<b>sqlmonsz()</b>	estimate the size of the data that would be returned for a particular invocation of sqlmonss()

The Command Line Processor (CLP) provides a convenient command-based front-end to the snapshot APIs. For example, the GET SNAPSHOT command invokes the sqlmonss() API. Appendix A, "Database System Monitor Interfaces" on page 213 contains detailed information on the commands and APIs associated with the database system monitor.

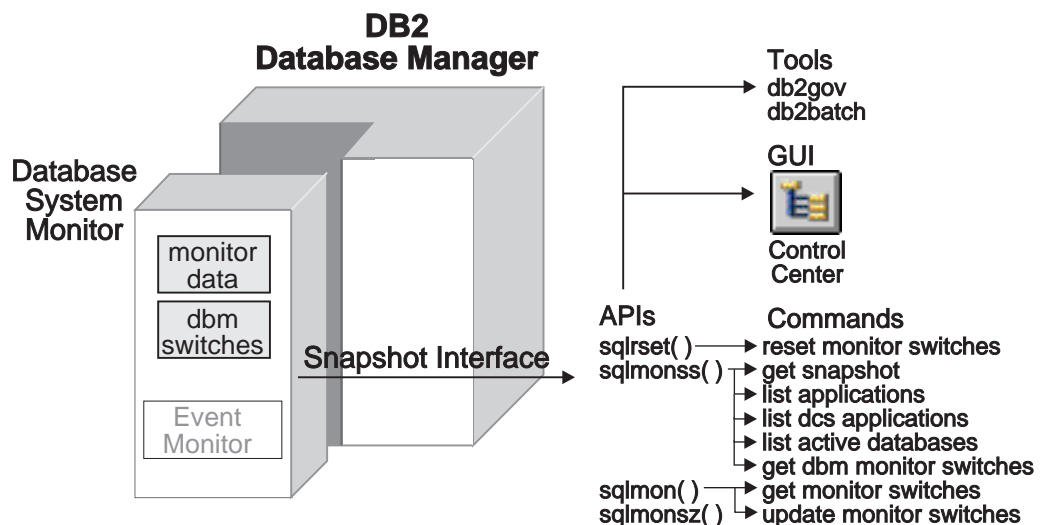


Figure 2. Snapshot Monitoring Interfaces

## Information Available by Taking Snapshots

The following table lists all the supported snapshot request types. For certain request types, some information is returned only if the associated monitor switch is set ON. See Chapter 3, "Database System Monitor Data Elements" on page 31 to determine if a required counter is under switch control.

In the table, the API Request type column identifies the value that is supplied as input to the SQLMA input structure in the sqlmonss() Snapshot API routine.

API request type	CLP command	Information returned
<b>List of connections</b>		
SQLMA_APPLINFO_ALL	list applications [show detail]	Application identification information for all applications currently connected to a database that is managed by the DB2 instance on the node where snapshot is taken.
SQLMA_DBASE_APPLINFO	list applications for database <i>dbname</i> [show detail]	Same as SQLMA_APPLINFO_ALL for each application currently connected to the specified database.
SQLMA_DCS_APPLINFO_ALL	list dcs applications	Application identification information for all DCS applications currently connected to a database that is managed by the DB2 instance on the node where snapshot is taken.
<b>Database manager snapshot</b>		
SQLMA_DB2	get snapshot for dbm	Database manager level information, including internal monitor switch settings.
	get dbm monitor switches	Returns internal monitor switch settings.
<b>Database snapshot</b>		
SQLMA_DBASE	get snapshot for database on <i>dbname</i>	Database level information and counters for a database. Information is returned only if there is at least one application connected to the database.
SQLMA_DBASE_ALL	get snapshot for all databases	Same as SQLMA_DBASE for each database active on the node.
	list active databases	The number of connections to each active database. Includes databases that were started using the ACTIVATE DATABASE command, but have no connections.
<b>Application snapshot</b>		
SQLMA_APPL	get snapshot for application applid <i>appl-id</i>	Application level information, includes cumulative counters, status information, and most recent SQL statement executed (if statement switch is set).
SQLMA_AGENT_ID	get snapshot for application agentid <i>appl-handle</i>	Same as SQLMA_APPL.
SQLMA_DBASE_APPLS	get snapshot for applications on <i>dbname</i>	Same as SQLMA_APPL, for each application that is connected to the database on the node.
SQLMA_APPL_ALL	get snapshot for all applications	Same as SQLMA_APPL, for each application that is active on the node.
<b>Table snapshot</b>		

API request type	CLP command	Information returned
SQLMA_DBASE_TABLES	get snapshot for tables on <i>dbname</i>	Table activity information at the database and application level for each application connected to the database, and at the table level for each table that <b>was accessed</b> by an application connected to the database. Requires the table switch.
<b>Lock snapshot</b>		
SQLMA_APPL_LOCKS	get snapshot for locks for application applid <i>appl-id</i>	List of locks held by the application. Also, lock wait information if any and the lock switch is ON.
SQLMA_APPL_LOCKS_AGENT_ID	get snapshot for locks for application agentid <i>appl-handle</i>	Same as SQLMA_APPL_LOCKS.
SQLMA_DBASE_LOCKS	get snapshot for locks on <i>dbname</i>	Lock information at the database level, and application level for each application connected to the database. Requires the lock switch.
<b>Table space snapshot</b>		
SQLMA_DBASE_TABLESPACES	get snapshot for tablespace on <i>dbname</i>	Information about table space activity at the database level, the application level for each application connected to the database, and the table space level for each table space that has been accessed by an application connected to the database. Requires the buffer pool switch.
<b>Buffer pool snapshot</b>		
SQLMA_BUFFERPOOLS_ALL	get snapshot for all bufferpools	Buffer pool activity counters. Requires the buffer pool switch.
SQLMA_DBASE_BUFFERPOOLS	get snapshot for bufferpools on <i>dbname</i>	Same as SQLMA_BUFFERPOOLS_ALL, but for specified database only.

## Snapshot Uses an Instance Connection

**You do not need to be connected to a database in order to use the snapshot APIs.** They are performed under an **instance connection**, which is a connection between an application and an instance of the DB2 database manager.

The instance attachment is usually done implicitly to the instance specified by the DB2INSTANCE environment variable when the first database system monitor API is invoked by the application. It can also be done explicitly, using the ATTACH TO NODE command.

Once an application is attached, all system monitor requests that it invokes are directed to that instance. This allows a client to monitor a remote server, by simply attaching to the instance on it.

## Availability of Snapshot Monitor Data

If all applications disconnect from a database, then the system monitor data for that database is no longer available. To obtain monitor information for all database activity during a given period you may want to use an event monitor. Alternatively, you can keep the database active until your final snapshot has been taken, either by starting it with the `ACTIVATE DATABASE` command, or by maintaining a permanent connection to the database.

---

## Event Monitors

In contrast to taking a point in time snapshot, an event monitor writes out database system monitor data to either a file or a named pipe, when one of the following events occurs:

- end of a transaction
- end of a statement
- a deadlock
- start of a connection
- end of a connection
- database activation
- database deactivation
- end of a statement's subsection (when a database is partitioned)

An event monitor effectively provides the ability to obtain a **trace** of the activity on a database.

For example, you can request that DB2 logs the occurrence of deadlocks between connections to a database. First, you must create and activate a `DEADLOCK` event monitor:

For UNIX systems

### Monitor Session

```
db2 connect to sample
db2 "create event monitor dlockmon for
deadlocks write to file '/tmp/dlocks'"
mkdir /tmp/dlocks
db2 "set event monitor dlockmon state 1"
```

For OS/2 and Windows systems

### Monitor Session

```
db2 connect to sample
db2 "create event monitor dlockmon for
deadlocks write to file 'c:\tmp\dlocks'"
mkdir c:\tmp\dlocks
db2 "set event monitor dlockmon state 1"
```

Now, two applications using the database enter a deadlock. That is, each one is holding a lock that the other one needs in order to continue processing. The deadlock is eventually detected and resolved by the DB2 deadlock detector component, which will rollback one of the transactions. The following figures illustrate this scenario.

#### Application 1

```
db2 connect to sample
db2 +c "insert into staff values (1, 'Ofer',
1, 'Mgr', 0, 0, 0)"
DB20000I The SQL command completed
successfully.
```

**Note:** The +c option turns autocommit off for CLP.

Application 1 is now holding an exclusive lock on a row of the staff table.

#### Application 2

```
db2 connect to sample
db2 +c "insert into department values ('1',
'System Monitor', '1', 'A00', NULL)"
DB20000I The SQL command completed
successfully.
```

Application 2 now has an exclusive lock on a row of the department table.

#### Application 1

```
db2 +c select deptname from department
```

Assuming cursor stability, Application 1 needs a share lock on each row of the department table as the rows are fetched, but a lock on the last row cannot be obtained because Application 2 has an exclusive lock on it. Application 1 enters a LOCK WAIT state, while it waits for the lock to be released.

#### Application 2

```
db2 +c select name from staff
```

Application 2 also enters a LOCK WAIT state, while waiting for Application 1 to release its exclusive lock on the last row of the staff table.

These applications are now in a deadlock. This waiting will never be resolved because each application is holding a resource that the other one needs to continue. Eventually, the deadlock detector checks for deadlocks (see the *dlchktme* database manager configuration parameter in the *Administration Guide*) and chooses a victim to rollback:

**Application 2**

```
SQLN0991N The current transaction has been
rolled back because of a deadlock or timeout.
Reason code "2". SQLSTATE=40001
```

At this point the event monitor logs a deadlock event to its target. Application 1 can now continue:

**Application 1**

```
DEPTNAME
-----
PLANNING
INFORMATION CENTER
. . .
SOFTWARE SUPPORT
SYSTEM MONITOR

9 record(s) selected
```

Because an event monitor buffers its output and this scenario did not generate enough event records to fill a buffer, the event monitor is turned off to force it to flush its buffers:

**Monitor Session**

```
db2 "set event monitor dlockmon state 0"
DB20000I The SQL command completed
successfully.
```

The event trace is written as a binary file. It that can now be formatted using the `db2evmon` tool:

**Monitor Session**

```
db2evmon -path /tmp/dlocks
Reading /tmp/dlocks/00000000.evt . . .
```

This will format and print to *stdout*, a trace similar to the following:

-----  
EVENT LOG HEADER

Event Monitor name: DLOCKMON  
Server Product ID: SQL05000  
Version of event monitor data: 5  
Byte order: BIG ENDIAN  
Number of nodes in db2 instance: 1  
Codepage of database: 850  
Country code of database: 1  
Server instance name: bourbon  
-----

-----  
Database Name: SAMPLE  
Database Path: /home/bourbon/bourbon/NODE0000/SQL00002/  
First connection timestamp: 06-03-1997 13:31:13.607548  
Event Monitor Start time: 06-03-1997 13:32:11.676071  
-----

3) Connection Header Event ...

Appl Handle: 0  
Appl Id: \*LOCAL.bourbon.970603173114 - Monitor session  
Appl Seq number: 0001  
DRDA AS Correlation Token: \*LOCAL.bourbon.970603173113  
Program Name : db2bp\_32  
Authorization Id: BOURBON  
Execution Id : bourbon  
Codepage Id: 850  
Country code: 1  
Client Process Id: 63590  
Client Database Alias: sample  
Client Product Id: SQL05000  
Client Platform: AIX  
Client Communication Protocol: Local  
Client Network Name:  
Connect timestamp: 06-03-1997 13:31:13.607548

4) Connection Header Event ...

Appl Handle: 1 - Application 1  
Appl Id: \*LOCAL.bourbon.970603173330  
Appl Seq number: 0001  
DRDA AS Correlation Token: \*LOCAL.bourbon.970603173329  
Program Name : db2bp\_32  
Authorization Id: BOURBON  
Execution Id : bourbon  
Codepage Id: 850  
Country code: 1  
Client Process Id: 119710  
Client Database Alias: sample  
Client Product Id: SQL05000  
Client Platform: AIX  
Client Communication Protocol: Local  
Client Network Name:  
Connect timestamp: 06-03-1997 13:33:29.518568

5) Connection Header Event ...

Appl Handle: 2  
Appl Id: \*LOCAL.bourbon.970603173409 - Application 2  
Appl Seq number: 0001  
DRDA AS Correlation Token: \*LOCAL.bourbon.970603173408  
Program Name : db2bp\_32  
Authorization Id: BOURBON  
Execution Id : bourbon  
Codepage Id: 850  
Country code: 1  
Client Process Id: 33984  
Client Database Alias: sample

```
Client Product Id: SQL05000
Client Platform: AIX
Client Communication Protocol: Local
Client Network Name:
Connect timestamp: 06-03-1997 13:34:08.972643
```

- 6) Deadlock Event ...  
Number of applications deadlocked: 2 - Deadlock  
Deadlock detection time: 06-03-1997 13:36:48.817786  
Rolled back Appl Id: : \*LOCAL.bourbon.970603173409  
Rolled back Appl seq number: : 0001
- 7) Deadlocked Connection ...  
Appl Id: \*LOCAL.bourbon.970603173409  
Appl Seq number: 0001  
Appl Id of connection holding the lock: \*LOCAL.bourbon.970603173330  
Seq. no. of connection holding the lock:  
Lock wait start time: 06-03-1997 13:36:43.251687  
Deadlock detection time: 06-03-1997 13:36:48.817786  
Table of lock waited on : STAFF  
Schema of lock waited on : BOURBON  
Tablespace of lock waited on : USERSPACE1  
Type of lock: Row  
Mode of lock: X  
Lock object name: 39
- 8) Deadlocked Connection ...  
Appl Id: \*LOCAL.bourbon.970603173330  
Appl Seq number: 0001  
Appl Id of connection holding the lock: \*LOCAL.bourbon.970603173409  
Seq. no. of connection holding the lock:  
Lock wait start time: 06-03-1997 13:35:32.227521  
Deadlock detection time: 06-03-1997 13:36:48.817786  
Table of lock waited on : DEPARTMENT  
Schema of lock waited on : BOURBON  
Tablespace of lock waited on : USERSPACE1  
Type of lock: Row  
Mode of lock: X  
Lock object name: 15

This event monitor trace shows that there was 1 application connected to the database when the event monitor was activated. This is indicated by the first *Connection Header Event* record in the output (record number 3). A Connection Event Header is generated for each active connection when an event monitor is turned on, and for each subsequent connection, once it becomes active. The other two Connection Headers, (records 4 and 5) were generated when the two applications connected.

The trace also shows that a deadlock occurred (record number 6). It shows which locks on which tables caused this deadlock (record numbers 7 and 8), and which application the deadlock detector chose to roll back (record number 6).

The **db2eva** graphical tool can also be used for formatting a trace. It is particularly useful for handling file traces that are too large to be read with db2evmon. It displays collected information in a tabular format. It includes a number of different view options, which allows you to filter unwanted records and drill down to the periods of interest in the trace. For instance, you can decide to display only the transaction events for a given connection. It also allows you to view the statement text for static SQL that it automatically fetches from the DB catalog (the text is only available for dynamic SQL in the event monitor trace).



You can invoke this tool with the db2eva command (see “db2eva - Event Analyzer Command” on page 222), or by selecting the Event Analyzer from the GUI.

**Note:** The files must be available on the machine where you invoked db2eva.

The db2eva tool is available on OS/2 and Windows systems.

## Authority Required for Event Monitoring

To define and use an event monitor on a database, you must have at least DBADM authority on that database.

## Using Event Monitors

As illustrated in the sample scenario, collecting system monitor data with an event monitor is a three step process:

1. Create the event monitor
2. Activate the event monitor
3. Read the trace produced.

### Create the event monitor.

Specify the events to be monitored. An event monitor is created and activated by using SQL statements. Unlike snapshot monitoring, where data can be collected at the database manager level, an event monitor only gathers data for a single database.

Creating an event monitor stores its definition in the event monitor database system catalogs:

SYSCAT.EVENTMONITORS	event monitors defined for the database
SYSCAT.EVENTS	events monitored for the database

It is necessary to connect to the database when defining an event monitor.

### Activate the event monitor.

Activating an event monitor starts a process or thread, which records monitor data to either a named pipe or a file as events occur. You may want an event monitor to be activated as soon as the database is started, so that all activity is monitored from start-up. This can be done by creating an AUTOSTART event monitor:

```
db2 "create event monitor DLOCKMON  
for deadlocks write to file '/tmp/dlocks'  
AUTOSTART"
```

This event monitor will be automatically started every time the database is activated, either by using the ACTIVATE DATABASE command, or when the first application connects. Note that creating an AUTOSTART event monitor does not activate it. This event

monitor will be activated the next time the database is stopped and re-activated. An event monitor that has not been automatically started must be manually started:

```
db2 set event monitor dlockmon state 1
```

All event monitors for a database are stopped when the database is deactivated.

**Read the trace produced.**

Reading a trace can be done using the **db2evmon** applet, or by writing your own application (see Chapter 4, “Event Monitor Output” on page 195). The Control Center and Event Analyzer (parts of the DB2 GUI) can be used to create and activate event monitors, and to read the traces produced by FILE event monitors.

Figure 3 illustrates the process and interface for using event monitors.

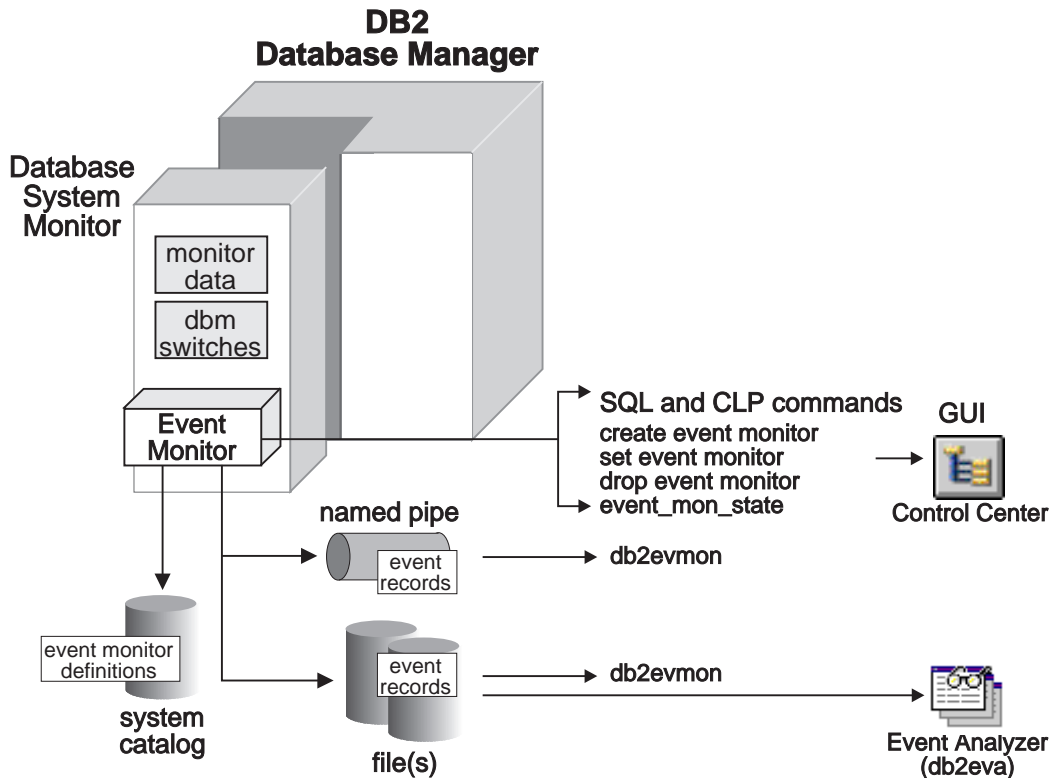


Figure 3. Event Monitoring Interfaces

As illustrated in Figure 3, event monitors are created and activated using the following SQL statements:

- CREATE EVENT MONITOR stores the event monitor definition in the database system catalogs for event monitors.
- SET EVENT MONITOR activates the event monitor, starting an **output thread** that will WRITE monitor data to either a file or named pipe. The trace produced can be formatted by the *db2evmon* or *db2eva* tools.
- DROP EVENT MONITOR deletes the event monitor definition from the database system catalogs for event monitors. An active event monitor cannot be dropped.

### Querying the State of an Event Monitor

You can determine if an event monitor is active by using the SQL function EVENT\_MON\_STATE:

```
db2 connect to sample
db2 "select evmonname, EVENT_MON_STATE(evmonname)
from syscat.eventmonitors"
```

NAME	2
-----	-----
DLOCKMON	0

1 record(s) selected

A returned value of 0 indicates that the event monitor is inactive.

### Information Available from Event Monitors

Event monitors return information that is similar to the information available using the snapshot API. In these cases, it is an event that controls when the snapshot is taken. For example, a connection event monitor basically takes an application snapshot just before the connection is terminated.

## Event Types

When you define an event monitor you must declare the event types that will be monitored. The following table lists the event types supported and indicates the information returned. **Note:** an event monitor can be defined for more than one event type.

Event type	When data is collected	Information returned
Deadlock	Detection of a deadlock	The applications involved and locks in contention.
Statements	End of SQL statement	Statement start/stop time, CPU used, text of dynamic SQL, SQLCA (return code of SQL statement), and other metrics such as fetch count.
	End of subsection	For partitioned databases: CPU consumed, execution time, table and tablequeue information.
Transactions	End of unit of work	Unit of work start/stop time, previous UOW time, CPU consumed, locking and logging metrics.
Connections	End of connection	All application level counters.
Database	Database deactivation or last connect reset	All database level counters.
Buffer pools		Counters for buffer pool, prefetchers, page cleaners and direct I/O for each buffer pool.
Table spaces		Counters for buffer pool, prefetchers, page cleaners and direct I/O for each table space.
Tables		Rows read/written for each table.

**Note:** In addition to the above information, all event monitors trace the establishment of connections to the database, by generating a *connection header record* for each active connection when the event monitor is turned ON, and for each subsequent connection, thereafter.

See “Output Stream Format” on page 195 for a list of the records generated for each event type.

## Using Pipe Event Monitors

A pipe event monitor allows you to process event records in real time. Another important advantage to using pipe event monitors is that your application can ignore unwanted data as it reads it off the pipe, giving the opportunity to considerably reduce storage requirements. It also allows an application to store event monitor data, in real-time, into an SQL database.

When you direct data to a pipe, I/O is always blocked and the only buffering is that performed by the pipe. It is the responsibility of the monitoring application to promptly read the data from the pipe as the event monitor writes the event data. If the event monitor is unable to write the data to the pipe (for example, because the pipe is full), monitor data will be lost.

The steps for using pipe event monitors are essentially the same on all operating systems. However, implementation can be different. The following section describes the basic steps, and highlights the differences between UNIX based systems, Windows NT, and OS/2.

1. Define the event monitor

```
db2 connect to sample
On AIX, and other UNIX platforms:
db2 create event monitor STMT2 for statements
write to PIPE '/tmp/evmpipe1'
On Windows NT:
db2 create event monitor STMT2 for statements
write to PIPE '\\.\pipe\evmpipe1'
On OS/2:
db2 create event monitor STMT2 for statements
write to PIPE '\pipe\evmpipe1'
```

2. Create the named pipe

In UNIX (this includes AIX environments), use the `mkfifo()` function or `mkfifo` command. In OS/2, use the `DosCreateNPipe()` function. In Windows NT, use the `CreateNamedPipe()` function. The pipe name must be the same as the target path specified on the CREATE EVENT MONITOR statement.

3. Open the named pipe

In UNIX, use the `open()` function. In OS/2, use the `DosConnectNPipe()` function. In Windows NT, use the `ConnectNamedPipe()` function.

You can also use the `db2evmon` application, specifying the database and pipe name, for example:

```
db2evmon -db sample -evm STMT2
```

This will open the named pipe and wait for the event monitor to write to it.

4. Activate the event monitor

If the event monitor is started automatically, you do not need to take any specific action to start it unless the database is already active (however, the pipe must already be opened).

```
db2 set event monitor stmt2 state 1
```

5. Read data from the named pipe

In UNIX, use the `read()` function. In OS/2, use the `DosRead()` function. In Windows NT, use the `ReadFile()` function. Your application may stop reading data from the pipe at any time. When it reads an EOF, there is no more monitor data. See Chapter 4, "Event Monitor Output" on page 195 for how to read the event monitor data.

6. Deactivate the event monitor

```
db2 set event monitor stmt2 state 0
```

This statement can be used to stop any event monitor, even one that was started automatically. If you do not explicitly stop an event monitor, it will be stopped when:

- The last application disconnects from the database
- It experiences an error while writing to the named pipe: for example, the monitoring application closes the pipe before deactivating the event monitor. In this case, the event monitor will turn itself off and log a system-error-level message in the diagnostic log, db2diag.log.

7. Close the named pipe.

In UNIX, use the `close()` function. In OS/2, use the `DosDisConnectNPIPE()` function. In Windows NT, use the `DisconnectNamedPipe()` function.

8. Delete the named pipe.

In UNIX, use the `unlink()` function. In OS/2, use the `DosClose()` function. In Windows NT, use the `CloseHandle()` function.

For UNIX-based operating systems, named pipes are like files, so you do not have to delete them and create them again before each use.

### Pipe Overflows

In addition, there must be enough space in the named pipe. If the application does not read the data fast enough from the named pipe, the pipe will fill up and overflow. Pipe overflows can also occur on platforms (such as OS/2) where the creator of the pipe can define the size of the named pipe buffer. The smaller the buffer, the greater the chance of an overflow occurring. When a pipe overflow occurs, the monitor creates overflow event records indicating that an overflow has occurred. The event monitor is not turned off, but monitor data is lost. If there are outstanding overflow event records when the monitor is deactivated, a diagnostic message will be logged. Otherwise, the overflow event records will be written to the pipe when possible.

If your operating system allows you to define the size of the pipe buffer, use a pipe buffer of at least 32K. For high-volume event monitors, you should set the monitoring application's process priority equal to or higher (lower nice value on AIX) than the agent process priority (see the section on Priority of Agents in the *Administration Guide*).

---

## When Counters are Initialized

The data collected by the database manager includes several accumulating counters. These counters are incremented during the operation of the database, for example, every time an application commits a transaction.

Counters are initialized when their applicable object becomes active. For example, the number of buffer pool pages read for a database (a basic element) is set to zero when the database is activated.

Counters under switch control are reset to zero when their associated switch is turned on.

Counters returned by event monitors are reset to zero when the event monitor is activated.

---

## Resetting Monitor Data

Each event monitor and any application using the snapshot monitor APIs has its own logical view of the DB2 monitor data and switches. This means that when counters are reset or initialized, it only affects the event monitor or application that reset or initialized them.

Event monitor data cannot be reset, except by turning the monitor off, and then on again.

An application taking snapshots can reset its view of the counters at any time by using the RESET MONITOR command.

When issuing its first snapshot API, an application inherits the default settings from the database manager configuration. For example, assuming that the statement switch was set in the database manager configuration file:

```
db2 update dbm cfg using DFT_MON_STMT on
db2start
```

Issuing a GET MONITOR SWITCHES command will show that the statement switch is ON.

```
db2 get monitor switches
```

Monitor Recording Switches		
Buffer Pool Activity Information (BUFFERPOOL)	= OFF	
Lock Information (LOCK)	= OFF	
Sorting Information (SORT)	= OFF	
SQL Statement Information (STATEMENT)	= ON	05-25-1997 10:44:34.820446
Table Activity Information (TABLE)	= OFF	
Unit of Work Information (UOW)	= OFF	

Turning OFF the statement switch from the command line will only affect the application issuing the command. The statement switch will still be ON for other applications (unless they have also turned it OFF). For example:

```
db2 update monitor switches using STATEMENT OFF
DB20000I The UPDATE MONITOR SWITCHES command completed successfully
```

Then query your application's switches.

```
db2 get monitor switches
```

#### Monitor Recording Switches

```
Buffer Pool Activity Information (BUFFERPOOL) = OFF
Lock Information (LOCK) = OFF
Sorting Information (SORT) = OFF
SQL Statement Information (STATEMENT) = OFF
Table Activity Information (TABLE) = OFF
Unit of Work Information (UOW) = OFF
```

Querying the database manager switches will show that the update did not affect its settings:

```
db2 get database manager monitor switches
```

#### DBM System Monitor Information Collected

```
Buffer Pool Activity Information (BUFFERPOOL) = OFF
Lock Information (LOCK) = OFF
Sorting Information (SORT) = OFF
SQL Statement Information (STATEMENT) = ON    05-25-1997 10:44:34
Table Activity Information (TABLE) = OFF
Unit of Work Information (UOW) = OFF
```

When a monitoring application turns off a monitor switch or resets a data element counter, the DB2 server does not reset its own internal counters. Instead, it re-initialize the private **logical view** for that user. Other monitoring applications or event monitors are not affected.

You must have SYSADM, SYSCTRL, or SYSMANT authority to use the UPDATE MONITOR SWITCHES command. See “UPDATE MONITOR SWITCHES Command” on page 266 for information on this command.

The database manager keeps track of all the applications using the snapshot monitor APIs and their switch settings. If a switch is set in its configuration, then the database manager always collects that monitor data. If a switch is OFF in the configuration, then the database manager will collect data as long as there is at least one application with this switch turned ON.

Chapter 3, “Database System Monitor Data Elements” on page 31 shows the data elements associated with each switch group.

Internally, event monitors also use switches to instruct the engine as to which data should be collected. However, this is an implementation issue, and the switch settings for a particular event monitor cannot be queried.

An actual DBMS monitor switch is set as long as at least one application or event monitor needs it, or if it is set in the configuration file.



---

## System Monitor Memory Requirements - (mon\_heap\_sz)

The memory required for maintaining the private views of the database system monitor system monitor data is allocated from the monitor heap. Its size is controlled by the *mon\_heap\_sz* configuration parameter. The amount of memory required for monitoring activity varies widely depending on the number of monitoring applications and event monitors, the switches set, and the level of database activity. The following formula provides an approximation of the number of pages required for the monitor heap.

$$\begin{aligned} & (\text{number of monitoring applications} + 1) * \\ & (\text{number of databases} * \\ & \quad (800 + (\text{number of tables accessed} * 20) + \\ & \quad \quad ((\text{number of applications connected} + 1) * \\ & \quad \quad \quad (600 + (\text{number of table spaces} * 100)))))) \\ & / 4096 \end{aligned}$$

You may need to experiment with this value, increasing it if monitor commands occasionally fail with an SQLCODE of -973, when the database manager switches are on.

---

## Partitioned Database Considerations

The database system monitor interface is the same for all types of systems, whether they use single partition or multiple partition databases and whether intra-query parallelism is used. All the commands and APIs are exactly the same. The only difference is the output; more complex systems generally return more information.

## Taking a Snapshot on Multi-node Systems

On systems that use inter-partition parallelism, taking a snapshot only returns monitor data from the instance where the application is attached. For example, assuming a table that is located in two database partitions, that is some of its rows are stored on one node (Node 100) and others are stored on another node (Node 200).

### Node 100

```
db2 connect to sample
db2 list applications
```

Auth Id	Appl Name	Appl Handle	Application Id	DB Name	# of Agents
BOURBON	db2bp_32	6553638	*LOCAL.bourbon.970414221746	SAMPLE	1

Taking a snapshot on Node 200 initially returns no data:

**Note:** The LIST APPLICATION command uses the database system monitor. Invoking it actually calls the the snapshot API `sqlmonss()` with a request of type `SQLMA_APPLINFO_ALL`.

Node 200

```
db2 list applications
SQL1611W No data was returned from Database System Monitor.
```

Now, issuing a query from Node 100 will result in a secondary connection to Node 200 to fetch the rows that reside in that partition:

Node 100

```
db2 +c select lastname from employee

Huras
Ofer
Bourbonnais
Musker
Cartwright
```

Now there is a subagent for the application running on Node 200:

Node 200

```
db2 list applications

Auth Id      Appl      Appl      Application Id      DB      # of
Name        Name      Handle    *LOCAL.bourbon.970414221746  Name    Agents
-----
BOURBON db2bp_32  6553638  *LOCAL.bourbon.970414221746  SAMPLE  1
```

And there are now two agents running on Node 100; the coordinator agent and a subagent:

Node 100

```
db2 list applications

Auth Id      Appl      Appl      Application Id      DB      # of
Name        Name      Handle    *LOCAL.bourbon.970414221746  Name    Agents
-----
BOURBON db2bp_32  6553638  *LOCAL.bourbon.970414221746  SAMPLE  2
```

On the non-coordinating node, you can determine where the coordinator resides, and check if the application originated on the node that issued the snapshot, using:

```
db2 list application show detail
```

Appl Handle	Application Id	Coordinating Node Number	Coordinator pid/thread
6553638	*LOCAL.bourbon.970414221746	100	66204

The *Application Handle* returned, 6553638 is unique across all nodes. The *node number* corresponds to one of the nodes listed in the *db2nodes.cfg* configuration file (see the *Administration Guide*).

Using the application handle, you can request monitor information on any node by issuing a GET SNAPSHOT FOR APPLICATION, which will return data if the application is connected on that node. You can also FORCE the application, which will work from any node:

```
db2 force application (6553638)
DB20000I The FORCE APPLICATION command completed successfully.
DB221024I This command is asynchronous and may not be effective
immediately.
```

## Using Event Monitors on Multi-node Systems

An event monitor uses an operating system process or a thread to write its trace. The node where this process or thread runs is called the **monitor node**. An event monitor can be monitoring events as they occur locally on the monitor node, or globally as they occur on any node where the DB2 database manager is running. A global event monitor writes a single trace that contains activity from all nodes.

Whether an event monitor is local or global is referred to as its **monitoring scope**. Both the monitor node and monitor scope are part of an event monitor's definition. For example:

```
db2 connect to sample
db2 "create event monitor DLOCKS for
deadlocks write to file '/tmp/dlocks'
ON NODE 5 GLOBAL"
```

This global event monitor will report deadlocks that involve any nodes in the system. Its I/O component will physically run on Node 5, writing its records to files in the /tmp/dlocks directory on that node.

You can look at the definition for this monitor in the system catalog:

```
db2 "select evmonname,nodenum, monscope
from syscat.eventmonitors"
-----
EVMONNAME      NODENUM      MONSCOPE
-----
DLOCKS          5             G

1 record(s) selected
```

The returned information shows event monitor DLOCKS is defined as global and its monitor node is 5.

**Note:** In DB2 Version 5, only deadlock event monitors can be defined as global, all other event monitors must be defined as local.

## Monitoring Subsections

On systems that use inter-partition parallelism, the SQL compiler partitions the access plan for an SQL statement into **subsections**. Each subsection is executed by a different DB2 agent.

The access plan for an SQL statement generated by the DB2 code generator during compilation can be obtained using the **db2expln** or **dynexpln** commands (see the *Command Reference*). As an example, selecting all the rows from a table that is partitioned across several nodes might result in an access plan having two subsections:

1. Subsection 0, the coordinator subsection, whose role is to collect rows fetched by the other DB2 agents (subagents) and return them to the application.
2. Subsection 1, whose role is to perform a table scan and return the rows to the coordinating agent.

In this simple example, subsection 1 would be distributed across all the database partitions. There would be a subagent executing this subsection on each physical node of the **nodegroup** to which this table belongs. See *Administration Guide* for more information on these concepts.

The database system monitor allows you to correlate run-time information with the access plan, which is compile-time information. With inter-partition parallelism, it breaks information down to the subsection level. For example, when the statement monitor switch is ON, a GET SNAPSHOT FOR APPLICATION will return information for each subsection executing on this node, as well as totals for the statement.

The subsection information returned for an application snapshot includes:

- the number of table rows read/written
- CPU consumption
- elapsed time

- the number of tablequeue rows sent and received from other agents working on this statement. This allows you to track the execution of a long running query by taking a series of snapshots.
- subsection status. If the subsection is in a WAIT state, because it is waiting for another agent to send or receive data, then the information also identifies the node or nodes preventing the subsection from progressing in its execution. You may then take a snapshot on these nodes to investigate the situation.

The information logged by a statement event monitor for each subsection after it has finished executing includes: CPU consumption, total execution, time, and several other counters.

---

## DB2 Productivity Tools

The database system monitor is a very powerful function of the DB2 database manager. It can be exploited to develop productivity tools for the database administrator (DBA) and database developer. The following are a few examples of productivity tools that use the function of the database system monitor, and are included with the DB2 product:

- db2batch

An application that uses snapshot monitoring to collect metrics for tuning SQL queries. It can be found in `sqllib/misc/db2batch`. See the *Administration Guide* for more information.

- db2gov

The DB2 governor is an application that uses snapshot monitoring to supervise the load and usage of the database manager. It provides the functions to FORCE or change the run-time priority of applications exceeding certain limits. These limits are specified by the DBA in the db2gov configuration file. Application limits and privileges can be expressed using several different parameters, for example maximum amount of CPU. It can be found in `sqllib/adm/db2gov`. See *Administration Guide* for more information.

- db2evmon

An application that formats the data stream created by an event monitor. It can be found in `sqllib/misc/db2evmon`.

- Control Center

A GUI for snapshot and event monitoring. For snapshots, it allows you to define performance variables in terms of the metrics returned by the database system monitor and graph them over time. For example, you can request that it take a snapshot and graph the progression of a performance variable over the last eight hours. **Alerts** can be set to notify the DBA when certain threshold are reached. For event monitors, it allows you to create, activate, start, stop, and delete event monitors. See the online help for the Control Center for more information.

- Event Analyzer

A GUI for viewing file event monitor traces. Information collected on connections, deadlocks, overflows, transactions, statements, and subsections is organized and displayed in a tabular format. See the online help for the Event Analyzer for more information.

---

## System Monitor Definitions

- Data Element

A piece of information collected by the database system monitor. Snapshot and Event Monitors are the two different interfaces for accessing data elements. Some elements are only accessible with the event monitor (for example, connections involved in a deadlock), and others are only available with snapshot monitoring (for example, data maintained at the database manager level - total number of data-base agents).

- Basic Data Element

Performance data that is always collected or maintained by the DB2 server, even when all monitor switches are off.

- Database System Monitor

The function of the DB2 database manager that maintains information about its operation and performance.

- Event Analyzer

A function of the DB2 administration GUI for viewing event monitor traces.

- Event Monitoring

Using an event monitor to monitor database activity and performance by obtaining traces of the data collected when specific events occur (for example, the end of a transaction). It is requested from the DB2 server using an SQL interface.

- Monitor Switch

Instructs the DB2 database manager to collect those data elements that involve a non-negligible processing or memory overhead for a group of monitor data.

- Monitored Level or Object

An object about which the database system monitor can return performance, informational, or status data: for example, the database manager, a database, a connection, or a transaction.

- Monitoring Application

An application using the snapshot API or an event monitor. Each monitoring application has its own logical view of the data collected by the DB2 server. This means that if a monitoring application resets its counters, it does not affect the counters collected by other monitoring applications. The reset API resets counters for snapshot monitoring. Turning off/on an event monitor resets its counters.

Counters and switches can only be reset for an event monitor by deactivating and reactivating it.

- Snapshot Monitoring

Using the system monitor API to monitor database activity and performance by sampling the monitor data maintained by the DB2 server. This API consists of:

<code>sqlmon()</code>	set or query monitor switch settings
<code>sqlmonrset()</code>	reset system monitor counters
<code>sqlmonss()</code>	take a snapshot
<code>sqlmonsz()</code>	estimate the size of the data that would be returned for a particular invocation of <code>sqlmonss()</code>

- Snapshot Monitor

A function of the Control Center GUI for snapshot monitoring. It provides the ability to follow trends, define performance variables in terms of data elements, and set alerts when thresholds for these variables are reached.





---

## Chapter 3. Database System Monitor Data Elements

This chapter describes the information that is available from the database system monitor. The information returned by database system monitor falls into the following categories:

- identification for the database manager, an application, or a database connection being monitored.
- data primarily intended to help you to configure the system.
- database activity at various levels including application, table, or statement. This information can be used for activity monitoring, problem determination, and performance analysis. But it can also be used for configuration.

In this chapter, data elements are organized by their primary use category. When applicable, elements that have multiple uses may be referred to by associated elements in other categories. Multi-use data element information only appears in its main category, it is not duplicated in other categories. Refer to *data elements* in the Index, if you have trouble finding a data element.

The information is grouped as follows:

- “Server Identification and Status” on page 33
- “Database Identification and Status” on page 39
- “Application Identification and Status” on page 44
  - “DB2 Agent Information” on page 66
- “Database Manager Configuration” on page 66
  - “Agents and Connections” on page 67
  - “Sort” on page 78
  - “Fast Communication Manager” on page 84
- “Database Configuration” on page 90
  - “Buffer Pool Activity” on page 90
    - “Extended Storage” on page 110
  - “Non-buffered I/O Activity” on page 114
  - “Catalog Cache” on page 119
  - “Package Cache” on page 122
  - “Database Heap” on page 126
  - “Logging” on page 127
- “Database and Application Activity” on page 130
  - “Locks and Deadlocks” on page 130
  - “Lock Wait Information” on page 140
  - “Rollforward Monitoring” on page 147
  - “Table Activity” on page 149
  - “SQL Cursors” on page 160
  - “SQL Statement Activity” on page 164
  - “SQL Statement Details” on page 174
  - “Subsection Details” on page 183
  - “Intra-query Parallelism” on page 189
  - “CPU Usage” on page 191
  - “Snapshot Monitoring Elements” on page 192

**Note:** For Extended Enterprise Edition users, snapshot elements only apply to the partition where the snapshot was issued.

---

## How to Read the Data Element Tables

The section for each data element begins with a table that lists standard information. An example is shown in Figure 4, followed by an explanation of each part of the table.

---

<b>① Snapshot Information Level</b> Database Application	<b>② API Structure(s)</b> sqlm_dbase sqlm_appl	<b>③ Monitor Switch</b> Sort Sort
<b>④ Resettable</b>	Yes	
<b>⑤ Event Type</b> Database Connection Statement	<b>⑥ Event Record(s)</b> sqlm_db_event sqlm_conn_event sqlm_stmt_event	
<b>⑦ API Element Name</b> <b>Element Type</b>	total_sorts counter	
<b>⑧ Related Information</b>	See Resettable See Switches See Sort Overflows	

---

Figure 4. Sample Element Table

1. The level of information that can be captured by the snapshot monitor.
2. The API structures where captured snapshot information is returned.
3. The snapshot monitor switch that must be set to obtain this information.
4. Whether or not the counter can be reset (snapshot monitor only).
5. The event monitor must be created with this event type to collect this information.
6. The data structure where captured event information is returned.
7. The name and type of element, as returned in the API structure or event record.
8. References to related data elements.

This table is followed by a description of the element and information on how you can use it when monitoring your database.

## Element Types

Data elements are classified by the following categories:

- Counter  
A counter counts the number of times an activity occurs. Counter values increase during monitoring. Most are resettable.
- Gauge  
A gauge indicates the current value for an item. This value can go up and down depending on database activity (for example, the number of locks held).
- Water mark  
A water mark indicates the highest (maximum) or lowest (minimum) value an element reached since monitoring was started. These are not resettable.
- Information  
An information element provides reference-type details of your monitoring activities. This can include items such as node names, aliases, and path details.
- Timestamp  
A timestamp indicates the date and time that an activity took place, by providing the number of seconds and microseconds that have elapsed since January 1, 1970. In the C language, for example, this can be converted to calendar date and time using the `ctime()` function.
- Time  
Time returns the number of seconds and microseconds spent on an activity.

---

## Server Identification and Status

The following elements provide identification and status information about the server:

- “Start Database Manager Timestamp” on page 34
- “Configuration NNAME at Monitoring (Server) Node” on page 34
- “Server Instance Name” on page 35
- “Database Manager Type at Monitored (Server) Node” on page 35
- “Server Product/Version ID” on page 36
- “Server Version” on page 36
- “Service Level” on page 37
- “Server Operating System” on page 37
- “Product Name” on page 38
- “Product Identification” on page 38
- “Status of DB2 Instance” on page 39

## Start Database Manager Timestamp

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_db2	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	db2start_time timestamp	
<b>Related Information</b>	• “Snapshot Time” on page 193	

**Description:** The date and time that the database manager was started using the db2start command.

**Usage:** This element may be used with the *Snapshot Time* monitor element to calculate the elapsed time since the database manager was started up until the snapshot was taken.

## Configuration NNAME at Monitoring (Server) Node

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_collected	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	server_nname information	
<b>Related Information</b>	• “Configuration NNAME of Client” on page 53	

**Description:** The name of the node being monitored by the database system monitor.

**Usage:** This element can be used to identify the database server node you are monitoring. This information can be useful if you are saving your monitor output in a file or database for later analysis and you need to differentiate the data from different database server nodes. This node name is determined based on the *nname* configuration parameter.

If you are using the database system monitor APIs, note that the API constant SQLM\_IDENT\_SZ is used to define the length of this element. Only the first 8 characters are currently used.

## Server Instance Name

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_collected	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>Event Type</b> Event Log Header	<b>Event Record(s)</b> sqlm_event_log_header	
<b>API Element Name</b> <b>Element Type</b>	server_instance_name information	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Configuration NNAME at Monitoring (Server) Node” on page 34</li> </ul>	

**Description:** The name of the database manager instance for which the snapshot was taken.

**Usage:** If more than one instance of the database manager is present on the same system, this data item is used to uniquely identify the instance for which the snapshot call was issued. Along with *Configuration NNAME at Monitoring (Server) Node*, this information can be useful if you are saving your monitor output in a file or database for later analysis, and you need to differentiate the data from different instances of the database manager.

If you are using the database system monitor APIs, note that the API constant SQLM\_IDENT\_SZ is used to define the length of this element. Only the first 8 characters are currently used.

## Database Manager Type at Monitored (Server) Node

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_collected	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	server_db2_type information	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Configuration NNAME at Monitoring (Server) Node” on page 34</li> </ul>	

**Description:** Identifies the type of database manager being monitored.

**Usage:** It contains one of the following types of configurations for the database manager:

<b>API Symbolic Constant</b>	<b>Command Line Processor Output</b>
sqlf_nt_server	Database Server with local and remote clients
sqlf_nt_stand_req	Database Server with local clients

The API symbolic constants are defined in the include file *sqlutil.h*.

## Server Product/Version ID

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_collected	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>Event Type</b> Database Manager	<b>Event Record(s)</b> sqlm_event_log_header	
<b>API Element Name</b> <b>Element Type</b>	server_prdid information	
<b>Related Information</b>	• “Client Product/Version ID” on page 53	

**Description:** The product and version that is running on the server.

**Usage:** It is in the form PPPVVRRM, where:

PPP is SQL  
VV identifies a 2-digit version number (with high-order 0 in the case of a 1-digit version)  
RR identifies a 2-digit release number (with high-order 0 in the case of a 1-digit release)  
M identifies a 1-digit modification level

If you are using the database system monitor APIs, note that the API constant SQLM\_IDENT\_SZ is used to define the length of this element. Only the first 8 characters are currently used.

## Server Version

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_collected	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	server_version information	
<b>Related Information</b>	• “Server Product/Version ID” on page 36	

**Description:** The version of the server returning the information.

**Usage:** This field identifies the level of the database server collecting database system monitor information. This allows applications to interpret the data based on the level of the server returning the data. Valid values are:

SQLM\_VERSION1 Data was returned by DB2 Version 1  
SQLM\_VERSION2 Data was returned by DB2 Version 2  
SQLM\_VERSION5 Data was returned by DB2 Universal Database Version 5

## Service Level

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_db2	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	service_level information	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Product Identification” on page 38</li></ul>	

**Description:** This is the current corrective service level of the server.

**Usage:** Used to provide information when requesting service or reporting a problem with DB2 on OS/2. This element will be blank for non-OS/2 systems.

**Note:** This element is similar to the *corr\_serv\_lvl* field in the *sqlestat* output. See Appendix C, “DB2 Version 1 sqlestat Users” on page 271 for more information on sqlestat equivalent data elements.

## Server Operating System

<b>Snapshot Information Level</b> Database	<b>API Structure(s)</b> sqlm_dbase	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>Event Type</b> Database	<b>Event Record(s)</b> sqlm_db_event	
<b>API Element Name</b> <b>Element Type</b>	server_platform information	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Client Operating Platform” on page 58</li><li>• “Database Location” on page 43</li></ul>	

**Description:** The operating system running the database server.

**Usage:** This element can be used for problem determination for remote applications. Values for this field can be found in the header file *sqlmon.h*.

**Note:** This element is similar to the *db\_type* field in the *sqlestat* output. See Appendix C, “DB2 Version 1 sqlestat Users” on page 271 for more information on sqlestat equivalent data elements.

## Product Name

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_db2	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	product_name information	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Product Identification” on page 38</li><li>• “Service Level” on page 37</li></ul>	

**Description:** Details of the version of the server that is running.

**Usage:** Used to provide information when requesting service or reporting a problem with DB2 on OS/2. This element will be blank for non-OS/2 systems.

**Note:** This element is similar to the *product\_name* field in the *sqlstat* output. See Appendix C, “DB2 Version 1 *sqlstat* Users” on page 271 for more information on *sqlstat* equivalent data elements.

## Product Identification

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_db2	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	component_id information	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Product Name” on page 38</li><li>• “Service Level” on page 37</li></ul>	

**Description:** Details of the type of the server that is running.

**Usage:** Used to provide information when requesting service or reporting a problem with DB2 on OS/2. This element will be blank for non-OS/2 systems.

**Note:** This element is similar to the *component\_id* field in the *sqlstat* output. See Appendix C, “DB2 Version 1 *sqlstat* Users” on page 271 for more information on *sqlstat* equivalent data elements.



## Status of DB2 Instance

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_db2	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	db2_status information	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>“Status of Database” on page 42</li> </ul>	

**Description:** The current status of the instance of the database manager.

**Usage:** You can use this element to determine the state of your database manager instance.

The value returned is always SQLM\_DB2\_ACTIVE.

## Database Identification and Status

The following elements provide identification and status information about the database:

- “Database Name”
- “Database Path” on page 40
- “Database Activation Timestamp” on page 41
- “Database Deactivation Timestamp” on page 41
- “Status of Database” on page 42
- “Catalog Node Network Name” on page 42
- “Database Location” on page 43
- “Catalog Node Number” on page 43
- “Last Backup Timestamp” on page 44

## Database Name

<b>Snapshot Information Level</b> Database Application Table Space	<b>API Structure(s)</b> sqlm_dbase sqlm_appl_id_info sqlm_tablespace_header sqlm_bufferpool	<b>Monitor Switch</b> Basic Basic Buffer Pool Buffer Pool
Table Lock DCS Application	sqlm_table_header sqlm_dbase_lock sqlm_dcs_applinfo	Table Basic Basic
<b>Resettable</b>	No	
<b>Event Type</b> Database	<b>Event Record(s)</b> sqlm_dbheader_event	
<b>API Element Name</b> <b>Element Type</b>	db_name information	

---

**Related Information**

- “Resetting Monitor Data” on page 21
  - “Last Reset Timestamp” on page 192
  - “Input Database Alias” on page 193
  - “Database Alias Used by Application” on page 54
  - “Database Path” on page 40
- 

**Description:** The real name of the database for which information is collected or to which the application is connected. This is the name the database was given when created.

**Usage:** You may use this element to identify the specific database to which the data applies.

For applications that are not using DDCS to connect to a DRDA host database, you can use this element in conjunction with the Database Path monitor element to uniquely identify the database and help relate the different levels of information provided by the monitor.

If you are using the database system monitor APIs, note that the API constant `SQLM_IDENT_SZ` is used to define the length of this element. Only the first 8 characters are currently used for DB2 databases, and the first 18 characters are used for DRDA host databases.

## Database Path

---

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
Application	sqlm_appl_id_info	Basic
Table Space	sqlm_tablespace_header	Buffer Pool
	sqlm_bufferpool	Buffer Pool
Table	sqlm_table_header	Table
Lock	sqlm_dbase_lock	Basic

---

<b>Resettable</b>	No
-------------------	----

---

<b>Event Type</b>	<b>Event Record(s)</b>
Database	sqlm_dbheader_event

---

<b>API Element Name</b>	db_path
<b>Element Type</b>	information

---

<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “Input Database Alias” on page 193</li><li>• “Database Name” on page 39</li></ul>
----------------------------	---

---

**Description:** The full path of the location where the database is stored on the monitored system.

**Usage:** This element can be used with the *Database Name* monitor element to identify the specific database to which the data applies.

## Database Activation Timestamp

<b>Snapshot Information Level</b>	<b>API Structure(s)</b>	<b>Monitor Switch</b>
Database	sqlm_dbase	Basic
Table Space	sqlm_tablespace_header	Buffer Pool
Table	sqlm_table_header	Basic
<b>Resettable</b>	No	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Connection	sqlm_dbheader_event sqlm_connheader_event	
<b>API Element Name</b>	db_conn_time	
<b>Element Type</b>	conn_time timestamp	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “Connection Request Start Timestamp” on page 62</li> <li>• “Snapshot Time” on page 193</li> </ul>	

**Description:** The date and time of the connection to the database (at the database level, this is the first connection to the database), or when the activate database was issued.

**Usage:** Use this element with the Database Deactivation Timestamp monitor element to calculate the total connection time.

## Database Deactivation Timestamp

<b>Event Type</b>	<b>Event Record(s)</b>
Database	sqlm_db_event
Connection	sqlm_conn_event
<b>API Element Name</b>	disconn_time
<b>Element Type</b>	timestamp
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• None</li> </ul>

**Description:** The date and time that the application disconnected from the database (at the database level, this is the time the last application disconnected).

**Usage:** Use this element to calculate the elapsed time since:

- The database was active (for information at the database level)
- The connection was active (for information at the connection level).

## Status of Database

<b>Snapshot Information Level</b> Database	<b>API Structure(s)</b> sqlm_dbase	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	db_status information	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>“Status of DB2 Instance” on page 39</li> </ul>	

**Description:** The current status of the database.

**Usage:** You can use this element to determine the state of your database.

Values for this field are:

API Constant	Description
SQLM_DB_ACTIVE	The database is active.
SQLM_DB QUIESCE_PEND	The database is in quiesce-pending state. New connections to the database are <b>not</b> permitted and new units of work <b>cannot</b> be started. Depending on the quiesce request, active units of work will be allowed to complete or will be rolled back immediately.
SQLM_DB QUIESCED	The database has been quiesced. New connections to the database are <b>not</b> permitted and new units of work <b>cannot</b> be started.
SQLM_DB_ROLLFWD	A rollforward is in progress on the database.

## Catalog Node Network Name

<b>Snapshot Information Level</b> Database	<b>API Structure(s)</b> sqlm_dbase	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>Event Type</b> Database	<b>Event Record(s)</b> sqlm_db_event	
<b>API Element Name</b> <b>Element Type</b>	catalog_node_name information	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>None</li> </ul>	

**Description:** The network name of the catalog node. On OS/2, the netbios name of the server where the database is located.

**Usage:** Use this element to determine the location of a database.

**Note:** This element is similar to the *node* field in the *sqlestat* output. See Appendix C, “DB2 Version 1 sqlestat Users” on page 271 for more information on sqlestat equivalent data elements.

## Database Location

<b>Snapshot Information Level</b> Database	<b>API Structure(s)</b> sqlm_dbase	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	db_location information	
<b>Related Information</b>	• “Server Operating System” on page 37	

**Description:** The location of the database in relation to the application.

**Usage:** Determine the relative location of the database server with respect to the application taking the snapshot. Values are:

- SQLM\_LOCAL
- SQLM\_REMOTE

**Note:** This element is similar to the *location* field in the *sqlstat* output. See Appendix C, “DB2 Version 1 sqlstat Users” on page 271 for more information on sqlstat equivalent data elements.

## Catalog Node Number

<b>Snapshot Information Level</b> Database	<b>API Structure(s)</b> sqlm_dbase	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>Event Type</b> Database	<b>Event Record(s)</b> sqlm_db_event	
<b>API Element Name</b> <b>Element Type</b>	catalog_node information	
<b>Related Information</b>	• None	

**Description:** The node number of the node where the database catalog tables are stored.

**Usage:** The catalog node is the node where all system catalog tables are stored. All access to system catalog tables must go through this node. See the *Administration Guide* for information on system catalog tables.

## Last Backup Timestamp

<b>Snapshot Information Level</b> Database	<b>API Structure(s)</b> sqlm_dbase	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	last_backup timestamp	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• None</li></ul>	

**Description:** The date and time that the latest database backup was completed.

**Usage:** You may use this element to help you identify a database that has not been backed up recently, or to identify which database backup file is the most recent. If the database has never been backed up, this timestamp is initialized to zero.

---

## Application Identification and Status

The following elements provide information about databases and their related applications.

- “Application Handle (agent ID)” on page 45
- “Application Status” on page 46
- “ID of Code Page Used by Application” on page 48
- “Application Status Change Time” on page 48
- “Application Name” on page 49
- “Application ID” on page 50
- “Sequence Number” on page 52
- “Authorization ID” on page 52
- “Configuration NNAME of Client” on page 53
- “Client Product/Version ID” on page 53
- “Database Alias Used by Application” on page 54
- “Host Product/Version ID” on page 55
- “Outbound Application ID” on page 55
- “Outbound Sequence Number” on page 56
- “User Login ID” on page 56
- “DRDA Correlation Token” on page 57
- “Client Process ID” on page 57
- “Client Operating Platform” on page 58
- “Client Communication Protocol” on page 58
- “Database Country Code” on page 59
- “Application Agent Priority” on page 59
- “Application Priority Type” on page 60
- “User Authorization Level” on page 60
- “Coordinating Node” on page 61
- “Connection Request Start Timestamp” on page 62
- “Connection Request Completion Timestamp” on page 62
- “Previous Unit of Work Completion Timestamp” on page 62
- “Unit of Work Start Timestamp” on page 63

- “Unit of Work Stop Timestamp” on page 64
- “Unit of Work Completion Status” on page 65
- “Previous Transaction Stop Time” on page 65
- “Application Idle Time” on page 66

## Application Handle (agent ID)

Snapshot Information Level	API Structure(s)	Monitor Switch
Application	sqlm_appl_id_info	Basic
Lock	sqlm_appl_lock	Basic
DCS Application	sqlm_dcs_applinfo	Basic
<b>Resettable</b>	No	
<b>Event Type</b>	<b>Event Record(s)</b>	
Connection	sqlm_connheader_event	
Statement	sqlm_stmt_event	
	sqlm_subsection_event	
<b>API Element Name</b>	agent_id	
<b>Element Type</b>	information	
<b>Related Information</b>	• None	

**Description:** A system-wide **unique** ID for the application. On multi-node systems, where a database is partitioned, this ID will be the same on every node where the application may make a secondary connection.

**Usage:** The application handle can be used to uniquely identify an active application (application handle is synonymous with agent Id).

**Note:** The name agent\_id is still used in the APIs, so that old applications can still be compiled with the Version 5 header files.

It can be used as input to GET SNAPSHOT commands that require an agent Id.

When reading event traces, it can be used to match event records with a given application.

It can be used as input to the FORCE APPLICATION command or API. On multi-node systems this command can be issued from any node where the application has a connection. Its effect is global

## Application Status

<b>Snapshot Information Level</b>	<b>API Structure(s)</b>	<b>Monitor Switch</b>
Application	sqlm_appl_id_info	Basic
Lock	sqlm_appl_lock	Basic
<b>Resettable</b>	No	
<b>API Element Name</b>	appl_status	
<b>Element Type</b>	information	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Application Status Change Time” on page 48</li> <li>• “Statement Operation” on page 176</li> </ul>	

**Description:** The current status of the application.

**Usage:** This element can help you diagnose potential application problems. Values for this field are:

API Constant	Description
SQLM_CONNECTPEND	<b>Database Connect Pending:</b> The application has initiated a database connection but the request has not yet completed.
SQLM_CONNECTED	<b>Database Connect Completed:</b> The application has initiated a database connection and the request has completed.
SQLM_UOWEXEC	<b>Unit of Work Executing:</b> The database manager is executing requests on behalf of the unit of work.
SQLM_UOWWAIT	<b>Unit of Work waiting:</b> The database manager is waiting on behalf of the unit of work in the application. This status typically means that the system is executing in the application's code.
SQLM_LOCKWAIT	<b>Lock Wait:</b> The unit of work is waiting for a lock. After the lock is granted, the status is restored to its previous value.
SQLM_COMMIT_ACT	<b>Commit Active:</b> The unit of work is committing its database changes.
SQLM_ROLLBACK_ACT	<b>Rollback Active:</b> The unit of work is rolling back its database changes.
SQLM_RECOMP	<b>Recompiling:</b> The database manager is recompiling (that is, rebinding) a plan on behalf of the application.
SQLM_COMP	<b>Compiling:</b> The database manager is compiling an SQL statement or precompiling a plan on behalf of the application.
SQLM_INTR	<b>Request Interrupted:</b> An interrupt of a request is in progress.



API Constant	Description
SQLM_DISCONNECTPEND	<b>Database Disconnect Pending:</b> The application has initiated a database disconnect but the command has not yet completed executing. The application may not have explicitly executed the database disconnect command. The database manager will disconnect from a database if the application ends without disconnecting.
SQLM_TPREP	<b>Transaction Prepared:</b> The unit of work is part of a global transaction that has entered the prepared phase of the two-phase commit protocol.
SQLM_THCOMT	<b>Transaction Heuristically Committed:</b> The unit of work is part of a global transaction that has been heuristically committed.
SQLM_THABRT	<b>Transaction Heuristically Rolled Back:</b> The unit of work is part of a global transaction that has been heuristically rolled-back.
SQLM_TEND	<b>Transaction Ended:</b> The unit of work is part of a global transaction that has ended but has not yet entered the prepared phase of the two-phase commit protocol.
SQLM_CREATE_DB	<b>Creating Database:</b> The agent has initiated a request to create a database and that request has not yet completed.
SQLM_RESTART	<b>Restarting Database:</b> The application is restarting a database in order to perform crash recovery.
SQLM_RESTORE	<b>Restoring Database:</b> The application is restoring a backup image to the database.
SQLM_BACKUP	<b>Backing Up Database:</b> The application is performing a backup of the database.
SQLM_LOAD	<b>Data Fast Load:</b> The application is performing a “fast load” of data into the database.
SQLM_UNLOAD	<b>Data Fast Unload:</b> The application is performing a “fast unload” of data from the database.
SQLM_IOERROR_WAIT	<b>Wait to Disable Table space:</b> The application has detected an I/O error and is attempting to disable a particular table space. The application has to wait for all other active transactions on the table space to complete before it can disable the table space.
SQLM_QUIESCE_TABLESPACE	<b>Quiescing a Table space:</b> The application is performing a quiesce table space request.

## ID of Code Page Used by Application

Snapshot Information Level	API Structure(s)	Monitor Switch
Application	sqlm_appl_id_info	Basic
Lock	sqlm_appl_lock	Basic
DCS Application	sqlm_dcs_applinfo	Basic
<b>Resettable</b>	No	
Event Type	Event Record(s)	
Event Log Header	sqlm_event_log_header	
Connection	sqlm_connheader_event	
API Element Name	codepage_id	
Element Type	information	
Related Information	<ul style="list-style-type: none"><li>• None</li></ul>	

**Description:** The code page identifier.

**Usage:** For snapshot monitor data, this is the code page at the node where the monitored application started. This identifier may be used for problem determination for remote applications. You may use this information to ensure that data conversion is supported between the application code page and the database code page (or for DRDA host databases, the host CCSID). For information about supported code pages, see the *Administration Guide*.

For event monitor data, this is the code page of the database for which event data is collected. You can use this element to determine whether your event monitor application is running under a different code page from that used by the database. Data written by the event monitor uses the database code page. If your event monitor application uses a different code page, you may need to perform some character conversion to make the data readable.

## Application Status Change Time

Snapshot Information Level	API Structure(s)	Monitor Switch
Application	sqlm_appl_id_info	Unit of Work
Lock	sqlm_appl_lock	Unit of Work
<b>Resettable</b>	No	
API Element Name	status_change_time	
Element Type	timestamp	
Related Information	<ul style="list-style-type: none"><li>• "Resetting Monitor Data" on page 21</li><li>• "Application Status" on page 46</li></ul>	

**Description:** The date and time the application entered its current status.

**Usage:** This element allows you to determine how long an application has been in its current status. If it has been in the same status for a long period of time, this may indicate that it has a problem.

## Application Name

<b>Snapshot Information Level</b>	<b>API Structure(s)</b>	<b>Monitor Switch</b>
Application	sqlm_appl_id_info	Basic
Lock	sqlm_appl_lock	Basic
DCS Application	sqlm_dcs_applinfo	Basic
<b>Resettable</b>	No	
<b>Event Type</b>	<b>Event Record(s)</b>	
Connection	sqlm_connheader_event	
<b>API Element Name</b>	appl_name	
<b>Element Type</b>	information	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Application ID” on page 50</li><li>• “ID of Code Page Used by Application” on page 48</li></ul>	

**Description:** The name of the application running at the client as known to the database manager or DDCS.

**Usage:** This element may be used with *Application ID* to relate data items with your application.

In a client/server environment, this name is passed from the client to the server to establish the database connection. For DRDA-AS connections, this name is the DRDA external name.

**Note:** Although the name of the application that is executing could be more than 20 bytes, only the first 20 bytes after the last path separator are used to set this element.

The application name is not available for applications running on the following down-level database client products:

- DB2 Version 1
- IBM Extended Services for OS/2

In situations where the client application code page is different from the code page under which the database system monitor is running, you can use *ID of Code Page Used by Application* to help translate *Application Name*.

## Application ID

<b>Snapshot Information Level</b>	<b>API Structure(s)</b>	<b>Monitor Switch</b>
Application	sqlm_appl_id_info	Basic
DCS Application	sqlm_dcs_applinfo	Basic
Lock	sqlm_appl_lock	Basic
<b>Resettable</b>	No	
<b>Event Type</b>	<b>Event Record(s)</b>	
Connection	sqlm_conn_event sqlm_connheader_event	
Statement	sqlm_stmt_event	
Transaction	sqlm_xaction_event	
Deadlock	sqlm_dlconn_event	
<b>API Element Name</b>	appl_id	
<b>Element Type</b>	information	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Outbound Application ID” on page 55</li> <li>• “Client Communication Protocol” on page 58</li> </ul>	

**Description:** This identifier is generated when the application connects to the database at the database manager or when DDCS receives a request to connect to a DRDA database.

**Usage:** This ID is known on both the client and server, so you can use it to correlate the client and server parts of the application. For DDCS applications, you will also need to use Outbound Application ID to correlate the client and server parts of the application.

This identifier is unique across the network. There are different formats for the application ID, which are dependent on the communication protocol between the client and the server machine on which the database manager and/or DDCS are running. Each of the formats consists of three parts separated by periods.

### 1. APPC

**Format** Network.LU Name.Application instance

**Example** CAIBMTOR.OSFDBX0.930131194520

**Details** This application ID is the displayable format of an actual SNA LUWID (Logical Unit-of-Work ID) that flows on the network when an APPC conversation is allocated. APPC-generated application IDs are made up by concatenating the network name, the LU name, and the LUWID instance number, which create a unique label for the client/server application. The network name and LU name can each be a maximum of 8 characters. The application instance corresponds to the 12-decimal-character LUWID instance number.

### 2. TCP/IP

**Format** \*TCP/IP.IPAddr.Application instance

**Example** \*TCP/IP.A12CF9E8.930131214645

**Details** A TCP/IP-generated application ID is made up by concatenating the string “\*TCPIP,” the IP address in hexadecimal characters, and a unique identifier for the instance of this application. The IP address is a 32-bit number displayed as a maximum of 8 hexadecimal characters.

### 3. IPX/SPX

**Format** Netid.nodeid.Application instance

**Example** C11A8E5C.400011528250.0131214645

**Details** An IPX/SPX-generated application ID is made up by concatenating a character network ID (8 hexadecimal characters), a node id (12 hexadecimal characters), and a unique identifier for the instance of the application. The application instance corresponds to a 10-decimal-character time stamp of the form MMDDHHMMSS.

### 4. NetBIOS

**Format** \*NETBIOS.nname.Application instance

**Example** \*NETBIOS.SBOIVIN.930131214645

**Details** A NetBIOS application ID is made up by concatenating the string “\*NETBIOS,” the nname defined in the client's database configuration file, and a unique identifier for the instance of this application.

### 5. Local Applications

**Format** \*LOCAL.DB2 instance.Application instance

**Example** \*LOCAL.DB2INST1.930131235945

**Details** The application ID generated for a local application is made up by concatenating the string \*LOCAL, the name of the DB2 instance, and a unique identifier for the instance of this application.

Use *Client Communication Protocol* to determine which communications protocol the connection is using and, as a result, the format of the *application ID*.

## Sequence Number

<b>Snapshot Information Level</b>	<b>API Structure(s)</b>	<b>Monitor Switch</b>
Application	sqlm_appl_id_info	Basic
DCS Application	sqlm_dcs_applinfo	Basic
<b>Resettable</b>	No	
<b>Event Type</b>	<b>Event Record(s)</b>	
Connection	sqlm_conn_event sqlm_connheader_event	
Statement	sqlm_stmt_event	
Transaction	sqlm_xaction_event	
Deadlock	sqlm_dlconn_event	
<b>API Element Name</b>	sequence_no	
<b>Element Type</b>	information	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• None</li></ul>	

**Description:** This element is reserved for future use. In this release, its value always be “0001.” It may contain different values in future releases of the product.

## Authorization ID

<b>Snapshot Information Level</b>	<b>API Structure(s)</b>	<b>Monitor Switch</b>
Application	sqlm_appl_id_info	Basic
Lock	sqlm_appl_lock	Basic
DCS Application	sqlm_dcs_applinfo	Basic
<b>Resettable</b>	No	
<b>Event Type</b>	<b>Event Record(s)</b>	
Connection	sqlm_connheader_event	
<b>API Element Name</b>	auth_id	
<b>Element Type</b>	information	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Application Name” on page 49</li></ul>	

**Description:** The authorization ID of the user who invoked the application that is being monitored. On a DDCS gateway node, this is the user's authorization ID on the host.

**Usage:** You can use this element to determine who is performing the monitoring, and to uniquely identify the application that is running.

If you are using the database system monitor APIs, note that the API constant `SQLM_IDENT_SZ` is used to define the length of this element. Only the first 8 characters are currently used.

## Configuration NNAME of Client

<b>Snapshot Information Level</b>	<b>API Structure(s)</b>	<b>Monitor Switch</b>
Application	sqlm_appl_id_info	Basic
DCS Application	sqlm_dcs_applinfo	Basic
<b>Resettable</b>	No	
<b>Event Type</b>	<b>Event Record(s)</b>	
Connection	sqlm_connheader_event	
<b>API Element Name</b>	client_nname	
<b>Element Type</b>	information	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Configuration NNAME at Monitoring (Server) Node” on page 34</li> </ul>	

**Description:** The *nname* in the database manager configuration file at the client node.

**Usage:** You can use this element to identify the client node that is running the application.

If you are using the database system monitor APIs, note that the API constant SQLM\_IDENT\_SZ is used to define the length of this element. Only the first 8 characters are currently used.

## Client Product/Version ID

<b>Snapshot Information Level</b>	<b>API Structure(s)</b>	<b>Monitor Switch</b>
Application	sqlm_appl_id_info	Basic
DCS Application	sqlm_dcs_applinfo	Basic
<b>Resettable</b>	No	
<b>Event Type</b>	<b>Event Record(s)</b>	
Connection	sqlm_connheader_event	
<b>API Element Name</b>	client_prdid	
<b>Element Type</b>	information	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Server Product/Version ID” on page 36</li> </ul>	

**Description:** The product and version that is running on the client.

**Usage:** You can use this element to identify the product and code version of the database client. It is in the form PPPVRRM, where:

- PPP identifies the product, which is “SQL” for the DB2 products
- VV identifies a 2-digit version number (with high-order 0 in the case of a 1-digit version)
- RR identifies a 2-digit release number (with high-order 0 in the case of a 1-digit release)
- M identifies a 1-digit modification level.

If you are using the database system monitor APIs, note that the API constant `SQLM_IDENT_SZ` is used to define the length of this element. Only the first 8 characters are currently used.

## Database Alias Used by Application

Snapshot Information Level	API Structure(s)	Monitor Switch
Application	sqlm_appl_id_info	Basic
Lock	sqlm_appl_lock	Basic
DCS Application	sqlm_dcs_applinfo	Basic
<b>Resettable</b>	No	
Event Type	Event Record(s)	
Connection	sqlm_connheader_event	
API Element Name	client_db_alias	
Element Type	information	
Related Information	<ul style="list-style-type: none"> <li>• All other database-level information</li> <li>• All other application-level information</li> <li>• “Last Reset Timestamp” on page 192</li> <li>• “Input Database Alias” on page 193</li> <li>• “Database Name” on page 39</li> </ul>	

**Description:** The alias of the database provided by the application to connect to the database.

**Usage:** This element can be used to identify the actual database that the application is accessing. The mapping between this name and *Database Name* could be done by using the database directories at the client node and the database manager server node.

This is the alias defined within the database manager where the database connection request originated.

If you are using the database system monitor APIs, note that the API constant `SQLM_IDENT_SZ` is used to define the length of this element. Only the first 8 characters are currently used.

This element can also be used to help you determine the authentication type, since different database aliases can have different authentication types.



## Host Product/Version ID

<b>Snapshot Information Level</b> DCS Application	<b>API Structure(s)</b> sqlm_dcs_applinfo	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	host_prdid information	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• None</li> </ul>	

**Description:** The product and version that is running on the server.

**Usage:** Used to identify the product and code version of the DRDA host database product. It is in the form PPPVRRM, where:

- PPP identifies the host DRDA product
  - ARI for DB2 for VSE and VM
  - DSN for DB2 for MVS/ESA
  - QSQ for DB2 for AS/400
  - SQL for other DB2 products.
- VV identifies a 2-digit version number (with high-order 0 in the case of a 1-digit version)
- RR identifies a 2-digit release number (with high-order 0 in the case of a 1-digit release)
- M identifies a 1-digit modification level

This field is defined as SQLM\_IDENT\_SZ characters long to allow for future expansion, but only the first 8 characters are presently used.

## Outbound Application ID

<b>Snapshot Information Level</b> DCS Application	<b>API Structure(s)</b> sqlm_dcs_applinfo	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	outbound_appl_id information	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Application ID” on page 50</li> </ul>	

**Description:** This identifier is generated when the application connects to the DRDA host database. It is used to connect the DDCS gateway to the host, while the *Application ID* is used to connect a client to the DDCS gateway.

**Usage:** You may use this element in conjunction with *Application ID* to correlate the client and server parts of the application information.

This identifier is unique across the network.

**Format** Network.LU Name.Application instance

**Example** CAIBMTOR.OSFDBM0.930131194520

**Details** This application ID is the displayable format of an actual SNA LUWID (Logical Unit-of-Work ID) that flows on the network when an APPC conversation is allocated. APPC-generated application IDs are made up by concatenating the network name, the LU name, and the LUWID instance number, which creates a unique label for the client/server application. The network name and LU name can each be a maximum of 8 characters. The application instance corresponds to the 12-decimal-character LUWID instance number.

## Outbound Sequence Number

<b>Snapshot Information Level</b> DCS Application	<b>API Structure(s)</b> sqlm_dcs_applinfo	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	outbound_sequence_no information	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>None</li> </ul>	

**Description:** This element is reserved for future use. In this release, its value will always be "0001." It may contain different values in future releases of the product.

## User Login ID

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_applinfo sqlm_appl	<b>Monitor Switch</b> Basic Basic
<b>Resettable</b>	No	
<b>Event Type</b> Connection	<b>Event Record(s)</b> sqlm_connheader_event	
<b>API Element Name</b> <b>Element Type</b>	execution_id information	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>"Authorization ID" on page 52</li> </ul>	

**Description:** The ID that the user specified when logging in to the operating system. This ID is distinct from Authorization ID, which the user specifies when connecting to the database.

**Usage:** You can use this element to determine who is running the application that you are monitoring.

For operating systems such as OS/2 that do not support the concept of "logging in," this field will be the same as Authorization ID .

If you are using the database system monitor APIs, note that the API constant SQLM\_IDENT\_SZ is used to define the length of this element. Only the first 8 characters are currently used.

## DRDA Correlation Token

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_applinfo sqlm_appl	<b>Monitor Switch</b> Basic Basic
<b>Resettable</b>	No	
<b>Event Type</b> Connection	<b>Event Record(s)</b> sqlm_connheader_event	
<b>API Element Name</b> <b>Element Type</b>	corr_token information	
<b>Related Information</b>	<ul style="list-style-type: none"><li>None</li></ul>	

**Description:** The DRDA AS correlation token.

**Usage:** The DRDA correlation token is used for correlating the processing between the application server and the application requester. It is an identifier dumped into logs when errors arise, that you can use to identify the conversation that is in error. In some cases, it will be the LUWID of the conversation.

If communications are not using DRDA, this element is blank.

If you are using the database system monitor APIs, note that the API constant SQLM\_APPLID\_SZ is used to define the length of this element.

## Client Process ID

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_applinfo sqlm_appl	<b>Monitor Switch</b> Basic Basic
<b>Resettable</b>	No	
<b>Event Type</b> Connection	<b>Event Record(s)</b> sqlm_connheader_event	
<b>API Element Name</b> <b>Element Type</b>	client_pid information	
<b>Related Information</b>	<ul style="list-style-type: none"><li>None</li></ul>	

**Description:** The process ID of the client application that made the connection to the database.

**Usage:** You can use this element to correlate monitor information such as CPU and I/O time to your client application.

In the case of a DRDA AS connection, this element will be set to 0.

## Client Operating Platform

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_applinfo sqlm_appl	<b>Monitor Switch</b> Basic Basic
<b>Resettable</b>	No	
<b>Event Type</b> Connection	<b>Event Record(s)</b> sqlm_connheader_event	
<b>API Element Name</b> <b>Element Type</b>	client_platform information	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>“Server Operating System” on page 37</li> </ul>	

**Description:** The operating system on which the client application is running.

**Usage:** This element can be used for problem determination for remote applications. Values for this field can be found in the header file *sqlmon.h*.

## Client Communication Protocol

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_applinfo sqlm_appl	<b>Monitor Switch</b> Basic Basic
<b>Resettable</b>	No	
<b>Event Type</b> Connection	<b>Event Record(s)</b> sqlm_connheader_event	
<b>API Element Name</b> <b>Element Type</b>	client_protocol information	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>None</li> </ul>	

**Description:** The communication protocol that the client application is using to communicate with the server.

**Usage:** This element can be used for problem determination for remote applications. Values for this field are:

API Constant	Communication Protocol
SQLM_PROT_UNKNOWN	(note 1)
SQLM_PROT_LOCAL	none (note 2)
SQLM_PROT_APPC	APPCC
SQLM_PROT_TCPIP	TCP/IP
SQLM_PROT_IPXSPX	IPX/SPX
SQLM_PROT_NETBIOS	NETBIOS

**Notes:**

1. The client is communicating using an unknown protocol. This value will only be returned if future clients connect with a down-level server.
2. The client is running on the same node as the server and no communications protocol is in use.

**Database Country Code**

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_applinfo sqlm_appl	<b>Monitor Switch</b> Basic Basic
<b>Resettable</b>	No	
<b>Event Type</b> Event Log Header Connection	<b>Event Record(s)</b> sqlm_event_log_header sqlm_connheader_event	
<b>API Element Name</b> <b>Element Type</b>	country_code information	
<b>Related Information</b>	• None	

**Description:** The country code of the database for which the monitor data is collected. Country code information is recorded in the database configuration file (see the *Administration Guide*).

For DRDA AS connections, this element will be set to 0.

**Application Agent Priority**

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_appl	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>Event Type</b> Connection	<b>Event Record(s)</b> sqlm_conn_event	
<b>API Element Name</b> <b>Element Type</b>	appl_priority information	
<b>Related Information</b>	• “Application Priority Type” on page 60	

**Description:** The priority of the agents working for this application.

**Usage:** You can use this element to check if applications are running with the expected priorities. Application priorities can be set by an administrator. They can be changed by the governor utility (**db2gov**).

The governor is used by DB2 to monitor and change the behavior of applications running against a database. This information is used to schedule applications and balance system resources.

A governor daemon collects statistics about the applications by taking snapshots. It checks these statistics against the rules governing applications running on that data-base. If the governor detects a rule violation, it takes the appropriate action. These rules and actions were specified by you in the governor configuration file.

If the action associated with a rule is to change an application's priority, the governor changes the priority of the agents in the partition where the violation was detected.

See the *Administration Guide* for more information on the governor.

## Application Priority Type

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_appl	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>Event Type</b> Connection	<b>Event Record(s)</b> sqlm_conn_event	
<b>API Element Name</b> <b>Element Type</b>	appl_priority_type information	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Query Cost Estimate” on page 183</li> <li>• “Application Agent Priority” on page 59</li> </ul>	

**Description:** Operating system priority for the application that is running.

**Usage:** Dynamic priority is recalculated by the operating system based on usage. Static priority does not change.

## User Authorization Level

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_appl sqlm_applinfo	<b>Monitor Switch</b> Basic Basic
<b>Resettable</b>	No	
<b>Event Type</b> Connection	<b>Event Record(s)</b> sqlm_conn_event	
<b>API Element Name</b> <b>Element Type</b>	authority_lvl information	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• None</li> </ul>	

**Description:** The highest authority level granted to an application.

**Usage:** The operations allowed by an application are granted either directly or indirectly in the *sql.h*.

These are the authorizations granted explicitly to a user:

- SQL\_SYSADMIN
- SQL\_DBADM

- SQL\_CREATETAB
- SQL\_BINDADD
- SQL\_CONNECT
- SQL\_CREATE\_NOT\_FENC
- SQL\_SYSCTRL
- SQL\_SYSMANT

The following are indirect authorizations inherited from group or public:

- SQL\_SYSADM\_GRP
- SQL\_DBADM\_GRP
- SQL\_CREATETAB\_GRP
- SQL\_BINDADD\_GRP
- SQL\_CONNECT\_GRP
- SQL\_CREATE\_NOT\_FENC\_GRP
- SQL\_SYSCTRL\_GRP
- SQL\_SYSMANT\_GRP

See the *Administration Guide* for detailed information on authority levels.

## Coordinating Node

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_appl	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>Event Type</b> Connection	<b>Event Record(s)</b> sqlm_conn_event	
<b>API Element Name</b> <b>Element Type</b>	coord_node information	
<b>Related Information</b>	• None	

**Description:** In a multi-node system, the node number of the node where the application connected or attached to the instance.

**Usage:** Each connected application is served by one coordinator node.

## Connection Request Start Timestamp

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_appl	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	appl_con_time timestamp	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Connection Request Completion Timestamp” on page 62</li><li>• All information related to the application</li></ul>	

**Description:** The date and time that an application started a connection request.

**Usage:** Use this element to determine when the application started its connection request to the database.

## Connection Request Completion Timestamp

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_appl	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	conn_complete_time timestamp	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Connection Request Start Timestamp” on page 62</li><li>• All information related to the application</li></ul>	

**Description:** The date and time that a connection request was granted.

**Usage:** Use this element to determine when a connection request to the database was granted.

## Previous Unit of Work Completion Timestamp

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_appl	<b>Monitor Switch</b> Unit of Work
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	prev_uow_stop_time timestamp	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “Unit of Work Start Timestamp” on page 63</li><li>• “Unit of Work Stop Timestamp” on page 64</li><li>• “Connection Request Completion Timestamp” on page 62</li></ul>	

**Description:** This is the time the unit of work completed.

**Usage:** You may use this element with *Unit of Work Stop Timestamp* to calculate the total elapsed time between COMMIT/ROLLBACK points, and with *Unit of Work Start*



*Timestamp* to calculate the time spent in the application between units of work. The time of one of the following:

- For applications currently within a unit of work, this is the time that the latest unit of work completed.
- For applications not currently within a unit of work (the application has completed a unit of work, but not yet started a new one), this is the stop time of the last unit of work that completed prior to the one that just completed. The stop time of the one just completed is indicated “Unit of Work Stop Timestamp” on page 64.
- For applications within their first unit of work, this is the database connection request completion time.

## Unit of Work Start Timestamp

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_appl	<b>Monitor Switch</b> Unit of Work
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	uow_start_time timestamp	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “Unit of Work Stop Timestamp” on page 64</li> <li>• “Previous Unit of Work Completion Timestamp” on page 62</li> <li>• “Connection Request Completion Timestamp” on page 62</li> </ul>	

**Description:** The date and time that the unit of work first required database resources.

**Usage:** This resource requirement occurs at the first SQL statement execution of that unit of work:

- For the first unit of work, it is the time of the first database request (SQL statement execution) after *Connection Request Completion Timestamp* .
- For subsequent units of work, it is the time of the first database request (SQL statement execution) after the previous COMMIT or ROLLBACK.

**Note:** The *SQL Reference* defines the boundaries of a unit of work as the COMMIT or ROLLBACK points.

The database system monitor excludes the time spent between the COMMIT/ROLLBACK and the next SQL statement from its definition of a unit of work. This measurement method reflects the time spent by the database manager in processing database requests, separate from time spent in application logic before the first SQL statement of that unit of work. The unit of work elapsed time does include the time spent running application logic between SQL statements within the unit of work.

You may use this element with *Unit of Work Stop Timestamp* to calculate the total elapsed time of the unit of work and with *Previous Unit of Work Completion Timestamp* to calculate the time spent in the application between units of work.

You can use the *Unit of Work Stop Timestamp* and the *Previous Unit of Work Completion Timestamp* to calculate the elapsed time for the SQL Reference's definition of a unit of work.

## Unit of Work Stop Timestamp

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_appl	<b>Monitor Switch</b> Unit of Work
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	uow_stop_time timestamp	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “Unit of Work Start Timestamp” on page 63</li> <li>• “Previous Unit of Work Completion Timestamp” on page 62</li> <li>• “Connection Request Completion Timestamp” on page 62</li> </ul>	

**Description:** The date and time that the most recent unit of work completed, which occurs when database changes are committed or rolled back.

**Usage:** You may use this element with *Previous Unit of Work Completion Timestamp* to calculate the total elapsed time between COMMIT/ROLLBACK points, and with *Unit of Work Start Timestamp* to calculate the elapsed time of the latest unit of work.

The timestamp contents will be set as follows:

- When the application has completed a unit of work and has not yet started a new one (as defined in *Unit of Work Start Timestamp*), this element will be a valid, non-zero timestamp
- When the application is currently executing a unit of work, this element will contain zeros
- When the application first connects to the database, this element is set to *Connection Request Completion Timestamp*.

As a new unit of work is started, the contents of this element are moved to *Previous Unit of Work Completion Timestamp*.

## Unit of Work Completion Status

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_appl	<b>Monitor Switch</b> Unit of Work
<b>Resettable</b>	No	
<b>Event Type</b> Transaction	<b>Event Record(s)</b> sqlm_xaction_event	
<b>API Element Name</b>	uow_comp_status (Snapshot) status (Event)	
<b>Element Type</b>	information	
<b>Related Information</b>	• “Resetting Monitor Data” on page 21	

**Description:** The status of the unit of work and how it stopped.

**Usage:** You may use this element to determine if the unit of work ended due to a deadlock or abnormal termination. It may have been:

- Committed due to a commit statement
- Rolled back due to a rollback statement
- Rolled back due to a deadlock
- Rolled back due to an abnormal termination
- Committed at normal application termination.

**Note:** API users should refer to the header file (*sqlmon.h*) containing definitions of database system monitor constants.

## Previous Transaction Stop Time

<b>Event Type</b> Transaction	<b>Event Record(s)</b> sqlm_xaction_event
<b>API Element Name</b> <b>Element Type</b>	prev_stop_time timestamp
<b>Related Information</b>	• None

**Description:** The completion time of the last unit of work.

**Usage:** You may use this element to calculate the time spent in the application between units of work.

This is the unit of work that completed prior to the unit of work for which this transaction event is generated.

For applications within their first unit of work, this is the database connection request completion time.

## Application Idle Time

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_appl	<b>Monitor Switch</b> Statement
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	appl_idle_time information	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Query Cost Estimate” on page 183</li><li>• “Application Agent Priority” on page 59</li><li>• “Application Priority Type” on page 60</li></ul>	

**Description:** Number of seconds since an application has issued any requests to the server. This includes applications that have not terminated a transaction, for example not issued a commit or rollback.

**Usage:** This information can be used to implement applications that force users that have been idle for a specified number of seconds.

## DB2 Agent Information

The following database system monitor elements provide information about agents:

- “Process or Thread ID”

### Process or Thread ID

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_appl sqlm_agent	<b>Monitor Switch</b> Basic Statement
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	coord_agent_pid agent_pid information	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• None</li></ul>	

**Description:** The process Id (UNIX systems) or thread Id (OS/2 or Windows systems) of a DB2 agent.

**Usage:** You can use this element to link database system monitor information to other sources of diagnostic information, such as system traces. You can also use it to monitor how agents working for a database application use system resources.

## Database Manager Configuration

The following elements provide database manager configuration information.

## Agents and Connections

Agents are how work gets done using the database manager. An agent is a process or thread that carries out the requests made by a client application. Each connected application is served by exactly 1 **coordinator agent** and possibly, a set of subordinator agents or **subagents**. Subagents are used for parallel SQL processing in partitioned databases and on SMP machines. Agents are classified as follows:

- **Coordinator agent**

This is the initial agent to which a local or remote application connects. There is one coordinator agent dedicated to each database connection or instance attachment. The maximum number of coordinating agents per partition is controlled by the *max\_coordagents* configuration parameter.

- **Subagent**

In partitioned databases, additional agents can be enlisted by the coordinator agent to speed up SQL processing. Subagents are selected from the agent pool and are returned there when their work is done. The size of the agent pool is controlled by the *num\_poolagents* configuration parameter.

- **Associated agent**

A coordinator or subagent that is doing work for an application is associated with that application. After it is finished an application's work, it goes into the agent pool as an idle agent, but it remains associated with that application. If the application attempts to do more work, DB2 will search the agent pool for an agent already associated with the application and assign the work to it. If none is found, DB2 will attempt to get an agent to satisfy the request by:

1. Choosing an idle agent that is not associated with an application.
2. Creating an agent, if an idle agent is not available.
3. Finding an agent that is associated with another application. For example, if an agent cannot be created because *maxagents* has been reached, DB2 will try to take an idle agent associated with another application. This is referred to as a **stolen agent**.

The *maxagents* configuration parameter defines the maximum number of agents, regardless of type, that can exist for an instance. The *maxagents* value does not create any agents. The initial number of agents that are created in the agent pool at DB2START is determined by the *num\_initagents* configuration parameter.

Assuming no idle agents, each connection creates a new agent, unless *max\_coordagents* has been reached. If subagents are not used, *max\_coordagents* equals *maxagents*. If subagents are used, some combination of coordinator agents and subagents could reach *maxagents*.

When an agent is assigned work, it attempts to obtain a token or permission to process the transaction. The database manager controls the number of tokens available using the *maxcagents* configuration parameter. If a token is not available, the agent will sleep until one becomes available, at which time the requested work will be processed. This

allows you to use *maxcagents* to control the load, or number of concurrently executing transactions, the server handles.

The following elements provide agent and connection information:

- “Remote Connections To Database Manager”
- “Remote Connections Executing in the Database Manager” on page 69
- “Local Connections” on page 69
- “Local Connections Executing in the Database Manager” on page 70
- “Local Databases with Current Connects” on page 71
- “Connects Since First Database Connect” on page 71
- “Applications Connected Currently” on page 72
- “Applications Executing in the Database Currently” on page 72
- “Agents Registered” on page 73
- “Agents Waiting for a Token” on page 73
- “Maximum Number of Agents Registered” on page 74
- “Maximum Number of Agents Waiting” on page 74
- “Number of Idle Agents” on page 75
- “Agents Assigned From Pool” on page 75
- “Agents Created Due to Empty Agent Pool” on page 76
- “Maximum Number of Coordinating Agents” on page 76
- “Stolen Agents” on page 77
- “Maximum Number of Associated Agents” on page 77
- “Committed Private Memory” on page 78
- “Secondary Connections” on page 78

## Remote Connections To Database Manager

<b>Snapshot Information Level</b>	<b>API Structure(s)</b>	<b>Monitor Switch</b>
Database Manager	sqlm_db2	Basic
<b>Resettable</b>	No	
<b>API Element Name</b>	rem_cons_in	
<b>Element Type</b>	gauge	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Remote Connections Executing in the Database Manager” on page 69</li> <li>• “Local Connections” on page 69</li> <li>• “Local Connections Executing in the Database Manager” on page 70</li> </ul>	

**Description:** The total number of current connections initiated from remote clients to the instance of the database manager that is being monitored.

**Usage:** Shows the number of connections from remote clients to databases in this instance. This value will change frequently, so you may need to sample it at specific intervals over an extended period of time to get a realistic view of system usage. This number does not include applications that were initiated from the same instance as the database manager.

When used in conjunction with the Local Connections monitor element, these elements can help you adjust the setting of the *max\_coordagents* configuration parameter, described in the *Administration Guide*.

## Remote Connections Executing in the Database Manager

Snapshot Information Level	API Structure(s)	Monitor Switch
Database Manager	sqlm_db2	Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	rem_cons_in_exec gauge	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Remote Connections To Database Manager” on page 68</li> <li>• “Local Connections” on page 69</li> <li>• “Local Connections Executing in the Database Manager” on page 70</li> </ul>	

**Description:** The number of remote applications that are currently connected to a database and are currently processing a unit of work within the database manager instance being monitored.

**Usage:** This number can help you determine the level of concurrent processing occurring on the database manager. This value will change frequently, so you may need to sample it at specific intervals over an extended period of time to get a realistic view of system usage. This number does not include applications that were initiated from the same instance as the database manager.

When used in conjunction with the Local Connections Executing in the Database Manager monitor element, this element can help you adjust the setting of the *maxcagents* configuration parameter, described in the *Administration Guide*.

## Local Connections

Snapshot Information Level	API Structure(s)	Monitor Switch
Database Manager	sqlm_db2	Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	local_cons gauge	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Remote Connections To Database Manager” on page 68</li> <li>• “Remote Connections Executing in the Database Manager” on page 69</li> <li>• “Local Connections Executing in the Database Manager” on page 70</li> </ul>	

**Description:** The number of local applications that are currently connected to a local database within the database manager instance being monitored.

**Usage:** This number can help you determine the level of concurrent processing occurring in the database manager. This value will change frequently, so you may need to sample it at specific intervals over an extended period of time to get a realistic view of system usage.

This number only includes applications that were initiated from the same instance as the database manager. The applications are connected, but may or may not be executing a unit of work in the database.

When used in conjunction with the Remote Connections To Database Manager monitor element, this element can help you adjust the setting of the *maxagents* configuration parameter, described in the *Administration Guide*.

## Local Connections Executing in the Database Manager

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_db2	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	local_cons_in_exec gauge	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Remote Connections To Database Manager” on page 68</li> <li>• “Remote Connections Executing in the Database Manager” on page 69</li> <li>• “Local Connections” on page 69</li> </ul>	

**Description:** The number of local applications that are currently connected to a local database within the database manager instance being monitored and are currently processing a unit of work.

**Usage:** This number can help you determine the level of concurrent processing occurring in the database manager. This value will change frequently, so you may need to sample it at specific intervals over an extended period of time to get a realistic view of system usage. This number only includes applications that were initiated from the same instance as the database manager.

When used in conjunction with the Remote Connections Executing in the Database Manager monitor element, this element can help you adjust the setting of the *maxcagents* configuration parameter, described in the *Administration Guide*.



## Local Databases with Current Connects

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_db2	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	con_local_dbases gauge	
<b>Related Information</b>	• None	

**Description:** The number of local databases that have applications connected.

**Usage:** This value gives an indication of how many database information records you can expect when gathering data at the database level.

The applications can be running locally or remotely, and may or may not be executing a unit of work within the database manager

## Connects Since First Database Connect

<b>Snapshot Information Level</b> Database	<b>API Structure(s)</b> sqlm_dbase	<b>Monitor Switch</b> Basic
<b>Resettable</b>	Yes	
<b>Event Type</b> Database	<b>Event Record(s)</b> sqlm_db_event	
<b>API Element Name</b> <b>Element Type</b>	total_cons counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “When Counters are Initialized” on page 20</li><li>• “Database Activation Timestamp” on page 41</li><li>• “Applications Connected Currently” on page 72</li><li>• “Applications Executing in the Database Currently” on page 72</li><li>• “Secondary Connections” on page 78</li></ul>	

**Description:** Indicates the number of connections to the database (coordinator agents).

**Usage:** You can use this element in conjunction with the Database Activation Timestamp and the Start Database Manager Timestamp monitor elements to calculate the frequency at which applications have connected to the database.

If the frequency of connects is low, you may want to do a dummy connect initially before connecting any other application, because of the extra overhead that sometimes accompanies the first connect to a database (for example, initial buffer pool allocation or autorestart). This will result in subsequent connects being processed at a higher rate. You could also use the ACTIVATE DATABASE command.

**Note:** When you reset this element, its value is set to the number of applications that are currently connected, not to zero.

## Applications Connected Currently

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
Lock	sqlm_dbase_lock	Basic
<b>Resettable</b>	No	
<b>API Element Name</b>	appls_cur_cons	
<b>Element Type</b>	gauge	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Applications Executing in the Database Currently” on page 72</li><li>• “Connects Since First Database Connect” on page 71</li><li>• “Remote Connections To Database Manager” on page 68</li><li>• “Local Connections” on page 69</li></ul>	

**Description:** Indicates the number of applications that are currently connected to the database.

**Usage:** You may use this element to help you understand the level of activity within a database and the amount of system resource being used.

It can help you adjust the setting of the *maxappls* and *max\_coordagents* configuration parameters, which are described in the *Administration Guide*. For example, its value is always the same as *maxappls*, you may want to increase the value of *maxappls*. See the *Remote Connections To Database Manager* and the *Local Connections* monitor elements for more information.

## Applications Executing in the Database Currently

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
<b>Resettable</b>	No	
<b>API Element Name</b>	appls_in_db2	
<b>Element Type</b>	gauge	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Applications Connected Currently” on page 72</li><li>• “Connects Since First Database Connect” on page 71</li><li>• “Remote Connections Executing in the Database Manager” on page 69</li><li>• “Local Connections Executing in the Database Manager” on page 70</li><li>• “Current Agents Waiting On Locks” on page 144</li></ul>	

**Description:** Indicates the number of applications that are currently connected to the database, and for which the database manager is currently processing a request.

**Usage:** You can use this element to understand how many of the database manager agent tokens are being used by applications connected to this database. If the sum of

*Remote Connections Executing in the Database Manager* and *Local Connections Executing in the Database Manager* is equal to the value of the *maxcagents* configuration parameter, you may want to increase the value of that parameter, as described in the *Administration Guide*.

## Agents Registered

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_db2	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	agents_registered gauge	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• None</li> </ul>	

**Description:** The number of agents registered in the database manager instance that is being monitored (coordinator agents and subagents).

**Usage:** You can use this element to help evaluate your setting for the *maxagents* configuration parameter.

## Agents Waiting for a Token

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_db2	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	agents_waiting_on_token gauge	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Agents Registered” on page 73</li> </ul>	

**Description:** The number of agents waiting for a token so they can execute a transaction in the database manager.

**Usage:** You can use this element to help evaluate your setting for the *maxcagents* configuration parameter.

Each application has a dedicated coordinator agent to process database requests within the database manager. Each agent has to get a token before it can execute a transaction. The maximum number of agents that can execute database manager transactions is limited by the configuration parameter *maxcagents*. For more information about this parameter, see the *Administration Guide*.

## Maximum Number of Agents Registered

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_db2	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	agents_registered_top water mark	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Agents Registered” on page 73</li><li>• “Maximum Number of Agents Waiting” on page 74</li></ul>	

**Description:** The maximum number of agents that the database manager has ever registered, at the same time, since it was started (coordinator agents and subagents).

**Usage:** You may use this element to help you evaluate your setting of the *maxagents* configuration parameter, described in the *Administration Guide*.

The number of agents registered at the time the snapshot was taken is recorded by Agents Registered.

## Maximum Number of Agents Waiting

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_db2	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	agents_waiting_top water mark	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Agents Waiting for a Token” on page 73</li><li>• “Maximum Number of Agents Registered” on page 74</li></ul>	

**Description:** The maximum number of agents that have ever been waiting for a token, at the same time, since the database manager was started.

**Usage:** You may use this element to help you evaluate your setting of the *maxcagents* configuration parameter, described in the *Administration Guide*.

The number of agents waiting for a token at the time the snapshot was taken is recorded by *Agents Waiting for a Token*.

If the *maxcagents* parameter is set to its default value (-1), no agents should wait for a token and the value of this monitor element should be zero.

## Number of Idle Agents

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_db2	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	idle_agents gauge	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Maximum Number of Agents Registered” on page 74</li><li>• “Maximum Number of Agents Waiting” on page 74</li><li>• “Agents Registered” on page 73</li></ul>	

**Description:** The number of agents in the agent pool that are currently unassigned to an application and are, therefore, “idle.”

**Usage:** You can use this element to help set the *num\_poolagents* configuration parameter. Having idle agents available to service requests for agents can improve performance. See the *Administration Guide* for more information.

## Agents Assigned From Pool

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_db2	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	agents_from_pool counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Agents Created Due to Empty Agent Pool” on page 76</li><li>• “Maximum Number of Coordinating Agents” on page 76</li></ul>	

**Description:** The number of agents assigned from the agent pool.

**Usage:** This element can be used with “*Agents Created Due to Empty Agent Pool*” on page 76 to determine how often an agent must be created because the pool is empty.

If the ratio of

$$\text{Agents Created Due to Empty Agent Pool} / \text{Agents Assigned From Pool}$$

is high, it may indicate that the *num\_poolagents* configuration parameter should be increased. A low ratio suggests that *num\_poolagents* is set too high, and that some of the agents in the pool are rarely used and are wasting system resources.

A high ratio can indicate that the overall workload for this node is too high. You can adjust the workload by lowering the maximum number of coordinating agents specified by the *maxcagents* configuration parameter, or by redistributing data among the nodes.

See the *Administration Guide* for more information on the Agent Pool Size (*num\_poolagents*) and Maximum Number of Concurrent Coordinating Agents (*maxcagents*) configuration parameters.

## Agents Created Due to Empty Agent Pool

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_db2	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	agents_created_empty_pool counter	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Agents Assigned From Pool” on page 75</li> <li>• “Maximum Number of Coordinating Agents” on page 76</li> </ul>	

**Description:** The number of agents created because the agent pool was empty.

**Usage:** In conjunction with Agents Assigned From Pool, you can calculate the ratio of  
 Agents Created Due to Empty Agent Pool / Agents Assigned From Pool

See “Agents Assigned From Pool” on page 75 for information on using this element.

## Maximum Number of Coordinating Agents

<b>Snapshot Information Level</b> Database Manager Database	<b>API Structure(s)</b> sqlm_db2 sqlm_dbase	Basic Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	coord_agents_top water mark	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Agents Assigned From Pool” on page 75</li> <li>• “Agents Created Due to Empty Agent Pool” on page 76</li> </ul>	

**Description:** The maximum number of coordinating agents working at one time.

**Usage:** If the peak number of coordinating agents represents too high a workload for this node, you can reduce the number that can be concurrently executing a transaction by changing the *maxcagents* configuration parameter.

See the *Administration Guide* for more information on the Maximum Number of Concurrent Coordinating Agents (*maxcagents*) configuration parameter.

## Stolen Agents

Snapshot Information Level	API Structure(s)	Monitor Switch
Database Manager	sqlm_db2	Basic
Application	sqlm_appl	Basic
<b>Resettable</b>	Yes	
<b>API Element Name</b>	agents_stolen	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Number of Agents Working on a Statement” on page 190</li></ul>	

**Description:** The number of times that agents are stolen from an application. Agents are stolen when an idle agent associated with an application is reassigned to work on a different application.

**Usage:** This element can be used in conjunction with “*Maximum Number of Associated Agents*” to evaluate the load that this application places on the system.

## Maximum Number of Associated Agents

Snapshot Information Level	API Structure(s)	Monitor Switch
Application	sqlm_appl	Basic
<b>Resettable</b>	No	
<b>API Element Name</b>	associated_agents_top	
<b>Element Type</b>	water mark	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Agents Assigned From Pool” on page 75</li><li>• “Agents Created Due to Empty Agent Pool” on page 76</li></ul>	

**Description:** The maximum number of subagents associated with this application.

**Usage:** If the peak number of subagents is close to the *num\_poolagents* configuration parameter, this might indicate too high a workload for this node.

See the *Administration Guide* for more information on the Agent Pool Size (*num\_poolagents*) configuration parameter.

## Committed Private Memory

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_db2	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	comm_private_mem gauge	
<b>Related Information</b>	• none	

**Description:** The amount of private memory that the instance of the database manager has currently committed at the time of the snapshot.

**Usage:** You can use this element to help set the *min\_priv\_mem* configuration parameter (see the *Administration Guide*) to ensure you have enough private memory available.

This element is only applicable to platforms where DB2 uses threads (such as OS/2).

## Secondary Connections

<b>Snapshot Information Level</b> Database	<b>API Structure(s)</b> sqlm_dbase	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	total_sec_cons counter	
<b>Related Information</b>	• “Connects Since First Database Connect” on page 71 • “Start Database Manager Timestamp” on page 34 • “Database Activation Timestamp” on page 41	

**Description:** The number of connections made by a subagent to the database at the node.

**Usage:** You can use this element in conjunction with the Connects Since First Database Connect, Database Activation Timestamp, and the Start Database Manager Timestamp monitor elements to calculate the frequency at which applications have connected to the database.

## Sort

The following elements provide information about the database manager sort work performed:

- “Total Sort Heap Allocated” on page 79
- “Post Threshold Sorts” on page 80
- “Piped Sorts Requested” on page 80
- “Piped Sorts Accepted” on page 81
- “Total Sorts” on page 82
- “Total Sort Time” on page 82



- “Sort Overflows” on page 83
- “Active Sorts” on page 84

## Total Sort Heap Allocated

Snapshot Information Level	API Structure(s)	Monitor Switch
Database Manager	sqlm_db2	Basic
Database	sqlm_dbase	Basic
<b>Resettable</b>	No	
<b>API Element Name</b>	sort_heap_allocated	
<b>Element Type</b>	gauge	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Total Sorts” on page 82</li> </ul>	

**Description:** The total number of allocated pages of sort heap space for all sorts at the level chosen and at the time the snapshot was taken.

**Usage:** The amount of memory allocated for each sort may be some or all of the available sort heap size. Sort heap size is the amount of memory available for each sort as defined in the database configuration parameter *sortheap*.

It is possible for a single application to have concurrent sorts active. For example, in some cases a SELECT statement with a subquery can cause concurrent sorts.

Information may be collected at two levels:

- At the database manager level, it represents the sum of sort heap space allocated for all sorts in all active databases in the database manager
- At the database level, it represents the sum of the sort heap space allocated for all sorts in a database.

Normal memory estimates do not include sort heap space. If excessive sorting is occurring, the extra memory used for the sort heap should be added to the base memory requirements for running the database manager. Generally, the larger the sort heap, the more efficient the sort. Appropriate use of indexes can reduce the amount of sorting required.

You may use the information returned at the database manager level to help you tune the *sheapthres* configuration parameter. If the element value is greater than or equal to *sheapthres*, it means that the sorts are not getting the full sort heap as defined by the *sortheap* parameter.

## Post Threshold Sorts

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_db2	<b>Monitor Switch</b> Sort
<b>Resettable</b>	Yes	
<b>API Element Name</b> <b>Element Type</b>	post_threshold_sorts counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “When Counters are Initialized” on page 20</li><li>• “Statement Sorts” on page 181</li><li>• “Active Sorts” on page 84</li></ul>	

**Description:** The number of sorts that have requested heaps after the sort heap threshold has been reached.

**Usage:** Under normal conditions, the database manager will allocate sort heap using the value specified by the *sortheap* configuration parameter. If the amount of memory allocated to sort heaps exceeds the sort heap threshold (*sheapthres* configuration parameter), the database manager will allocate sort heap using a value less than that specified by the *sortheap* configuration parameter.

Each active sort on the system allocates memory, which may result in sorting taking up too much of the system memory available. Sorts that start after the sort heap threshold has been reached may not receive an optimum amount of memory to execute, but, as a result, the entire system may benefit. By modifying the sort heap threshold and sort heap size configuration parameters, the performance of sort operations and/or the overall system can be improved. If this element's value is high, you can:

- Increase the sort heap threshold (*sheapthres*) or,
- Adjust applications to use fewer or smaller sorts via SQL query changes.

## Piped Sorts Requested

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_db2	<b>Monitor Switch</b> Basic
<b>Resettable</b>	Yes	
<b>API Element Name</b> <b>Element Type</b>	piped_sorts_requested counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “When Counters are Initialized” on page 20</li><li>• “Piped Sorts Accepted” on page 81</li><li>• “Post Threshold Sorts” on page 80</li></ul>	

**Description:** The number of piped sorts that have been requested.

**Usage:** Each active sort on the system allocates memory, which may result in sorting taking up too much of the available system memory.

The sort list heap (*sortheap*) and sort heap threshold (*sheapthres*) configuration parameters help to control the amount of memory used for sort operations. These parameters are also used to determine whether a sort will be piped.

Since piped sorts may reduce disk I/O, allowing more piped sorts can improve the performance of sort operations and possibly the performance of the overall system. A piped sort is not accepted if the sort heap threshold will be exceeded when the sort heap is allocated for the sort. See *Piped Sorts Accepted* for more information if you are experiencing piped sort rejections.

The SQL EXPLAIN output will show whether the optimizer requests a piped sort. For more information on piped and non-piped sorts see the *Administration Guide*.

## Piped Sorts Accepted

Snapshot Information Level	API Structure(s)	Monitor Switch
Database Manager	sqlm_db2	Basic
<b>Resettable</b>	Yes	
<b>API Element Name</b>	pipedsorts_accepted	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “When Counters are Initialized” on page 20</li> <li>• “Piped Sorts Requested” on page 80</li> <li>• “Post Threshold Sorts” on page 80</li> </ul>	

**Description:** The number of piped sorts that have been accepted.

**Usage:** Each active sort on the system allocates memory, which may result in sorting taking up too much of the available system memory.

When the number of accepted piped sorts is low compared to the number requested, you can improve sort performance by adjusting one or both of the following configuration parameters:

- sortheap
- sheapthres

If piped sorts are being rejected, you might consider decreasing your sort heap or increasing your sort heap threshold. You should be aware of the possible implications of either of these options. If you increase the sort heap threshold, then there is the possibility that more memory will remain allocated for sorting. This could cause the paging of memory to disk. If you decrease the sort heap, you might require an extra merge phase that could slow down the sort.

See the *Administration Guide* for more information on sorts.

## Total Sorts

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Sort
Application	sqlm_appl	Sort
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
Statement	sqlm_stmt_event	
API Element Name	total_sorts	
Element Type	counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “When Counters are Initialized” on page 20</li><li>• “Sort Overflows” on page 83</li><li>• “Total Sort Time” on page 82</li></ul>	

**Description:** The total number of sorts that have been executed.

**Usage:** At a database or application level, use this value with *Sort Overflows* to calculate the percentage of sorts that need more heap space. You can also use it with *Total Sort Time* to calculate the average sort time.

If the number of sort overflows is small with respect to the total sorts, then increasing the sort heap size may have little impact on performance, unless this buffer size is increased substantially.

At a statement level, use this element to identify statements which are performing large numbers of sorts. These statements may benefit from additional tuning to reduce the number of sorts. You can also use the SQL EXPLAIN statement to identify the number of sorts a statement performs. See the *Administration Guide* for more information.

## Total Sort Time

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Sort
Application	sqlm_appl	Sort
	sqlm_stmt	Sort
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
Statement	sqlm_stmt_event	
API Element Name	total_sort_time	
Element Type	counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “When Counters are Initialized” on page 20</li><li>• “Total Sorts” on page 82</li><li>• “Total Sorts” on page 82</li></ul>	

**Description:** The total elapsed time (in milliseconds) for all sorts that have been executed.

**Usage:** At a database or application level, use this element with *Total Sorts* to calculate the average sort time, which can indicate whether or not sorting is an issue as far as performance is concerned.

At a statement level, use this element to identify statements that spend a lot of time sorting. These statements may benefit from additional tuning to reduce the sort time.

This count also includes sort time of temporary tables created during related operations. It provides information for one statement, one application, or all applications accessing one database.

When using data elements providing elapsed times, you should consider:

1. Elapsed times are affected by system load, so the more processes you have running, the higher this elapsed time value.
2. To calculate this data element at a database level, the database system monitor sums the application-level times. This can result in double counting elapsed times at a database level, since more than one application process can be running at the same time.

To provide meaningful data from the database level, you should normalize the data to a lower level. For example:

```
total sort time / total sorts
```

provides information about the average elapsed time for each sort.

## Sort Overflows

<b>Snapshot Information Level</b>	<b>API Structure(s)</b>	<b>Monitor Switch</b>
Database	sqlm_dbase	Sort
Application	sqlm_appl sqlm_stmt	Sort Sort
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
Statement	sqlm_stmt_event	
<b>API Element Name</b>	sort_overflows	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “When Counters are Initialized” on page 20</li><li>• “Total Sorts” on page 82</li></ul>	

**Description:** The total number of sorts that ran out of sort heap and may have required disk space for temporary storage.

**Usage:** At a database or application level, use this element in conjunction with *Total Sorts* to calculate the percentage of sorts that had to overflow to disk. If this percentage is high, you may want adjust the database configuration by increasing the value of *sortheap*.

At a statement level, use this element to identify statements that require large sorts. These statements may benefit from additional tuning to reduce the amount of sorting required.

When a sort overflows, additional overhead will be incurred because the sort will require a merge phase and can potentially require more I/O, if data needs to be written to disk.

This element provides information for one statement, one application, or all applications accessing one database.

### Active Sorts

<b>Snapshot Information Level</b> Database	<b>API Structure(s)</b> sqlm_dbase	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	active_sorts counter	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Total Sort Heap Allocated” on page 79</li> <li>• “Total Sorts” on page 82</li> </ul>	

**Description:** The number of sorts in the database that currently have a sort heap allocated.

**Usage:** Use this value in conjunction with *Total Sort Heap Allocated* to determine the average sort heap space used by each sort. If the *sortheap* configuration parameter is substantially larger than the average sort heap used, you may be able to lower the value of this parameter. (See the *Administration Guide* for more details.)

This value includes heaps for sorts of temporary tables that were created during relational operations.

## Fast Communication Manager

The following database system monitor elements provide information about the Fast Communication Manager (FCM):

- “FCM Buffers Currently Free” on page 85
- “Minimum FCM Buffers Free” on page 85
- “Message Anchors Currently Free” on page 86
- “Minimum Message Anchors” on page 86
- “Connection Entries Currently Free” on page 87
- “Minimum Connection Entries” on page 87
- “Request Blocks Currently Free” on page 88
- “Minimum Request Blocks” on page 88

- “Number of Nodes” on page 89
- “Connection Status” on page 89
- “Total FCM Buffers Sent” on page 90
- “Total FCM Buffers Received” on page 90

### FCM Buffers Currently Free

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_fcm	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	buff_free gauge	
<b>Related Information</b>	• “Minimum FCM Buffers Free” on page 85	

**Description:** This element indicates the number of FCM buffers currently free.

**Usage:** Use the number of FCM buffers currently free in conjunction with the *fcm\_num\_buffers* configuration parameter to determine the current FCM buffer pool utilization. You can use this information to tune *fcm\_num\_buffers*.

### Minimum FCM Buffers Free

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_fcm	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	buff_free_bottom water mark	
<b>Related Information</b>	• “FCM Buffers Currently Free” on page 85	

**Description:** The lowest number of free FCM buffers reached during processing.

**Usage:** Use this element in conjunction with the *fcm\_num\_buffers* configuration parameter to determine the maximum FCM buffer pool utilization. If *buff\_free\_bottom* is low, you should increase *fcm\_num\_buffers* to ensure that operations do not run out of FCM buffers. If *buff\_free\_bottom* is high, you can decrease *fcm\_num\_buffers* to conserve system resources.

## Message Anchors Currently Free

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_fcm	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	MA_free gauge	
<b>Related Information</b>	• “Minimum Message Anchors” on page 86	

**Description:** This element indicates the number of message anchors currently free.

**Usage:** Use the number of message anchors currently free in conjunction with the *fcm\_num\_anchors* configuration parameter to determine the current message anchor utilization. You can use this information to tune *fcm\_num\_anchors*.

## Minimum Message Anchors

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_fcm	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	MA_free_bottom water mark	
<b>Related Information</b>	• “Message Anchors Currently Free” on page 86	

**Description:** The lowest number of free message anchors reached during processing.

**Usage:** Use this element in conjunction with the *fcm\_num\_anchors* configuration parameter to determine the maximum message anchors utilization. If *MA\_free\_bottom* is low, you should increase *fcm\_num\_anchors* to ensure that operations do not run out of message anchors. If *MA\_free\_bottom* is high, you can decrease *fcm\_num\_anchors* to conserve system resources.



## Connection Entries Currently Free

Snapshot Information Level	API Structure(s)	Monitor Switch
Database Manager	sqlm_fcm	Basic
<b>Resettable</b>	No	
<b>API Element Name</b>	CE_free	
<b>Element Type</b>	gauge	
<b>Related Information</b>	• “Minimum Connection Entries” on page 87	

**Description:** This element indicates the number of connection entries currently free.

**Usage:** Use the number of connection entries currently free in conjunction with the *fcm\_num\_connect* configuration parameter to determine the current connection entry utilization. You can use this information to tune *fcm\_num\_connect*.

## Minimum Connection Entries

Snapshot Information Level	API Structure(s)	Monitor Switch
Database Manager	sqlm_fcm	Basic
<b>Resettable</b>	No	
<b>API Element Name</b>	CE_free_bottom	
<b>Element Type</b>	water mark	
<b>Related Information</b>	• “Connection Entries Currently Free” on page 87	

**Description:** The lowest number of free connection entries reached during processing.

**Usage:** Use this element in conjunction with the *fcm\_num\_connect* configuration parameter to determine the maximum connection entry utilization. If *CE\_free\_bottom* is low, you should increase *fcm\_num\_connect* to ensure that operations do not run out of connection entries. If *CE\_free\_bottom* is high, you can decrease *fcm\_num\_connect* to conserve system resources.

## Request Blocks Currently Free

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_fcm	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	RB_free gauge	
<b>Related Information</b>	• “Request Blocks Currently Free” on page 88	

**Description:** This element indicates the number of request blocks currently free.

**Usage:** Use the number of request blocks currently free in conjunction with the *fcm\_num\_rqb* configuration parameter to determine the current request block utilization. You can use this information to tune *fcm\_num\_rqb*.

## Minimum Request Blocks

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_fcm	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	RB_free_bottom water mark	
<b>Related Information</b>	• “Request Blocks Currently Free” on page 88	

**Description:** The lowest number of free request blocks reached during processing.

**Usage:** Use this element in conjunction with the *fcm\_num\_rqb* configuration parameter to determine the maximum request block utilization. If *RB\_free\_bottom* is low, you should increase *fcm\_num\_rqb* to ensure that operations do not run out of request blocks. If *RB\_free\_bottom* is high, you can decrease *fcm\_num\_rqb* to conserve system resources.

## Number of Nodes

Snapshot Information Level	API Structure(s)	Monitor Switch
Database Manager	sqlm_fcm	Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	number_nodes information	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• None</li></ul>	

**Description:** The number of nodes in the current configuration.

**Usage:** Use this element to determine the number of *sqlm\_fcm\_node* structures that will be returned.

## Connection Status

Snapshot Information Level	API Structure(s)	Monitor Switch
Database Manager	sqlm_fcm_node	Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	connection_status information	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Total FCM Buffers Sent” on page 90</li><li>• “Total FCM Buffers Received” on page 90</li></ul>	

**Description:** This element indicates the status of the communication connection status between the node issuing the GET SNAPSHOT command and other nodes listed in the *db2nodes.cfg* file.

**Usage:** The connection values are :

SQLM_FCM_CONNECT_INACTIVE	No current connection
SQLM_FCM_CONNECT_ACTIVE	Connection is active
SQLM_FCM_CONNECT_CONGESTED	Connection is congested

Two nodes can be active, but the communication connection between them will remain inactive, unless there is some communication between those nodes.

## Total FCM Buffers Sent

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_fcm_node	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	total_buffers_sent counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Connection Status” on page 89</li><li>• “Total FCM Buffers Received” on page 90</li></ul>	

**Description:** The total number of FCM buffers that have been sent from the node issuing the GET SNAPSHOT command to the node identified by the *node\_number* (see the *db2nodes.cfg* file).

**Usage:** You can use this element to measure the level of traffic between the current node and the remote node. If the total number of FCM buffers sent to this node is high, you may want to redistribute the database, or move tables to reduce the inter-node traffic.

## Total FCM Buffers Received

<b>Snapshot Information Level</b> Database Manager	<b>API Structure(s)</b> sqlm_fcm_node	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	total_buffers_rcvd counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Connection Status” on page 89</li><li>• “Total FCM Buffers Sent” on page 90</li></ul>	

**Description:** The total number of FCM buffers received by the node issuing the GET SNAPSHOT command from the node identified by the *node\_number* (see the *db2nodes.cfg* file).

**Usage:** You can use this element to measure the level of traffic between the current node and the remote node. If the total number of FCM buffers received from this node is high, you may want to redistribute the database, or move tables to reduce the inter-node traffic.

---

## Database Configuration

The following elements provide information particularly helpful for database configuration.

## Buffer Pool Activity

The database server reads and updates all data from a buffer pool. Data is copied from disk to a buffer pool as it is required by applications.

Pages are placed in a buffer pool:

- by the agent. This is synchronous I/O.
- by the I/O servers (prefetchers). This is asynchronous I/O.

Pages are written to disk from a buffer pool:

- by the agent, synchronously
- by page cleaners, asynchronously

If the server needs to read a page of data, and that page is already in the buffer pool, then the ability to access that page is much faster than if the page had to be read from disk. It is desirable to **hit** as many pages as possible in the buffer pool. Avoiding disk I/O is the main issue when trying to improve the performance of your server. And so, proper configuration of the buffer pools are probably the most important consideration for performance tuning.

### Buffer Pool Hit Ratio

The buffer pool hit ratio indicates the percentage of time that the database manager did not need to load a page from disk in order to service a page request. That is, the page was already in the buffer pool. The greater the buffer pool hit ratio, the lower the frequency of disk I/O.

The buffer pool hit ratio can be calculated as follows:

$$(1 - ((\text{pool\_data\_p\_reads} + \text{pool\_index\_p\_reads}) / (\text{pool\_data\_l\_reads} + \text{pool\_index\_l\_reads}))) * 100\%$$

This calculation takes into account all of the pages (index and data) that are cached by the buffer pool.

For a large database, increasing the buffer pool size may have minimal effect on the buffer pool hit ratio. Its number of data pages may be so large, that the statistical chances of a hit are not improved increasing its size. But you might find that tuning the index buffer pool hit ratio achieves the desired result. This can be achieved using two methods:

1. Split the data and indices into two different buffer pools and tune them separately.
2. Use one buffer pool, but increase its size until the index hit ratio stops increasing.

The index buffer pool hit ratio can be calculated as follows:

$$(1 - ((\text{pool\_index\_p\_reads}) / (\text{pool\_index\_l\_reads}))) * 100\%$$

The first method is often more effective, but because it requires indices and data to reside in different tablespaces, it may not be an option for existing databases. It also requires tuning two buffer pools instead of one, which can be a more difficult task, particularly when memory is constrained.

## Prefetchers

You should also consider the impact that prefetchers may be having on the hit ratio. Prefetchers read data pages into the buffer pool anticipating their need by an application (asynchronously). In most situations, these pages are read just before they are needed (the desired case). However, prefetchers can cause unnecessary I/O by reading pages into the buffer pool that will not be used. For example, an application starts reading through a table. This is detected and prefetching starts, but the application fills an application buffer and stops reading. Meanwhile, prefetching has been done for a number of additional pages. I/O has occurred for pages that will not be used and the buffer pool is partially taken up with those pages.

## Page Cleaners

Page cleaners monitor the buffer pool and asynchronously write pages to disk. Their goals are:

- Ensure that agents will always find free pages in the buffer pool. If an agent does not find free pages in the buffer pool, it must clean them itself, and the associated application will have a poorer response.
- Speed database recovery, if a system crash occurs. The more pages that have been written to disk, the smaller the number of log file records that must be processed to recover the database.

Although dirty pages are written out to disk, the pages are not removed from the buffer pool right away, unless the space is needed to read in new pages.

**Note:** Buffer pool information is typically gathered at a table space level, but the facilities of the database system monitor can roll this information up to the buffer pool and database levels. Depending on your type of analysis, you may need to examine this data at any or all of these levels.

The following elements provide information about buffer pool activity. For an overview how the database manager uses buffer pools, see the *Administration Guide*.

- “Buffer Pool Data Logical Reads” on page 93
- “Buffer Pool Data Physical Reads” on page 94
- “Buffer Pool Data Writes” on page 95
- “Buffer Pool Index Logical Reads” on page 96
- “Buffer Pool Index Physical Reads” on page 97
- “Buffer Pool Index Writes” on page 98
- “Total Buffer Pool Physical Read Time” on page 99
- “Total Buffer Pool Physical Write Time” on page 100
- “Database Files Closed” on page 100
- “Buffer Pool Asynchronous Data Reads” on page 101
- “Buffer Pool Asynchronous Data Writes” on page 102
- “Buffer Pool Asynchronous Index Writes” on page 103
- “Buffer Pool Asynchronous Index Reads” on page 104
- “Buffer Pool Asynchronous Read Time” on page 105
- “Buffer Pool Asynchronous Write Time” on page 106
- “Buffer Pool Asynchronous Read Requests” on page 107
- “Buffer Pool Log Space Cleaners Triggered” on page 107

- “Buffer Pool Victim Page Cleaners Triggered” on page 108
- “Buffer Pool Threshold Cleaners Triggered” on page 109
- “Buffer Pool Information” on page 109
- “Time Waited for Prefetch” on page 110

## Buffer Pool Data Logical Reads

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Buffer Pool
Table Space	sqlm_tablespace	Buffer Pool
	sqlm_bp_info	Buffer Pool
Application	sqlm_appl	Buffer Pool
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Table Space	sqlm_tablespace_event	
Connection	sqlm_conn_event	
API Element Name	pool_data_l_reads	
Element Type	counter	
Related Information	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “When Counters are Initialized” on page 20</li> <li>• “Database Activation Timestamp” on page 41</li> <li>• “Connection Request Start Timestamp” on page 62</li> <li>• “Buffer Pool Data Physical Reads” on page 94</li> <li>• “Buffer Pool Data Writes” on page 95</li> <li>• “Buffer Pool Index Logical Reads” on page 96</li> <li>• “Buffer Pool Index Physical Reads” on page 97</li> </ul>	

**Description:** Indicates the number of logical read requests for data pages that have gone through the buffer pool.

**Usage:** This count includes accesses to data that is:

- Already in the buffer pool when the database manager needs to process the page
- Read into the buffer pool before the database manager can process the page.

In conjunction with *Buffer Pool Data Physical Reads*, you can calculate the data page hit ratio for the buffer pool using the following formula:

$$1 - (\text{buffer pool data physical reads} / \text{buffer pool data logical reads})$$

In conjunction with *Buffer Pool Data Physical Reads*, *Buffer Pool Index Physical Reads*, and *Buffer Pool Index Logical Reads*, you can calculate the overall buffer pool hit ratio using the following formula:

$$1 - ((\text{buffer pool data physical reads} + \text{buffer pool index physical reads}) / (\text{buffer pool data logical reads} + \text{buffer pool index logical reads}))$$

Increasing buffer pool size will generally improve the hit ratio, but you will reach a point of diminishing return. Ideally, if you could a buffer pool large enough to store your entire database, then once the system is up and running you would get a hit ratio of 100%.

However, this is realistic in most cases. The significance of the hit ratio really depends on the size of your data, and the way it is accessed. A very large database where data is accessed evenly would have a poor hit ratio. There is little you can do with very large tables. In such a case, you would focus your attention on smaller, frequently accessed tables, and on the indices. Perhaps, assigning them to individual buffer pools, for which you can aim for higher hit ratios.

## Buffer Pool Data Physical Reads

<b>Snapshot Information Level</b>	<b>API Structure(s)</b>	<b>Monitor Switch</b>
Database	sqlm_dbase	Buffer Pool
Table Space	sqlm_tablespace	Buffer Pool
	sqlm_bp_info	Buffer Pool
Application	sqlm_appl	Buffer Pool
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Table Space	sqlm_tablespace_event	
Connection	sqlm_conn_event	
<b>API Element Name</b>	pool_data_p_reads	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “When Counters are Initialized” on page 20</li> <li>• “Database Activation Timestamp” on page 41</li> <li>• “Connection Request Start Timestamp” on page 62</li> <li>• “Buffer Pool Data Logical Reads” on page 93</li> <li>• “Buffer Pool Data Writes” on page 95</li> <li>• “Buffer Pool Index Logical Reads” on page 96</li> <li>• “Buffer Pool Index Physical Reads” on page 97</li> <li>• “Buffer Pool Asynchronous Data Reads” on page 101</li> </ul>	

**Description:** The number of read requests that required I/O to get data pages into the buffer pool.

**Usage:** See Buffer Pool Data Logical Reads and Buffer Pool Asynchronous Data Reads for information about how to use this element.



## Buffer Pool Data Writes

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Buffer Pool
Table Space	sqlm_tablespace	Buffer Pool
Application	sqlm_bp_info sqlm_appl	Buffer Pool Buffer Pool
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Table Space	sqlm_tablespace_event	
Connection	sqlm_conn_event	
API Element Name	pool_data_writes	
Element Type	counter	
Related Information	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “When Counters are Initialized” on page 20</li><li>• “Database Activation Timestamp” on page 41</li><li>• “Connection Request Start Timestamp” on page 62</li><li>• “Buffer Pool Data Logical Reads” on page 93</li><li>• “Buffer Pool Data Physical Reads” on page 94</li><li>• “Total Buffer Pool Physical Write Time” on page 100</li><li>• “Buffer Pool Asynchronous Data Writes” on page 102</li></ul>	

**Description:** Indicates the number of times a buffer pool data page was physically written to disk.

**Usage:** If a buffer pool data page is written to disk for a high percentage of the Buffer Pool Data Physical Reads, you may be able to improve performance by increasing the number of buffer pool pages available for the database.

A buffer pool data page is written to disk for the following reasons:

- To free a page in the buffer pool so another page can be read
- To flush the buffer pool.

The system does not always write a page to make room for a new one. If the page has not been updated, it can simply be replaced. This replacement is not counted for this element.

The data page can be written by an asynchronous page-cleaner agent before the buffer pool space is required. These asynchronous page writes are included in the value of this element in addition to synchronous page writes (see Buffer Pool Asynchronous Data Writes).

When calculating this percentage, disregard the number of physical reads required to initially fill the buffer pool. To determine the number of pages written:

1. Run your application (to load the buffer)
2. Note the value of this element
3. Run your application again

4. Subtract the value recorded in step 2 from the new value of this element.

In order to prevent the buffer pool from being deallocated between the runnings of your application, you should either;

- activate the database with the `ACTIVATE DATABASE` command
- have an idle application connected to the database.

If all applications are updating the database, increasing the size of the buffer pool may not have much impact on performance since most of the buffer pool pages contain updated data, which must be written to disk. However, if the updated pages can be used by other units of work before being written out, the buffer pool can save a write and a read, which will improve your performance.

See the *Administration Guide* for more information about buffer pool size.

## Buffer Pool Index Logical Reads

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Buffer Pool
Table Space	sqlm_tablespace	Buffer Pool
Application	sqlm_bp_info sqlm_appl	Buffer Pool Buffer Pool
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Table Space	sqlm_tablespace_event	
Connection	sqlm_conn_event	
API Element Name	pool_index_l_reads	
Element Type	counter	
Related Information	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “When Counters are Initialized” on page 20</li> <li>• “Database Activation Timestamp” on page 41</li> <li>• “Connection Request Start Timestamp” on page 62</li> <li>• “Buffer Pool Index Physical Reads” on page 97</li> <li>• “Buffer Pool Index Writes” on page 98</li> <li>• “Buffer Pool Data Physical Reads” on page 94</li> <li>• “Buffer Pool Data Writes” on page 95</li> <li>• “Buffer Pool Data Logical Reads” on page 93</li> </ul>	

**Description:** Indicates the number of logical read requests for index pages that have gone through the buffer pool.

**Usage:** This count includes accesses to index pages that are:

- Already in the buffer pool when the database manager needs to process the page
- Read into the buffer pool before the database manager can process the page.

In conjunction with Buffer Pool Index Physical Reads, you can calculate the index page hit ratio for the buffer pool using one of the following:

$$1 - (\text{buffer pool index physical reads} / \text{buffer pool index logical reads})$$

To calculate the overall buffer pool hit ratio, see Buffer Pool Data Logical Reads.

If the hit ratio is low, increasing the number of buffer pool pages may improve performance. See the *Administration Guide* for more information about buffer pool size.

## Buffer Pool Index Physical Reads

<b>Snapshot Information Level</b>	<b>API Structure(s)</b>	<b>Monitor Switch</b>
Database	sqlm_dbase	Buffer Pool
Table Space	sqlm_tablespace	Buffer Pool
	sqlm_bp_info	Buffer Pool
Application	sqlm_appl	Buffer Pool
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Table Space	sqlm_tablespace_event	
Connection	sqlm_conn_event	
<b>API Element Name</b>	pool_index_p_reads	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “When Counters are Initialized” on page 20</li> <li>• “Database Activation Timestamp” on page 41</li> <li>• “Connection Request Start Timestamp” on page 62</li> <li>• “Buffer Pool Index Logical Reads” on page 96</li> <li>• “Buffer Pool Index Writes” on page 98</li> <li>• “Buffer Pool Data Logical Reads” on page 93</li> <li>• “Buffer Pool Data Physical Reads” on page 94</li> </ul>	

**Description:** Indicates the number of physical read requests to get index pages into the buffer pool.

**Usage:** See Buffer Pool Index Logical Reads for information about how to use this element.

## Buffer Pool Index Writes

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Buffer Pool
Table Space	sqlm_tablespace	Buffer Pool
	sqlm_bp_info	Buffer Pool
Application	sqlm_appl	Buffer Pool
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Table Space	sqlm_tablespace_event	
Connection	sqlm_conn_event	
API Element Name	pool_index_writes	
Element Type	counter	
Related Information	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “When Counters are Initialized” on page 20</li><li>• “Database Activation Timestamp” on page 41</li><li>• “Connection Request Start Timestamp” on page 62</li><li>• “Buffer Pool Index Logical Reads” on page 96</li><li>• “Buffer Pool Index Physical Reads” on page 97</li><li>• “Buffer Pool Asynchronous Index Writes” on page 103</li></ul>	

**Description:** Indicates the number of times a buffer pool index page was physically written to disk.

**Usage:** Like a data page, a buffer pool index page is written to disk for the following reasons:

- To free a page in the buffer pool so another page can be read
- To flush the buffer pool.

The system does not always write a page to make room for a new one. If the page has not been updated, it can simply be replaced. This replacement is not counted for this element.

The index page can be written by an asynchronous page-cleaner agent before the buffer pool space is required. These asynchronous index page writes are included in the value of this element in addition to synchronous index page writes (see Buffer Pool Asynchronous Index Writes).

If a buffer pool index page is written to disk for a high percentage of the *Buffer Pool Index Physical Reads*, you may be able to improve performance by increasing the number of buffer pool pages available for the database.

When calculating this percentage, disregard the number of physical reads required to initially fill the buffer pool. To determine the number of pages written:

1. Run your application (to load the buffer)
2. Note the value of this element
3. Run your application again

4. Subtract the value recorded in step 2 from the new value of this element.

In order to prevent the buffer pool from being deallocated between the runnings of your application, you should either:

- activate the database with the `ACTIVATE DATABASE` command
- have an idle application connected to the database.

If all applications are updating the database, increasing the size of the buffer pool may not have much impact on performance, since most of the pages contain updated data which must be written to disk.

See the *Administration Guide* for more information about buffer pool size.

### Total Buffer Pool Physical Read Time

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Buffer Pool
Table Space	sqlm_tablespace	Buffer Pool
Application	sqlm_bp_info sqlm_appl	Buffer Pool Buffer Pool
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Table Space	sqlm_tablespace_event	
Connection	sqlm_conn_event	
API Element Name	pool_read_time	
Element Type	counter	
Related Information	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “When Counters are Initialized” on page 20</li><li>• “Buffer Pool Data Physical Reads” on page 94</li><li>• “Buffer Pool Index Physical Reads” on page 97</li><li>• “Database Activation Timestamp” on page 41</li><li>• “Connection Request Start Timestamp” on page 62</li><li>• “Buffer Pool Asynchronous Read Time” on page 105</li></ul>	

**Description:** Provides the total amount of elapsed time spent processing read requests that caused data or index pages to be physically read from disk to buffer pool.

**Usage:** You can use this element with Buffer Pool Data Physical Reads and Buffer Pool Index Physical Reads to calculate the average page-read time. This average is important since it may indicate the presence of an I/O wait, which in turn may indicate that you should be moving data to a different device.

At the database and table space levels, this element includes the value of Buffer Pool Asynchronous Read Time.

## Total Buffer Pool Physical Write Time

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Buffer Pool
Table Space	sqlm_tablespace	Buffer Poo
	sqlm_bp_info	Buffer Pool
Application	sqlm_appl	Buffer Pool
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Table Space	sqlm_tablespace_event	
Connection	sqlm_conn_event	
API Element Name	pool_write_time	
Element Type	counter	
Related Information	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “When Counters are Initialized” on page 20</li> <li>• “Buffer Pool Data Writes” on page 95</li> <li>• “Buffer Pool Index Writes” on page 98</li> <li>• “Database Activation Timestamp” on page 41</li> <li>• “Connection Request Start Timestamp” on page 62</li> </ul>	

**Description:** Provides the total amount of time spent physically writing data or index pages from the buffer pool to disk.

**Usage:** You can use this element with Buffer Pool Data Writes and Buffer Pool Index Writes to calculate the average page-write time. This average is important since it may indicate the presence of an I/O wait, which in turn may indicate that you should be moving data to a different device.

At the database and table space levels, this element includes the value of Buffer Pool Asynchronous Write Time.

## Database Files Closed

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Buffer Pool
Table Space	sqlm_tablespace	Buffer Pool
	sqlm_bp_info	Buffer Pool
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Table Space	sqlm_tablespace_event	
API Element Name	files_closed	
Element Type	counter	
Related Information	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “When Counters are Initialized” on page 20</li> </ul>	

**Description:** The total number of database files closed.

**Usage:** The database manager opens files for reading and writing into and out of the buffer pool. The maximum number of database files open by an application at any time is controlled by the *maxfilop* configuration parameter. If the maximum is reached, one file will be closed before the new file is opened. Note that the actual number of files opened may not equal the number of files closed.

You can use this element to help you determine the best value for the *maxfilop* configuration parameter (see the *Administration Guide* for more information).

## Buffer Pool Asynchronous Data Reads

<b>Snapshot Information Level</b>	<b>API Structure(s)</b>	<b>Monitor Switch</b>
Database	sqlm_dbase	Buffer Pool
Table Space	sqlm_tablespace	Buffer Pool
	sqlm_bp_info	Buffer Pool
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Table Space	sqlm_tablespace_event	
<b>API Element Name</b>	pool_async_data_reads	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “When Counters are Initialized” on page 20</li> <li>• “Buffer Pool Asynchronous Read Time” on page 105</li> <li>• “Buffer Pool Data Physical Reads” on page 94</li> <li>• “Buffer Pool Asynchronous Read Requests” on page 107</li> <li>• “Direct Reads From Database” on page 115</li> </ul>	

**Description:** The number of pages read asynchronously into the buffer pool.

**Usage:** You can use this element with Buffer Pool Data Physical Reads to calculate the number of physical reads that were performed synchronously (that is, physical data page reads that were performed by database manager agents). Use the following formula:

$$\text{buffer pool data physical reads} - \text{buffer pool asynchronous data reads}$$

By comparing the ratio of asynchronous to synchronous reads, you can gain insight into how well the prefetchers are working. This element can be helpful when you are tuning the *num\_ioservers* configuration parameter (see the *Administration Guide*).

Asynchronous reads are performed by database manager prefetchers. For information about these prefetchers, see the *Administration Guide*.

## Buffer Pool Asynchronous Data Writes

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Buffer Pool
Table Space	sqlm_tablespace	Buffer Pool
	sqlm_bp_info	Buffer Pool
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Table Space	sqlm_tablespace_event	
API Element Name	pool_async_data_writes	
Element Type	counter	
Related Information	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “When Counters are Initialized” on page 20</li> <li>• “Buffer Pool Asynchronous Index Writes” on page 103</li> <li>• “Buffer Pool Data Writes” on page 95</li> <li>• “Buffer Pool Asynchronous Write Time” on page 106</li> <li>• “Buffer Pool Log Space Cleaners Triggered” on page 107</li> <li>• “Buffer Pool Victim Page Cleaners Triggered” on page 108</li> <li>• “Buffer Pool Threshold Cleaners Triggered” on page 109</li> <li>• “Direct Writes to Database” on page 116</li> </ul>	

**Description:** The number of times a buffer pool data page was physically written to disk by either an asynchronous page cleaner, or a prefetcher. A prefetcher may have written dirty pages to disk to make space for the pages being prefetched.

**Usage:** You can use this element with Buffer Pool Data Writes to calculate the number of physical write requests that were performed synchronously (that is, physical data page writes that were performed by database manager agents). Use the following formula:

$$\text{buffer pool data writes} - \text{buffer pool asynchronous data writes}$$

By comparing the ratio of asynchronous to synchronous writes, you can gain insight into how well the buffer pool page cleaners are performing. This ratio can be helpful when you are tuning the *num\_iocleaners* configuration parameter.

For more information about asynchronous page cleaners, see the *Administration Guide*.



## Buffer Pool Asynchronous Index Writes

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Buffer Pool
Table Space	sqlm_tablespace	Buffer Pool
	sqlm_bp_info	Buffer Pool
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Table Space	sqlm_tablespace_event	
API Element Name	pool_async_index_writes	
Element Type	counter	
Related Information	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “When Counters are Initialized” on page 20</li><li>• “Buffer Pool Asynchronous Data Writes” on page 102</li><li>• “Buffer Pool Asynchronous Index Reads” on page 104</li><li>• “Buffer Pool Index Writes” on page 98</li><li>• “Buffer Pool Asynchronous Write Time” on page 106</li><li>• “Buffer Pool Log Space Cleaners Triggered” on page 107</li><li>• “Buffer Pool Victim Page Cleaners Triggered” on page 108</li><li>• “Buffer Pool Threshold Cleaners Triggered” on page 109</li><li>• “Direct Writes to Database” on page 116</li></ul>	

**Description:** The number of times a buffer pool index page was physically written to disk by either an asynchronous page cleaner, or a prefetcher. A prefetcher may have written dirty pages to disk to make space for the pages being prefetched.

**Usage:** You can use this element with Buffer Pool Index Writes to calculate the number of physical index write requests that were performed synchronously. That is, physical index page writes that were performed by database manager agents. Use the following formula:

$$\text{buffer pool index writes} - \text{buffer pool asynchronous index writes}$$

By comparing the ratio of asynchronous to synchronous writes, you can gain insight into how well the buffer pool page cleaners are performing. This ratio can be helpful when you are tuning the *num\_iocleaners* configuration parameter.

For more information about asynchronous page cleaners, see the *Administration Guide*.

## Buffer Pool Asynchronous Index Reads

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Buffer Pool
Table Space	sqlm_tablespace	Buffer Pool
	sqlm_bp_info	Buffer Pool
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Table Space	sqlm_tablespace_event	
API Element Name	pool_async_index_reads	
Element Type	counter	
Related Information	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “When Counters are Initialized” on page 20</li><li>• “Buffer Pool Asynchronous Data Writes” on page 102</li><li>• “Buffer Pool Asynchronous Index Writes” on page 103</li><li>• “Buffer Pool Index Physical Reads” on page 97</li><li>• “Buffer Pool Asynchronous Read Time” on page 105</li><li>• “Buffer Pool Log Space Cleaners Triggered” on page 107</li><li>• “Buffer Pool Victim Page Cleaners Triggered” on page 108</li><li>• “Buffer Pool Threshold Cleaners Triggered” on page 109</li><li>• “Direct Reads From Database” on page 115</li></ul>	

**Description:** The number of index pages read asynchronously into the buffer pool by a prefetcher.

**Usage:** You can use this element with Buffer Pool Index Physical Reads to calculate the number of physical reads that were performed synchronously (that is, physical index page reads that were performed by database manager agents). Use the following formula:

```
buffer pool index physical reads - buffer pool asynchronous index reads
```

By comparing the ratio of asynchronous to synchronous reads, you can gain insight into how well the prefetchers are working. This element can be helpful when you are tuning the *num\_ioservers* configuration parameter (see the *Administration Guide*).

Asynchronous reads are performed by database manager prefetchers. For information about these prefetchers, see the *Administration Guide*.

## Buffer Pool Asynchronous Read Time

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Buffer Pool
Table Space	sqlm_tablespace	Buffer Pool
	sqlm_bp_info	Buffer Pool
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Table Space	sqlm_tablespace_event	
<b>API Element Name</b>	pool_async_read_time	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “When Counters are Initialized” on page 20</li><li>• “Buffer Pool Asynchronous Data Reads” on page 101</li><li>• “Total Buffer Pool Physical Read Time” on page 99</li><li>• “Buffer Pool Asynchronous Read Requests” on page 107</li><li>• “Direct Read Time” on page 118</li></ul>	

**Description:** The total elapsed time spent reading by database manager prefetchers.

**Usage:** You can use this element to calculate the elapsed time for synchronous reading, using the following formula:

$$\text{total buffer pool physical read time} - \text{buffer pool asynchronous read time}$$

You can also use this element to calculate the average asynchronous read time using the following formula:

$$\text{buffer pool asynchronous read time} / \text{buffer pool asynchronous data reads}$$

These calculations can be used to understand the I/O work being performed.

## Buffer Pool Asynchronous Write Time

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Buffer Pool
Table Space	sqlm_tablespace	Buffer Pool
	sqlm_bp_info	Buffer Pool
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Table Space	sqlm_tablespace_event	
API Element Name	pool_async_write_time	
Element Type	counter	
Related Information	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “When Counters are Initialized” on page 20</li><li>• “Buffer Pool Asynchronous Data Writes” on page 102</li><li>• “Buffer Pool Asynchronous Index Writes” on page 103</li><li>• “Total Buffer Pool Physical Write Time” on page 100</li><li>• “Buffer Pool Asynchronous Read Requests” on page 107</li><li>• “Direct Write Time” on page 118</li></ul>	

**Description:** The total elapsed time spent writing data or index pages from the buffer pool to disk by database manager page cleaners.

**Usage:** To calculate the elapsed time spent writing pages synchronously, use the following formula:

total buffer pool physical write time - buffer pool asynchronous write time

You can also use this element to calculate the average asynchronous read time using the following formula:

buffer pool asynchronous write time  
/ (buffer pool asynchronous data writes + buffer pool asynchronous index writes)

These calculations can be used to understand the I/O work being performed.

## Buffer Pool Asynchronous Read Requests

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Buffer Pool
Table Space	sqlm_tablespace	Buffer Pool
	sqlm_bp_info	Buffer Pool
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Table Space	sqlm_tablespace_event	
<b>API Element Name</b>	pool_async_data_read_reqs	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “When Counters are Initialized” on page 20</li><li>• “Buffer Pool Asynchronous Data Reads” on page 101</li></ul>	

**Description:** The number of asynchronous read requests.

**Usage:** To calculate the average number of data pages read per asynchronous request, use the following formula:

buffer pool asynchronous data reads / buffer pool asynchronous read requests

This average can help you determine the amount of asynchronous I/O done in each interaction with the prefetcher.

## Buffer Pool Log Space Cleaners Triggered

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Buffer Pool
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
<b>API Element Name</b>	pool_lsn_gap_clns	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “When Counters are Initialized” on page 20</li><li>• “Buffer Pool Victim Page Cleaners Triggered” on page 108</li><li>• “Buffer Pool Threshold Cleaners Triggered” on page 109</li></ul>	

**Description:** The number of times a page cleaner was invoked because the logging space used had reached a predefined criterion for the database.

**Usage:** This element can be used to help evaluate whether you have enough space for logging, and whether you need more log files or larger log files.

The page cleaning criterion is determined by the setting for the *softmax* configuration parameter. Page cleaners are triggered if the oldest page in the buffer pool contains an update described by a log record that is older than the current log position by the criterion value. See the *Administration Guide* for more information.

## Buffer Pool Victim Page Cleaners Triggered

<b>Snapshot Information Level</b> Database	<b>API Structure(s)</b> sqlm_dbase	<b>Monitor Switch</b> Buffer Pool
<b>Resettable</b>	Yes	
<b>Event Type</b> Database	<b>Event Record(s)</b> sqlm_db_event	
<b>API Element Name</b> <b>Element Type</b>	pool_drty_pg_steal_clns counter	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “When Counters are Initialized” on page 20</li> <li>• “Buffer Pool Log Space Cleaners Triggered” on page 107</li> <li>• “Buffer Pool Threshold Cleaners Triggered” on page 109</li> </ul>	

**Description:** The number of times a page cleaner was invoked because a synchronous write was needed during the victim buffer replacement for the database.

**Usage:** Using the following formula, you may calculate what percentage of all cleaner invocations are represented by this element:

$$\frac{\text{buffer pool victim page cleaners triggered}}{(\text{buffer pool victim page cleaners triggered} + \text{buffer pool threshold cleaners triggered} + \text{buffer pool log space cleaners triggered})}$$

If this ratio is low, it may indicate that you have defined too many page cleaners. If your *chnngpgs\_thresh* is set too low, you may be writing out pages that you will dirty later. Aggressive cleaning defeats one purpose of the buffer pool, that is to defer writing to the last possible moment.

If this ratio is high, it may indicate that you have too few page cleaners defined. Too few page cleaners will increase recovery time after failures (see the *Administration Guide*).

**Note:** Although dirty pages are written out to disk, the pages are not removed from the buffer pool right away, unless the space is needed to read in new pages.

## Buffer Pool Threshold Cleaners Triggered

<b>Snapshot Information Level</b> Database	<b>API Structure(s)</b> sqlm_dbase	<b>Monitor Switch</b> Buffer Pool
<b>Resettable</b>	Yes	
<b>Event Type</b> Database	<b>Event Record(s)</b> sqlm_db_event	
<b>API Element Name</b> <b>Element Type</b>	pool_drty_pg_thrsh_clns counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “When Counters are Initialized” on page 20</li><li>• “Buffer Pool Log Space Cleaners Triggered” on page 107</li></ul>	

**Description:** The number of times a page cleaner was invoked because a buffer pool had reached the dirty page threshold criterion for the database.

**Usage:** The threshold is set by the *chnpggs\_thresh* configuration parameter. It is a percentage applied to the buffer pool size. When the number of dirty pages in the pool exceeds this value, the cleaners are triggered.

If this value is set too low, pages might be written out too early, requiring them to be read back in. If set too high, then too many pages may accumulate, requiring users to write out pages synchronously. See the *Administration Guide* for more information.

## Buffer Pool Information

<b>Snapshot Information Level</b> Table Space	<b>API Structure(s)</b> sqlm_bufferpool	<b>Monitor Switch</b> Buffer Pool
<b>Resettable</b>	No	
<b>Event Type</b> Table Space	<b>Event Record(s)</b> sqlm_bufferpool_event	
<b>API Element Name</b> <b>Element Type</b>	bp_info information	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li></ul>	

**Description:** Data management counters for a buffer pool.

**Usage:** Activity performed for a bufferpool.

## Time Waited for Prefetch

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Buffer Pool
Application	sqlm_appl	Buffer Pool
<b>Resettable</b>	No	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
<b>API Element Name</b>	prefetch_wait_time	
<b>Element Type</b>	counter	
<b>Related Information</b>	• None	

**Description:** The time an application spent waiting for an I/O server (prefetcher) to finish loading pages into the buffer pool.

**Usage:** This element can be used to experiment with changing the number of I/O servers, and I/O server sizes.

## Extended Storage

Extended storage provides a secondary level of storage for buffer pools. This allows a user to access memory beyond the maximum allowed for each process. Extended storage consists of segments that will be allocated in addition to the buffer pools. The extended storage will assign pages to segments that are attached or detached, as needed. The number and size of segments are configurable. Attachment is allowed to only one segment at a given time.

There is one extended storage for all buffer pools, and each buffer pool can be configured to use it or not. See the *Administration Guide* for more information.

Extended storage should only be used on systems with very large amount of real memory. These are systems that have more memory than can be attached to by a single process.

**Using Extended Storage Counters:** If you have extended storage set on for a buffer pool, all pages removed from the buffer pool will be written to extended storage. Each of these writes has a cost associated with it. Some of these pages may never be required or they may be forced out of extended storage before they are ever read back into the buffer pool.

You can calculate the extended storage read/write ratio as follows:

$$\frac{(\text{data} + \text{index copied from extended storage})}{(\text{data} + \text{index copied to extended storage})}$$

Where the numerator in this equation is pages from extended storage to buffer pool and the denominator is pages from buffer pool to extended storage.



The top portion of this equation represents a performance saving. When a page is transferred from extended storage to buffer pool, you save a system I/O call. However, you still incur the cost of attaching to the extended memory segment, copying the page, and detaching from the segment. The bottom part represents the cost of transferring a page to extended storage, that is, attaching to the segment, copying the page, and detaching.

The higher the ratio, the more likely you are to benefit from extended storage. In general, extended storage is particularly useful if I/O activity is very high on your system.

There is a crossover point where the cost of copying pages to be removed from the buffer pool to extended storage equals the savings from reading pages from extended storage, instead of having to read them from disk. This crossover point is affected by:

- cost of an I/O on your system
- cost of copying data in memory and accessing shared memory segments

It is difficult to establish an exact crossover point. To establish a baseline, you must experiment by enabling extended storage for different buffer pools, and determine whether it improves your overall database performance. This can be measured by using application benchmarks. For instance, you may want to monitor transaction rates and execution time. See the *Administration Guide* for information on benchmarking.

Once you have established that extended storage is beneficial for some buffer pools. You want to measure the read/write ratio to obtain a baseline. This ratio is most important during database creation and initial setup. After that, you want to monitor this ratio to ensure that it is not deviating from the initial baseline.

The following elements provide information about buffer pools and extended storage. For more information on how the database manager uses extended storage, see the *Administration Guide*.

- “Buffer Pool Data Pages to Extended Storage” on page 112
- “Buffer Pool Index Pages to Extended Storage” on page 112
- “Buffer Pool Data Pages from Extended Storage” on page 113
- “Buffer Pool Index Pages from Extended Storage” on page 114

## Buffer Pool Data Pages to Extended Storage

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Buffer Pool
Table Space	sqlm_tablespace	Buffer Pool
	sqlm_bp_info	Buffer Pool
Application	sqlm_appl	Buffer Pool
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
Table Space	sqlm_tablespace_event	
API Element Name	pool_data_to_estore	
Element Type	counter	
Related Information	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “Buffer Pool Index Pages to Extended Storage” on page 112</li> <li>• “Buffer Pool Data Pages from Extended Storage” on page 113</li> <li>• “Buffer Pool Index Pages from Extended Storage” on page 114</li> </ul>	

**Description:** Number of buffer pool data pages copied to extended storage.

**Usage:** Pages are copied from the buffer pool to extended storage, when they are selected as victim pages. This copying is required to make space for new pages in the buffer pool.

## Buffer Pool Index Pages to Extended Storage

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Buffer Pool
Table Space	sqlm_tablespace	Buffer Pool
	sqlm_bp_info	Buffer Pool
Application	sqlm_appl	Buffer Pool
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
Table Space	sqlm_tablespace_event	
API Element Name	pool_index_to_estore	
Element Type	counter	
Related Information	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “Buffer Pool Data Pages to Extended Storage” on page 112</li> <li>• “Buffer Pool Data Pages from Extended Storage” on page 113</li> <li>• “Buffer Pool Index Pages from Extended Storage” on page 114</li> </ul>	

**Description:** Number of buffer pool index pages copied to extended storage.

**Usage:** Pages are copied from the buffer pool to extended storage, when they are selected as victim pages. This copying is required to make space for new pages in the buffer pool.

### Buffer Pool Data Pages from Extended Storage

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Buffer Pool
Table Space	sqlm_tablespace	Buffer Pool
	sqlm_bp_info	Buffer Pool
Application	sqlm_appl	Buffer Pool
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
Table Space	sqlm_tablespace_event	
API Element Name	pool_data_from_estore	
Element Type	counter	
Related Information	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “Buffer Pool Data Pages to Extended Storage” on page 112</li><li>• “Buffer Pool Index Pages to Extended Storage” on page 112</li><li>• “Buffer Pool Index Pages from Extended Storage” on page 114</li></ul>	

**Description:** Number of buffer pool data pages copied from extended storage.

**Usage:** Required pages are copied from extended storage to the buffer pool, if they are not in the buffer pool, but are in extended storage. This copying may incur the cost of connecting to the shared memory segment, but saves the cost of a disk read.

## Buffer Pool Index Pages from Extended Storage

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Buffer Pool
Table Space	sqlm_tablespace	Buffer Pool
	sqlm_bp_info	Buffer Pool
Application	sqlm_appl	Buffer Pool
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
Table Space	sqlm_tablespace_event	
API Element Name	pool_index_from_estore	
Element Type	counter	
Related Information	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “Buffer Pool Data Pages to Extended Storage” on page 112</li><li>• “Buffer Pool Index Pages to Extended Storage” on page 112</li><li>• “Buffer Pool Data Pages from Extended Storage” on page 113</li></ul>	

**Description:** Number of buffer pool index pages copied from extended storage.

**Usage:** Required index pages are copied from extended storage to the buffer pool, if they are not in the buffer pool, but are in extended storage. This copying may incur the cost of connecting to the shared memory segment, but saves the cost of a disk read.

## Non-buffered I/O Activity

The following elements provide information about I/O activity that does not use the buffer pool:

- “Direct Reads From Database” on page 115
- “Direct Writes to Database” on page 116
- “Direct Read Requests” on page 117
- “Direct Write Requests” on page 117
- “Direct Read Time” on page 118
- “Direct Write Time” on page 118

## Direct Reads From Database

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Buffer Pool
Table Space	sqlm_tablespace	Buffer Pool
Application	sqlm_bp_info sqlm_appl	Buffer Pool Buffer Pool
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
Table Space	sqlm_tablespace_event	
API Element Name	direct_reads	
Element Type	counter	
Related Information	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “When Counters are Initialized” on page 20</li><li>• “Direct Read Requests” on page 117</li><li>• “Direct Read Time” on page 118</li><li>• “Direct Writes to Database” on page 116</li></ul>	

**Description:** The number of read operations that do not use the buffer pool.

**Usage:** Use the following formula to calculate the average number of sectors that are read by a direct read:

$$\text{direct reads from database} / \text{direct read requests}$$

When using system monitors to track I/O, this data element helps you distinguish database I/O from non-database I/O on the device.

Direct reads are performed in units, the smallest being a 512-byte sector. They are used when:

- Reading LONG VARCHAR columns
- Reading LOB (large object) columns
- Performing a backup

## Direct Writes to Database

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Buffer Pool
Table Space	sqlm_tablespace	Buffer Pool
	sqlm_bp_info	Buffer Pool
Application	sqlm_appl	Buffer Pool
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
Table Space	sqlm_tablespace_event	
API Element Name	direct_writes	
Element Type	counter	
Related Information	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “When Counters are Initialized” on page 20</li><li>• “Direct Write Requests” on page 117</li><li>• “Direct Write Time” on page 118</li><li>• “Direct Reads From Database” on page 115</li></ul>	

**Description:** The number of write operations that do not use the buffer pool.

**Usage:** Use the following formula to calculate the average number of sectors that are written by a direct write.

$$\text{direct writes to database} / \text{direct write requests}$$

When using system monitors to track I/O, this data element helps you distinguish database I/O from non-database I/O on the device.

Direct writes are performed in units, the smallest being a 512-byte sector. They are used when:

- Writing LONG VARCHAR columns
- Writing LOB (large object) columns
- Performing a restore
- Performing a load.

## Direct Read Requests

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Buffer Pool
Table Space	sqlm_tablespace	Buffer Pool
Application	sqlm_bp_info sqlm_appl	Buffer Pool Buffer Pool
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
Table Space	sqlm_tablespace_event	
<b>API Element Name</b>	direct_read_reqs	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “When Counters are Initialized” on page 20</li> <li>• “Direct Reads From Database” on page 115</li> <li>• “Direct Read Time” on page 118</li> <li>• “Direct Write Requests” on page 117</li> </ul>	

**Description:** The number of requests to perform a direct read of one or more sectors of data.

**Usage:** Use the following formula to calculate the average number of sectors that are read by a direct read:

$$\text{direct reads from database} / \text{direct read requests}$$

## Direct Write Requests

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Buffer Pool
Table Space	sqlm_tablespace	Buffer Pool
Application	sqlm_bp_info sqlm_appl	Buffer Pool Buffer Pool
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
Table Space	sqlm_tablespace_event	
<b>API Element Name</b>	direct_write_reqs	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “When Counters are Initialized” on page 20</li> <li>• “Direct Writes to Database” on page 116</li> <li>• “Direct Write Time” on page 118</li> <li>• “Direct Read Requests” on page 117</li> </ul>	

**Description:** The number of requests to perform a direct write of one or more sectors of data.

**Usage:** Use the following formula to calculate the average number of sectors that are written by a direct write:

direct writes to database / direct write requests

## Direct Read Time

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Buffer Pool
Table Space	sqlm_tablespace	Buffer Pool
	sqlm_bp_info	Buffer Pool
Application	sqlm_appl	Buffer Pool
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
Table Space	sqlm_tablespace_event	
API Element Name	direct_read_time	
Element Type	counter	
Related Information	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “When Counters are Initialized” on page 20</li> <li>• “Direct Reads From Database” on page 115</li> <li>• “Direct Read Requests” on page 117</li> <li>• “Direct Write Time” on page 118</li> </ul>	

**Description:** The elapsed time (in milliseconds) required to perform the direct reads.

**Usage:** Use the following formula to calculate the average direct read time per sector:

direct read time / direct reads from database

A high average time may indicate an I/O conflict.

## Direct Write Time

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Buffer Pool
Table Space	sqlm_tablespace	Buffer Pool
	sqlm_bp_info	Buffer Pool
Application	sqlm_appl	Buffer Pool
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
Table Space	sqlm_tablespace_event	
API Element Name	direct_write_time	
Element Type	counter	



---

**Related Information**

- “Resetting Monitor Data” on page 21
  - “When Counters are Initialized” on page 20
  - “Direct Writes to Database” on page 116
  - “Direct Write Requests” on page 117
  - “Direct Read Time” on page 118
- 

**Description:** The elapsed time (in milliseconds) required to perform the direct writes.

**Usage:** Use the following formula to calculate the average direct write time per sector:

direct write time / direct writes to database

A high average time may indicate an I/O conflict.

## Catalog Cache

The catalog cache stores table descriptors for tables, views, and aliases. A descriptor stores information about a table, view, or alias in a condensed internal format. When a transaction references a table, it causes an insert of a table descriptor into the cache, so that subsequent transactions referencing that same table can use that descriptor and avoid reading from disk. (Transactions reference a table descriptor when compiling an SQL statement.)

The following database system monitor elements are used for catalog caches:

- “Catalog Cache Lookups”
- “Catalog Cache Inserts” on page 120
- “Catalog Cache Overflows” on page 121
- “Catalog Cache Heap Full” on page 122

## Catalog Cache Lookups

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
<b>API Element Name</b>	cat_cache_lookups	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Catalog Cache Inserts” on page 120</li><li>• “Catalog Cache Overflows” on page 121</li><li>• “Catalog Cache Heap Full” on page 122</li></ul>	

---

**Description:** The number of times that the catalog cache was referenced to obtain table descriptor information.

**Usage:** This element includes both successful and unsuccessful accesses to the catalog cache. The catalog cache is referenced whenever a table, view, or alias name is processed during the compilation of an SQL statement.

To calculate the catalog cache hit ratio use the following formula:

$$(1 - (\text{cat\_cache\_inserts} / \text{cat\_cache\_lookups}))$$

indicates how well the catalog cache is avoiding catalog accesses. If the ratio is high (more than 0.8), then the cache is performing well. A smaller ratio might suggest that the *catalogcache\_sz* should be increased. You should expect a large ratio immediately following the first connection to the database.

The execution of Data Definition Language (DDL) SQL statements involving a table, view, or alias will evict the table descriptor information for that object from the catalog cache causing it to be re-inserted on the next reference. Therefore, the heavy use of DDLs may also increase the ratio.

See the *Administration Guide* for more information on the Catalog Cache Size configuration parameter.

## Catalog Cache Inserts

<b>Snapshot Information Level</b>	<b>API Structure(s)</b>	<b>Monitor Switch</b>
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
<b>API Element Name</b>	cat_cache_inserts	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Catalog Cache Lookups” on page 119</li><li>• “Catalog Cache Overflows” on page 121</li><li>• “Catalog Cache Heap Full” on page 122</li><li>• “Data Definition Language (DDL) SQL Statements” on page 169</li></ul>	

**Description:** The number of times that the system tried to insert table descriptor information into the catalog cache.

**Usage:** Table descriptor information is usually inserted into the cache following a failed lookup to the catalog cache while processing a table, view, or alias reference in an SQL statement. The *catalog cache inserts* value includes attempts to insert table descriptor information that fail due to catalog cache overflow and heap full conditions.

See “Catalog Cache Lookups” on page 119 for more catalog cache information.

## Catalog Cache Overflows

<b>Snapshot Information Level</b>	<b>API Structure(s)</b>	<b>Monitor Switch</b>
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
<b>API Element Name</b>	cat_cache_overflows	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Catalog Cache Lookups” on page 119</li><li>• “Catalog Cache Inserts” on page 120</li></ul>	

**Description:** The number of times that an insert into the catalog cache failed due the catalog cache being full.

**Usage:** The catalog cache space is filled with table descriptor information.

The cache entries for transactions that compile SQL statements, either by issuing dynamic SQL statements or by binding a package, will not be eligible to be removed from the cache until that transaction has either been committed or rolled back. Catalog cache space is reclaimed by evicting table descriptor information for tables, views, or aliases that are not currently in use by any transaction. Once a transaction has experienced a catalog cache overflow, all subsequent attempts by the same transaction to insert table descriptor information into the catalog cache will also result in an overflow.

**Note:** A transaction involved in an overflow will proceed, but its descriptor information is not inserted into the cache.

If *catalog cache overflows* is large, the catalog cache may be too small for the workload. Enlarging the catalog cache may improve its performance. If the workload includes transactions which compile a large number of SQL statements referencing many tables, views, and aliases in a single unit of work, then compiling fewer SQL statements in a single transaction may improve the performance of the catalog cache. Or if it includes binding of packages containing many SQL statements referencing many tables, views or aliases, you can try splitting packages so that they include fewer SQL statements to improve performance.

## Catalog Cache Heap Full

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
<b>API Element Name</b>	cat_cache_heap_full	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Package Cache Inserts” on page 124</li><li>• “Data Definition Language (DDL) SQL Statements” on page 169</li><li>• “Dynamic SQL Statements Attempted” on page 165</li><li>• “Static SQL Statements Attempted” on page 164</li></ul>	

**Description:** The number of times that an insert into the catalog cache failed due to a heap-full condition in the database heap.

**Usage:** The catalog cache draws its storage dynamically from the database heap and even if the cache storage has not reached its limit, inserts into the catalog cache may fail due to a lack of space in the database heap.

If the catalog cache heap full count is not zero, then this insert failure condition can be corrected by increasing the database heap size or reducing the catalog cache size.

## Package Cache

The package and section information required for the execution of dynamic and static SQL statements are placed in the package cache as required. This information is required whenever a dynamic or static statement is being executed. The package cache exists at a database level. This means that agents with similar environments can share the benefits of another agent’s work. For static SQL statements, this can mean avoiding catalog access. For dynamic SQL statements, this can mean avoiding the cost of compilation.

The following database system monitor elements are used for package caches:

- “Package Cache Lookups” on page 123
- “Package Cache Inserts” on page 124
- “Section Lookups” on page 125
- “Section Inserts” on page 126

## Package Cache Lookups

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
<b>API Element Name</b>	pkg_cache_lookups	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Package Cache Inserts” on page 124</li><li>• “Section Lookups” on page 125</li><li>• “Section Inserts” on page 126</li><li>• “Static SQL Statements Attempted” on page 164</li><li>• “Dynamic SQL Statements Attempted” on page 165</li><li>• “Data Definition Language (DDL) SQL Statements” on page 169</li></ul>	

**Description:** The number of times that an application looked for a section or package in the package cache. At a database level, it indicates the overall number of references since the database was started, or monitor data was reset.

**Note:** This counter includes the cases where the section is already loaded in the cache and when the section has to be loaded into the cache.

**Usage:** To calculate the package cache hit ratio use the following formula:

$$1 - (\text{Package Cache Inserts} / \text{Package Cache Lookups})$$

The package cache hit ratio tells you whether or not the package cache is being used effectively. If the hit ratio is high (more than 0.8), the cache is performing well. A smaller ratio may indicate that the package cache should be increased.

You will need to experiment with the size of the package cache to find the optimal number for the *pckcachesz* configuration parameter. For example, you might be able to use a smaller package cache size if there is no increase in the *pkg\_cache\_inserts* data element when you decrease the size of the cache. Decreasing the package cache size frees up system resources for other work. It is also possible that you could improve overall system performance by increasing the size of the package cache if by doing so, you decrease the number of package cache inserts. This experimentation is best done under full workload conditions.

You can use this data element with *ddl\_sql\_stmts* to determine whether or not the execution of DDL statements is impacting the performance of the package cache. Sections for dynamic SQL statements can become invalid when DDL statements are executed. Invalid sections are implicitly prepared by the system when next used. The execution of a DDL statement could invalidate a number of sections and the resulting extra overhead incurred when preparing those sections could significantly impact performance. In this case, the package cache hit ratio reflects the implicit recompilation of invalid

sections and not the insertion of new sections into the cache, so increasing the size of the package cache will not improve overall performance. You might find it less confusing to tune the cache for an application on its own before working in the full environment.

It is necessary to determine the role that DDL statements are playing in the value of the package cache hit ratio before deciding on what action to take. If DDL statements rarely occur, then cache performance may be improved by increasing its size. If DDL statements are frequent, then improvements may require that you limit the use of DDL statements (possibly to specific time periods).

The *static\_sql\_stmts* and *dynamic\_sql\_stmts* counts can be used to help provide information on the quantity and type of sections being cached.

See the *Administration Guide* for more information on the Package Cache Size (*pckcachesz*) configuration parameter.

**Note:** You may want to use this information at the database level to calculate the average package cache hit ratio all each applications. You should look at this information at an application level to find out the exact package cache hit ratio for a given application. It may not be worthwhile to increase the size of the package cache in order to satisfy the cache requirements of an application that only executes infrequently.

## Package Cache Inserts

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
API Element Name	pkg_cache_inserts	
Element Type	counter	
Related Information	<ul style="list-style-type: none"> <li>• "Package Cache Lookups" on page 123</li> <li>• "Section Lookups" on page 125</li> <li>• "Section Inserts" on page 126</li> </ul>	

**Description:** The total number of times that a requested section was not available for use and had to be loaded into the package cache. This count includes any implicit prepares performed by the system.

**Usage:** In conjunction with "Package Cache Lookups", you can calculate the package cache hit ratio using the following formula:

$$1 - (\text{Package Cache Inserts} / \text{Package Cache Lookups})$$

See "Package Cache Lookups" on page 123 for information on using this element.

## Section Lookups

<b>Snapshot Information Level</b>	<b>API Structure(s)</b>	<b>Monitor Switch</b>
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
<b>API Element Name</b>	appl_section_lookups	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Package Cache Lookups” on page 123</li><li>• “Package Cache Inserts” on page 124</li><li>• “Section Inserts” on page 126</li></ul>	

**Description:** Lookups of SQL sections by an application from its SQL work area.

**Usage:** Each agent has access to a unique SQL work area where the working copy of any executable section is kept. In partitioned databases, this work area is shared by all non-SMP agents. In other environments and with SMP agents, each agent has its own unique SQL work area.

This counter indicates how many times the SQL work area was accessed by agents for an application. It is a cumulative total of all lookups on all SQL work heaps for agents working for this application.

You can use this element in conjunction with “*Section Inserts*” on page 126 to tune the size of the heap used for the SQL work area. In partitioned databases this size is controlled by the *app\_ctl\_heap\_sz* configuration parameter. SQL work area size in other database environments use the the *applheapsz* configuration parameter. The size of the SQL work area for SMP agents is controlled by *applheapsz* in all environments.

## Section Inserts

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
<b>API Element Name</b>	appl_section_inserts	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Package Cache Lookups” on page 123</li><li>• “Package Cache Inserts” on page 124</li><li>• “Section Lookups” on page 125</li></ul>	

**Description:** Inserts of SQL sections by an application from its SQL work area.

**Usage:** The working copy of any executable section is stored in a unique SQL work area. This is a count of when a copy was not available and had to be inserted. See “Section Lookups” on page 125 for more information on using sections.

## Database Heap

The following database system monitor elements are used for database heaps:

- “Maximum Database Heap Allocated”

### Maximum Database Heap Allocated

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
<b>Resettable</b>	No	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
<b>API Element Name</b>	db_heap_top	
<b>Element Type</b>	water mark	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• None</li></ul>	

**Description:** The largest amount of database heap allocated and used by the database, since the first application connected to the database (in bytes).

**Usage:** You may use this element to evaluate the setting of the *dbheap* configuration parameter, which is described in the *Administration Guide*. The *dbheap* parameter limits the amount of storage that can be allocated for database heap.

If the value of this element is the same as the *dbheap* parameter, it is quite likely that an application has received an error indicating that there was not enough storage available.



## Logging

The following database system monitor elements are used only when circular logging is being used. That is, they are not used if either the *logretain* or *userexit* configuration parameter is enabled.

- “Maximum Secondary Log Space Used”
- “Maximum Total Log Space Used” on page 128
- “Secondary Logs Allocated Currently” on page 129

The following database system monitor elements are used for all types of logging:

- “Number of Log Pages Read” on page 129
- “Number of Log Pages Written” on page 130
- “Unit of Work Log Space Used” on page 130

For more information about logging and log configuration parameters, see the *Administration Guide*.

### Maximum Secondary Log Space Used

<b>Snapshot Information Level</b> Database	<b>API Structure(s)</b> sqlm_dbase	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>Event Type</b> Database	<b>Event Record(s)</b> sqlm_db_event	
<b>API Element Name</b> <b>Element Type</b>	sec_log_used_top water mark	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Unit of Work Log Space Used” on page 130</li><li>• “Secondary Logs Allocated Currently” on page 129</li><li>• “Maximum Total Log Space Used” on page 128</li></ul>	

**Description:** The maximum amount of secondary log space used (in bytes).

**Usage:** You may use this element in conjunction with *Secondary Logs Allocated Currently* and *Maximum Total Log Space Used* to show your current dependency on secondary logs. If this value is high, you may need larger log files, or more primary log files, or more frequent COMMIT statements within your application.

As a result, you may need to adjust the following configuration parameters:

- logfilisz
- logprimary
- logsecond
- logretain

The value will be zero if the database does not have any secondary log files. This would be the case if there were none defined.

For more information, see the *Administration Guide*.

**Note:** While the database system monitor information is given in bytes, the configuration parameters are set in pages, which are each 4K bytes.

## Maximum Total Log Space Used

<b>Snapshot Information Level</b> Database	<b>API Structure(s)</b> sqlm_dbase	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>Event Type</b> Database	<b>Event Record(s)</b> sqlm_db_event	
<b>API Element Name</b> <b>Element Type</b>	tot_log_used_top water mark	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Unit of Work Log Space Used” on page 130</li> <li>• “Secondary Logs Allocated Currently” on page 129</li> <li>• “Maximum Secondary Log Space Used” on page 127</li> </ul>	

**Description:** The maximum amount of total log space used (in bytes).

**Usage:** You can use this element to help you evaluate the amount of primary log space that you have allocated. Comparing the value of this element with the amount of primary log space you have allocated can help you to evaluate your configuration parameter settings. Your primary log space allocation can be calculated using the following formula:

$$\text{logprimary} \times \text{logfilsiz} \times 4096 \text{ (see note below)}$$

You can use this element in conjunction with *Maximum Secondary Log Space Used* and *Secondary Logs Allocated Currently* to show your current dependency on secondary logs.

This value includes space used in both primary and secondary log files, and is only returned if circular logging is used. (That is, it is not returned if either the *logretain* or *userexit* configuration parameter is enabled.)

As a result, you may need to adjust the following configuration parameters:

- logfilsz
- logprimary
- logsecond
- logretain

For more information, see the *Administration Guide*.

**Note:** While the database system monitor information is given in bytes, the configuration parameters are set in pages, which are each 4K bytes.

## Secondary Logs Allocated Currently

<b>Snapshot Information Level</b> Database	<b>API Structure(s)</b> sqlm_dbase	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	sec_logs_allocated gauge	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Unit of Work Log Space Used” on page 130</li><li>• “Maximum Secondary Log Space Used” on page 127</li><li>• “Maximum Total Log Space Used” on page 128</li></ul>	

**Description:** The total number of secondary log files that are currently being used for the database.

**Usage:** You may use this element in conjunction with *Maximum Secondary Log Space Used* and *Maximum Total Log Space Used* to show your current dependency on secondary logs. If this value is consistently high, you may need larger log files, or more primary log files, or more frequent COMMIT statements within your application.

As a result, you may need to adjust the following configuration parameters:

- logfilisz
- logprimary
- logsecond
- logretain

For more information, see the *Administration Guide*.

## Number of Log Pages Read

<b>Snapshot Information Level</b> Database	<b>API Structure(s)</b> sqlm_dbase	<b>Monitor Switch</b> Basic
<b>Resettable</b>	Yes	
<b>Event Type</b> Database	<b>Event Record(s)</b> sqlm_db_event	
<b>API Element Name</b> <b>Element Type</b>	log_reads counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “When Counters are Initialized” on page 20</li><li>• “Number of Log Pages Written” on page 130</li></ul>	

**Description:** The number of log pages read from disk by the logger.

**Usage:** You can use this element with an operating system monitor to quantify the amount of I/O on a device that is attributable to database activity.

## Number of Log Pages Written

<b>Snapshot Information Level</b> Database	<b>API Structure(s)</b> sqlm_dbase	<b>Monitor Switch</b> Basic
<b>Resettable</b>	Yes	
<b>Event Type</b> Database	<b>Event Record(s)</b> sqlm_db_event	
<b>API Element Name</b> <b>Element Type</b>	log_writes counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “When Counters are Initialized” on page 20</li><li>• “Number of Log Pages Read” on page 129</li></ul>	

**Description:** The number of log pages written to disk by the logger.

**Usage:** You may use this element with an operating system monitor to quantify the amount of I/O on a device that is attributable to database activity.

## Unit of Work Log Space Used

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_appl	<b>Monitor Switch</b> Unit of Work
<b>Resettable</b>	No	
<b>Event Type</b> Transaction	<b>Event Record(s)</b> sqlm_xaction_event	
<b>API Element Name</b>	uow_log_space_used (Snapshot) log_space_used (Event)	
<b>Element Type</b>	gauge	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “Secondary Logs Allocated Currently” on page 129</li><li>• “Maximum Secondary Log Space Used” on page 127</li><li>• “Maximum Total Log Space Used” on page 128</li></ul>	

**Description:** The amount of log space (in bytes) used in the current unit of work of the monitored application.

**Usage:** You may use this element to understand the logging requirements at the unit of work level.

---

## Database and Application Activity

The following sections provide information on database and application activity.

### Locks and Deadlocks

The following elements provide information about locks and deadlocks:

- “Locks Held” on page 131
- “Total Lock List Memory In Use” on page 132

- “Deadlocks Detected” on page 133
- “Lock Escalations” on page 134
- “Exclusive Lock Escalations” on page 135
- “Lock Mode” on page 136
- “Lock Status” on page 137
- “Lock Object Type Waited On” on page 137
- “Lock Object Name” on page 138
- “Number of Lock Timeouts” on page 139
- “Maximum Number of Locks Held” on page 139
- “Connections Involved in Deadlock” on page 140

## Locks Held

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
Lock	sqlm_dbase_lock sqlm_appl_lock	Basic Basic
<b>Resettable</b>	No	
<b>API Element Name</b>	locks_held	
<b>Element Type</b>	gauge	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Lock Escalations” on page 134</li> <li>• “Exclusive Lock Escalations” on page 135</li> <li>• “Maximum Number of Locks Held” on page 139</li> </ul>	

**Description:** The number of locks currently held.

**Usage:** If the monitor information is at the database level, this is the total number of locks currently held by all applications in the database.

If it is at the application level, this is the total number of locks currently held by all agents for the application. How you use this element depends on the level of information being returned from the database system monitor.

- At the database level, you can use it in one of two ways:
  - This element can provide summary information about locking. For example, you can calculate the average number of locks per application by dividing the value of this element by *Applications Connected Currently*. If the resulting number is high, it may indicate that you can tune one of your applications to improve performance.
  - You can also compare the value of this element against the results of the following formula to determine the number of additional locks that may be requested. This comparison can help you determine if the configuration parameters need adjusting or your applications need tuning.

$$(\text{locklist} * 4096 / 36) - \text{locks held} = \# \text{ remaining}$$

where:

- *locklist* is the configuration parameter as described in the *Administration Guide*
- 4096 is the number of bytes in one 4K page
- 36 is the number of bytes required for each lock.

**Note:** You may also use “Total Lock List Memory In Use” in a similar fashion.

- At the application level, you can use this counter to determine if the application is approaching the maximum number of locks available to it, as defined by the *maxlocks* configuration parameter. This parameter indicates the percentage of the lock list that each application can use before lock escalations occur. Lock escalations can result in a decrease in concurrency between applications connected to a database. (See the *Administration Guide* for more information about this parameter.)

Since the *maxlocks* parameter is specified as a percentage and this element is a counter, you can compare the count provided by this element against the total number of locks that can be held by an application, as calculated using the following formula:

$$(\text{locklist} * 4096 / 36) * (\text{maxlocks} / 100)$$

If you have a large number of locks, you may need to perform more commits within your application so that some of the locks can be released.

## Total Lock List Memory In Use

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
<b>Resettable</b>	No	
<b>API Element Name</b>	lock_list_in_use	
<b>Element Type</b>	gauge	
<b>Related Information</b>	• None	

**Description:** The total amount of lock list memory (in bytes) that is in use.

**Usage:** This element may be used in conjunction with the *locklist* configuration parameter to calculate the lock list utilization. If the lock list utilization is high, you may want to consider increasing the size of that parameter. See the *Administration Guide* for more information.

**Note:** When calculating utilization, it is important to note that the *locklist* configuration parameter is allocated in pages of 4K bytes each, while this monitor element provides results in bytes.

## Deadlocks Detected

<b>Snapshot Information Level</b>	<b>API Structure(s)</b>	<b>Monitor Switch</b>
Database	sqlm_dbase	Basic
Application	sqlm_appl	Lock
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
<b>API Element Name</b>	deadlocks	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “When Counters are Initialized” on page 20</li><li>• “Database Activation Timestamp” on page 41</li><li>• “Connection Request Start Timestamp” on page 62</li><li>• “Lock Escalations” on page 134</li><li>• “Exclusive Lock Escalations” on page 135</li><li>• “Application ID Holding Lock” on page 146</li></ul>	

**Description:** The total number of deadlocks that have occurred.

**Usage:** This element can indicate that applications are experiencing contention problems. These problems could be caused by the following situations:

- Lock escalations are occurring for the database
- An application may be locking tables explicitly when system-generated row locks may be sufficient
- An application may be using an inappropriate isolation level when binding
- Catalog tables are locked for repeatable read
- Applications are getting the same locks in different orders, resulting in deadlock.

You may be able to resolve the problem by determining in which applications (or application processes) the deadlocks are occurring. You may then be able to modify the application to better enable it to execute concurrently. Some applications, however, may not be capable of running concurrently.

You can use the connection timestamp monitor elements (“*Last Reset Timestamp*” on page 192, *Database Activation Timestamp*, and *Connection Request Start Timestamp*) to determine the severity of the deadlocks. For example, 10 deadlocks in 5 minutes is much more severe than 10 deadlocks in 5 hours.

The descriptions for the related elements listed above may also provide additional tuning suggestions.

## Lock Escalations

<b>Snapshot Information Level</b>	<b>API Structure(s)</b>	<b>Monitor Switch</b>
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
Transaction	sqlm_xaction_event	
<b>API Element Name</b>	lock_escals	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “When Counters are Initialized” on page 20</li> <li>• “Database Activation Timestamp” on page 41</li> <li>• “Exclusive Lock Escalations” on page 135</li> <li>• “Maximum Number of Locks Held” on page 139</li> </ul>	

**Description:** The number of times that locks have been escalated from several row locks to a table lock.

**Usage:** A lock is escalated when the total number of locks held by an application reaches the maximum amount of lock list space available to the application, or the lock list space consumed by all applications is approaching the total lock list space. The amount of lock list space available is determined by the *maxlocks* and *locklist* configuration parameters.

When an application reaches the maximum number of locks allowed and there are no more locks to escalate, it will then use space in the lock list allocated for other applications. When the entire lock list is full, an error occurs.

This data item includes a count of all lock escalations, including exclusive lock escalations.

There are several possible causes for excessive lock escalations:

- The lock list size (*locklist*) may be too small for the number of concurrent applications
- The percent of the lock list usable by each application (*maxlocks*) may be too small
- One or more applications may be using an excessive number of locks.

To resolve these problems, you may be able to:

- Increase the *locklist* configuration parameter value. See the *Administration Guide* for a description of this configuration parameter.
- Increase the *maxlocks* configuration parameter value. See the *Administration Guide* for a description of this configuration parameter.
- Identify the applications with large numbers of locks (see *Maximum Number of Locks Held*), or those that are holding too much of the lock list, using the following formula:

$$(((locks\ held * 36) / (locklist * 4096)) * 100)$$



and comparing the value to `maxlocks`. These applications can also cause lock escalations in other applications by using too large a portion of the lock list. These applications may need to resort to using table locks instead of row locks, although table locks may cause an increase in “Lock Waits” on page 141 and “Time Waited On Locks” on page 142.

- Identify applications holding too much of the lock list, using the following formula:

## Exclusive Lock Escalations

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
Transaction	sqlm_xaction_event	
API Element Name	x_lock_escals	
Element Type	counter	
Related Information	<ul style="list-style-type: none"> <li>• “When Counters are Initialized” on page 20</li> <li>• “Database Activation Timestamp” on page 41</li> <li>• “Lock Escalations” on page 134</li> <li>• “Connection Request Start Timestamp” on page 62</li> <li>• “Maximum Number of Locks Held” on page 139</li> </ul>	

**Description:** The number of times that locks have been escalated from several row locks to one exclusive table lock, or the number of times an exclusive lock on a row caused the table lock to become an exclusive lock.

**Usage:** Other applications cannot access data held by an exclusive lock; therefore it is important to track exclusive locks since they can impact the concurrency of your data.

A lock is escalated when the total number of locks held by an application reaches the maximum amount of lock list space available to the application. The amount of lock list space available is determined by the *locklist* and *maxlocks* configuration parameters.

When an application reaches the maximum number of locks allowed and there are no more locks to escalate, it will then use space in the lock list allocated for other applications. When the entire lock list is full, an error occurs.

See “*Lock Escalations*” on page 134 for possible causes and resolutions to excessive exclusive lock escalations.

An application may be using exclusive locks when share locks are sufficient. Although share locks may not reduce the total number of lock escalations share lock escalations may be preferable to exclusive lock escalations.

## Lock Mode

<b>Snapshot Information Level</b>	<b>API Structure(s)</b>	<b>Monitor Switch</b>
Application	sqlm_appl	Lock
Lock	sqlm_lock	Lock
	sqlm_lock_wait	Lock
<b>Resettable</b>	No	
<b>Event Type</b>	<b>Event Record(s)</b>	
Deadlock	sqlm_dlconn_event	
<b>API Element Name</b>	lock_mode	
<b>Element Type</b>	information	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• Other lock information</li> </ul>	

**Description:** The type of lock being held.

**Usage:** This mode can help you determine the source of contention for resources.

This element indicates one of the following, depending on the type of monitor information being examined:

- The type of lock another application holds on the object that this application is waiting to lock (for application-monitoring and deadlock-monitoring levels)
- The type of lock held on the object by this application (for object-lock levels).

The values for this field are:

Mode	Type of Lock	API Constant
	No Lock	SQLM_LNON
IS	Intention Share Lock	SQLM_LOIS
IX	Intention Exclusive Lock	SQLM_LOIX
S	Share Lock	SQLM_LOOS
SIX	Share with Intention Exclusive Lock	SQLM_LSIX
X	Exclusive Lock	SQLM_LOOX
IN	Intent None	SQLM_LOIN
Z	Super Exclusive Lock	SQLM_LOOZ
U	Update Lock	SQLM_LOOU
NS	Next Key Share Lock	SQLM_LONS
NX	Next Key Exclusive Lock	SQLM_LONX
W	Weak Exclusive Lock	SQLM_LOOW
NW	Next Key Weak Exclusive Lock	SQLM_LONW

## Lock Status

Snapshot Information Level	API Structure(s)	Monitor Switch
Lock	sqlm_lock	Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	lock_status information	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Lock Mode” on page 136</li><li>• “Lock Object Type Waited On” on page 137</li><li>• “Table File ID” on page 159</li></ul>	

**Description:** Indicates the internal status of the lock.

**Usage:** This element can help explain what is happening when an application is waiting to obtain a lock on an object. While it may appear that the application already has a lock on the object it needs, it may have to wait to obtain a different type of lock on the same object.

The lock can be in one of the following statuses:

<b>Granted state</b>	indicates that the application has the lock in the state specified by “Lock Mode” on page 136.
<b>Converting state</b>	indicates that the application is trying to change the lock held to a different type; for example, changing from a share lock to an exclusive lock.

**Note:** API users should refer to the *sqlmon.h* header file containing definitions of database system monitor constants.

## Lock Object Type Waited On

Snapshot Information Level	API Structure(s)	Monitor Switch
Application	sqlm_appl sqlm_appl_lock	Lock Lock
Lock	sqlm_lock sqlm_lock_wait	Basic Lock
<b>Resettable</b>	No	
<b>Event Type</b>	<b>Event Record(s)</b>	
Deadlock	sqlm_dlconn_event	
<b>API Element Name</b> <b>Element Type</b>	lock_object_type information	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li></ul>	

**Description:** The type of object against which the application holds a lock (for object-lock-level information), or the type of object for which the application is waiting to obtain a lock (for application-level and deadlock-level information).

**Usage:** This element can help you determine the source of contention for resources.

The objects may be one of the following types:

- Table space
- Table
- Record (or row)
- Internal (another type of lock held internally by the database manager).

## Lock Object Name

<b>Snapshot Information Level</b>	<b>API Structure(s)</b>	<b>Monitor Switch</b>
Application	sqlm_appl	Lock
Lock	sqlm_appl_lock sqlm_lock	Lock Basic
<b>Resettable</b>	No	
<b>Event Type</b>	<b>Event Record(s)</b>	
Deadlock	sqlm_dlconn_event	
<b>API Element Name</b>	lock_object_name	
<b>Element Type</b>	information	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “Lock Object Type Waited On” on page 137</li> <li>• “Table Space Name” on page 143</li> <li>• “Table Name” on page 151</li> <li>• “Table Schema Name” on page 152</li> </ul>	

**Description:** This element is provided for informational purposes only. It is the name of the object for which the application holds a lock (for object-lock-level information), or the name of the object for which the application is waiting to obtain a lock (for application-level and deadlock-level information).

**Usage:** It is the name of the object for table-level locks is the file ID (FID) for SMS and DMS table spaces. For row-level locks, the object name is the row ID (RID). For table space locks, the object name is blank.

To determine the table holding the lock, use *Table Name* and *Table Schema Name* instead of the file ID, since the file ID may not be unique.

To determine the table space holding the lock, use *Table Space Name* .

## Number of Lock Timeouts

<b>Snapshot Information Level</b>	<b>API Structure(s)</b>	<b>Monitor Switch</b>
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
<b>API Element Name</b>	lock_timeouts	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “When Counters are Initialized” on page 20</li><li>• Other elements in “Locks and Deadlocks” on page 130</li></ul>	

**Description:** The number of times that a request to lock an object timed-out instead of being granted.

**Usage:** This element can help you adjust the setting for the *locktimeout* database configuration parameter. If the number of lock time-outs becomes excessive when compared to normal operating levels, you may have an application that is holding locks for long durations. In this case, this element may indicate that you should analyze some of the other elements related to “Locks and Deadlocks” on page 130 to determine if you have an application problem.

You could also have too few lock time-outs if your *locktimeout* database configuration parameter is set too high. In this case, your applications may wait excessively to obtain a lock. See the *Administration Guide* for more information.

## Maximum Number of Locks Held

<b>Event Type</b>	<b>Event Record(s)</b>
Transaction	sqlm_xaction_event
<b>API Element Name</b>	locks_held_top
<b>Element Type</b>	counter
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Locks Held” on page 131</li><li>• “Lock Escalations” on page 134</li><li>• “Exclusive Lock Escalations” on page 135</li></ul>

**Description:** The maximum number of locks held during this transaction.

**Usage:** You can use this element to determine if your application is approaching the maximum number of locks available to it, as defined by the *maxlocks* configuration parameter. This parameter indicates the percentage of the lock list that each application can use before lock escalations occur. Lock escalations can result in a decrease in concurrency between applications connected to a database. (See the *Administration Guide* for more information about this parameter.)

Since the *maxlocks* parameter is specified as a percentage and this element is a counter, you can compare the count provided by this element against the total number of locks that can be held by an application, as calculated using the following formula:

$$(\text{locklist} * 4096 / 36) * (\text{maxlocks} / 100)$$

If you have a large number of locks, you may need to perform more commits within your application so that some of the locks can be released.

## Connections Involved in Deadlock

<b>Event Type</b>	<b>Event Record(s)</b>
Deadlock	sqlm_deadlock_event
<b>API Element Name</b>	dl_conns
<b>Element Type</b>	gauge
<b>Related Information</b>	<ul style="list-style-type: none"> <li>None</li> </ul>

**Description:** The number of connections that are involved in the deadlock.

**Usage:** Use this element in your monitoring application to identify how many deadlock connection event records will follow in the event monitor data stream.

## Lock Wait Information

The following elements provide information is returned when a DB2 agent working on behalf of an application is waiting to obtain a lock:

- “Lock Waits” on page 141
- “Time Waited On Locks” on page 142
- “Table Space Name” on page 143
- “Current Agents Waiting On Locks” on page 144
- “Total Time Unit of Work Waited on Locks” on page 144
- “Lock Wait Start Timestamp” on page 145
- “Agent ID Holding Lock” on page 145
- “Application ID Holding Lock” on page 146
- “Sequence Number Holding Lock” on page 147
- “Rolled Back Application” on page 147

## Lock Waits

<b>Snapshot Information Level</b>	<b>API Structure(s)</b>	<b>Monitor Switch</b>
Database	sqlm_dbase	Lock
Application	sqlm_appl	Lock
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
<b>API Element Name</b>	lock_waits	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “When Counters are Initialized” on page 20</li><li>• “Connection Request Start Timestamp” on page 62</li><li>• “Time Waited On Locks” on page 142</li></ul>	

**Description:** The total number of times that applications or connections waited for locks.

**Usage:** At the database level, this is the total number of times that applications have had to wait for locks within this database.

At the application-connection level, this is the total number of times that this connection requested a lock but had to wait because another connection was already holding a lock on the data.

This element may be used with *Time Waited On Locks* to calculate, at the database level, the average wait time for a lock. This calculation can be done at either the database or the application-connection level.

If the average lock wait time is high, you should look for applications that hold many locks, or have lock escalations, with a focus on tuning your applications to improve concurrency, if appropriate. If escalations are the reason for a high average lock wait time, then the values of one or both of the *locklist* and *maxlocks* configuration parameters may be too low. See the *Administration Guide* for more information.

## Time Waited On Locks

<b>Snapshot Information Level</b>	<b>API Structure(s)</b>	<b>Monitor Switch</b>
Database	sqlm_dbase	Lock
Application	sqlm_appl sqlm_appl_lock	Lock
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
Transaction	sqlm_xaction_event	
<b>API Element Name</b>	lock_wait_time	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “When Counters are Initialized” on page 20</li> <li>• “Current Agents Waiting On Locks” on page 144</li> <li>• “Lock Waits” on page 141</li> </ul>	

**Description:** The total elapsed time waited for a lock.

**Usage:** At the database level, this is the total amount of elapsed time that all applications were waiting for a lock within this database.

At the application-connection and transaction levels, this is the total amount of elapsed time that this connection or transaction has waited for a lock to be granted to it.

This element may be used in conjunction with the Lock Waits monitor element to calculate the average wait time for a lock. This calculation can be performed at either the database or the application-connection level.

When using data elements providing elapsed times, you should consider:

- Elapsed times are affected by system load, so the more processes you have running, the higher this elapsed time value.
- To calculate this data element at the database level, the database system monitor sums the application-level times. This can result in double counting elapsed times at a database level, since more than one application process can be running at the same time.

To provide meaningful data, you can calculate the average wait time for a lock, as described above.



## Table Space Name

Snapshot Information Level	API Structure(s)	Monitor Switch
Table Space	sqlm_tablespace	Buffer Pool
Application	sqlm_appl_lock	Basic
Lock	sqlm_lock sqlm_lock_wait	Lock Lock
<b>Resettable</b>	No	
<b>Event Type</b>	<b>Event Record(s)</b>	
Deadlock	sqlm_dlconn_event	
Table Space	sqlm_tablespace_header	
<b>API Element Name</b>	tablespace_name	
<b>Element Type</b>	information	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “Lock Object Type Waited On” on page 137</li></ul>	

**Description:** The name of a table space.

**Usage:** This element can help you determine the source of contention for resources.

It is equivalent to the TBSpace column in the database catalog table SYSCAT.TABLESPACE. At the application level, application-lock level, and deadlock monitoring level, this is the name of the table space that the application is waiting to lock. Another application currently holds a lock on this table space.

At the lock level, this is the name of the table space against which the application currently holds a lock.

At the table space level (when the buffer pool monitor group is ON), this is the name of the table space for which information is returned.

If you are using the database system monitor APIs, note that the API constant SQLM\_IDENT\_SZ is used to define the length of this element. Only the first 18 characters are currently used.

## Current Agents Waiting On Locks

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
Lock	sqlm_dbase_lock	Basic
<b>Resettable</b>	No	
<b>API Element Name</b>	locks_waiting	
<b>Element Type</b>	gauge	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Applications Connected Currently” on page 72</li></ul>	

**Description:** Indicates the number of agents waiting on a lock.

**Usage:** When used in conjunction with *Applications Connected Currently*, this element indicates the percentage of applications waiting on locks. If this number is high, the applications may have concurrency problems, and you should identify applications that are holding locks or exclusive locks for long periods of time.

## Total Time Unit of Work Waited on Locks

Snapshot Information Level	API Structure(s)	Monitor Switch
Application	sqlm_appl	Unit of Work
<b>Resettable</b>	No	
<b>API Element Name</b>	uow_lock_wait_time	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• Application-level information on locks</li></ul>	

**Description:** The total amount of elapsed time this unit of work has spent waiting for locks.

**Usage:** This element can help you determine the severity of the resource contention problem.

## Lock Wait Start Timestamp

Snapshot Information Level	API Structure(s)	Monitor Switch
Application	sqlm_appl	Lock
Lock	sqlm_lock_wait	Lock
<b>Resettable</b>	No	
<b>Event Type</b>	<b>Event Record(s)</b>	
Deadlock	sqlm_dlconn_event	
<b>API Element Name</b>	lock_wait_start_time	
<b>Element Type</b>	timestamp	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “Agent ID Holding Lock” on page 145</li></ul>	

**Description:** The date and time that this application started waiting to obtain a lock on the object that is currently locked by another application.

**Usage:** This element can help you determine the severity of resource contention.

## Agent ID Holding Lock

Snapshot Information Level	API Structure(s)	Monitor Switch
Application	sqlm_appl	Lock
	sqlm_appl_lock	Lock
Lock	sqlm_lock_wait	Lock
<b>Resettable</b>	No	
<b>API Element Name</b>	agent_id_holding_lock	
<b>Element Type</b>	information	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “Lock Wait Start Timestamp” on page 145</li><li>• “Application ID Holding Lock” on page 146</li></ul>	

**Description:** The application handle of the agent holding a lock for which this application is waiting. The lock monitor group must be turned on to obtain this information.

**Usage:** This element can help you determine which applications are in contention for resources.

If this element is 0 (zero) and the application is waiting for a lock, this indicates that the lock is held by an indoubt transaction. You can use either “Application ID Holding Lock” on page 146 or the command line processor LIST INDOUBT TRANSACTIONS command (which displays the application ID of the CICS agent that was processing the transaction when it became indoubt) to determine the indoubt transaction, and then either commit it or roll it back.

Note that more than one application can hold a shared lock on an object for which this application is waiting. See “Lock Mode” on page 136 for information about the type of lock that the application holds. If you are taking an application snapshot, only one of the

agent IDs holding a lock on the object will be returned. If you are taking a lock snapshot, all of the agent IDs holding a lock on the object will be identified.

## Application ID Holding Lock

Snapshot Information Level	API Structure(s)	Monitor Switch
Application	sqlm_appl	Lock
	sqlm_appl_lock	Lock
Lock	sqlm_lock_wait	Lock
<b>Resettable</b>	No	
<b>Event Type</b>	<b>Event Record(s)</b>	
Deadlock	sqlm_dlconn_event	
<b>API Element Name</b>	appl_id_holding_lk	
<b>Element Type</b>	information	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “Agent ID Holding Lock” on page 145</li> <li>• “Deadlocks Detected” on page 133</li> </ul>	

**Description:** The application ID of the application that is holding a lock on the object that this application is waiting to obtain.

**Usage:** This element can help you determine which applications are in contention for resources. Specifically, it can help you identify the application handle (agent ID) and table ID that are holding the lock. Note that you may use the LIST APPLICATIONS command to obtain information to relate the application ID with an agent ID. However, it is a good idea to collect this type of information when you take the snapshot, as it could be unavailable if the application ends before you run the LIST APPLICATIONS command.

If you are using the database system monitor APIs, note that the API constant SQLM\_APPLID\_SZ is used to define the length of this element. Only the first 30 characters are currently used.

Note that more than one application can hold a shared lock on an object for which this application is waiting to obtain a lock. See “Lock Mode” on page 136 for information about the type of lock that the application holds. If you are taking an application snapshot, only one of the application IDs holding a lock on the object will be returned. If you are taking a lock snapshot, all of the application IDs holding a lock on the object will be returned.

## Sequence Number Holding Lock

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_appl sqlm_appl_lock	<b>Monitor Switch</b> Basic Basic
<b>Resettable</b>	No	
<b>Event Type</b> Deadlock	<b>Event Record(s)</b> sqlm_dlconn_event	
<b>API Element Name</b> <b>Element Type</b>	sequence_no_holding_lk information	
<b>Related Information</b>	• None	

**Description:** This element is reserved for future use. In this release, its value will always be “0001.” In future releases of the product, it may contain different values.

## Rolled Back Application

<b>Event Type</b> Deadlock	<b>Event Record(s)</b> sqlm_deadlock_event
<b>API Element Name</b> <b>Element Type</b>	rolled_back_appl_id information
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Service Level” on page 37</li><li>• “Maximum Number of Coordinating Agents” on page 76</li></ul>

**Description:** Application id that was rolled back when a deadlock occurred.

**Usage:** A system administrator can use this information to determine which application did not complete its updates, and determine which applications should be restarted

## Rollforward Monitoring

Recovering database changes can be a time consuming process. You can use the database system monitor to monitor the progression of a recovery. The following elements provide information about rollforward status:

- “Rollforward Timestamp” on page 148
- “Tablespace Being Rolled Forward” on page 148
- “Rollforward Type” on page 148
- “Log Being Rolled Forward” on page 149
- “Log Phase” on page 149
- “Number of Rollforward Table Spaces” on page 149

## Rollforward Timestamp

<b>Snapshot Information Level</b> Table Space	<b>API Structure(s)</b> sqlm_rollfwd_info	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	rf_timestamp timestamp	
<b>Related Information</b>	• “Tablespace Being Rolled Forward” on page 148	

**Description:** The timestamp of the log being processed.

**Usage:** If a rollforward is in progress, this is the timestamp of the log record being processed. This is an indicator of the data changes that will be recovered.

## Tablespace Being Rolled Forward

<b>Snapshot Information Level</b> Table Space	<b>API Structure(s)</b> sqlm_rollfwd_ts_info	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	ts_name information	
<b>Related Information</b>	• “Rollforward Timestamp” on page 148	

**Description:** The name of the table space currently rolled forward.

**Usage:** If a rollforward is in progress, this element identifies the table spaces involved.

## Rollforward Type

<b>Snapshot Information Level</b> Table Space	<b>API Structure(s)</b> sqlm_rollfwd_info	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	rf_type information	
<b>Related Information</b>	• None	

**Description:** The type of rollforward in progress.

**Usage:** An indicator of whether recovery is happening at a database or table space level. For more information on rollforward recovery at the database or table space level see the *Administration Guide*.

## Log Being Rolled Forward

Snapshot Information Level	API Structure(s)	Monitor Switch
Table Space	sqlm_rollfwd_info	Basic
Resettable	No	
API Element Name	rf_log_num	
Element Type	information	
Related Information	<ul style="list-style-type: none"><li>None</li></ul>	

**Description:** The log being processed.

**Usage:** If a rollforward is in progress, this element identifies the log involved.

## Log Phase

Snapshot Information Level	API Structure(s)	Monitor Switch
Table Space	sqlm_rollfwd_info	Basic
Resettable	No	
API Element Name	rf_status	
Element Type	information	
Related Information	<ul style="list-style-type: none"><li>None</li></ul>	

**Description:** The status of the recovery.

**Usage:** This element indicates the progression of a recovery. It indicates if the recovery is in an undo (rollback) or redo (rollforward) phase.

## Number of Rollforward Table Spaces

Snapshot Information Level	API Structure(s)	Monitor Switch
Table Space	sqlm_rollfwd_info	Basic
Resettable	No	
API Element Name	rf_num_tspaces	
Element Type	counter	
Related Information	<ul style="list-style-type: none"><li>None</li></ul>	

**Description:** The number of table spaces involved in a rollforward.

**Usage:** This is a counter of the table spaces involved in recovery.

## Table Activity

The following elements provide information about the tables:

- “Table Type” on page 150
- “Table Name” on page 151
- “Table Schema Name” on page 152
- “Rows Deleted” on page 153

- “Rows Inserted” on page 153
- “Rows Updated” on page 154
- “Rows Selected” on page 154
- “Rows Written” on page 155
- “Rows Read” on page 156
- “Accesses to Overflowed Records” on page 157
- “Internal Rows Deleted” on page 157
- “Internal Rows Updated” on page 158
- “Internal Rows Inserted” on page 159
- “Table File ID” on page 159

## Table Type

<b>Snapshot Information Level</b> Table	<b>API Structure(s)</b> sqlm_table	<b>Monitor Switch</b> Table
<b>Resettable</b>	No	
<b>Event Type</b> Table	<b>Event Record(s)</b> sqlm_table_event	
<b>API Element Name</b> <b>Element Type</b>	table_type information	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “Table File ID” on page 159</li> </ul>	

**Description:** The type of table for which information is returned.

**Usage:** You can use this element to help identify the table for which information is returned. If the table is a user table or a system catalog table, you can use *Table Name* and *Table Schema Name* to identify the table.

The type of table may be one of the following:

- User table.
- User table that has been dropped. The table type will only be updated after the changes are committed (either explicitly or implicitly).
- Temporary table. Information regarding temporary tables is returned, even though the tables are not kept in the database after being used. You may still find information about this type of table useful.
- System catalog table.
- Reorganization table. A table created and used by the database manager while performing a reorganization of another table.



## Table Name

Snapshot Information Level	API Structure(s)	Monitor Switch
Table	sqlm_table	Table
Application	sqlm_appl	Lock
	sqlm_appl_lock	Lock
Lock	sqlm_lock	Lock
	sqlm_lock_wait	Lock
<b>Resettable</b>	No	
<b>Event Type</b>	<b>Event Record(s)</b>	
Table	sqlm_table_event	
Deadlock	sqlm_dlconn_event	
<b>API Element Name</b>	table_name	
<b>Element Type</b>	information	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “Table Schema Name” on page 152</li><li>• “Lock Object Type Waited On” on page 137</li></ul>	

**Description:** The name of the table.

**Usage:** Along with *Table Schema Name*, this element can help you determine the source of contention for resources.

At the application-level, application-lock level, and deadlock-monitoring-level, this is the table that the application is waiting to lock, because it is currently locked by another application. For snapshot monitoring, this item is only valid when the “lock” monitor group information is turned on, and when *Lock Object Type Waited On* indicates that the application is waiting to obtain a table lock.

For snapshot monitoring at the object-lock level, this item is returned for table-level and row-level locks. The table reported at this level is the table against which this application holds these locks.

For snapshot and event monitoring at the table level, this is the table for which information has been collected. This element is blank for temporary tables, reorganization tables, and tables that were dropped. Table names are only provided for catalog and user tables. For snapshot monitoring, this element is only valid when the “table” monitor group information is turned on.

If you are using the database system monitor APIs, note that the API constant `SQLM_IDENT_SZ` is used to define the length of this element. Only the first 18 characters are currently used.

## Table Schema Name

Snapshot Information Level	API Structure(s)	Monitor Switch
Table	sqlm_table	Table
Application	sqlm_appl	Lock
	sqlm_appl_lock	Lock
Lock	sqlm_lock	Lock
	sqlm_lock_wait	Lock
<b>Resettable</b>	No	
Event Type	Event Record(s)	
Table	sqlm_table_event	
Deadlock	sqlm_dlconn_event	
API Element Name	table_schema	
Element Type	information	
Related Information	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “Table Name” on page 151</li> <li>• “Lock Object Type Waited On” on page 137</li> </ul>	

**Description:** The schema of the table.

**Usage:** Along with *Table Name*, this element can help you determine the source of contention for resources.

For application-level, application-lock-level, deadlock-monitoring-level, this is the schema of the table that the application is waiting to lock, because it is currently locked by another application. This element is only set if *Lock Object Type Waited On* indicates that the application is waiting to obtain a table lock. For snapshot monitoring at the application-level and application-lock levels, this item is only valid when the “lock” monitor group information is turned on.

For snapshot monitoring at the object-lock level, this item is returned for table and row level locks. The table reported at this level is the table against which this application holds these locks.

For snapshot and event monitoring at the table level, this element identifies the schema of the table for which information has been collected. This element is blank for temporary tables, reorganization tables, and tables that were dropped. Schema names are provided only for catalog and user tables. For snapshot monitoring, this element is valid only when the “table” monitor group information is turned on.

If you are using the database system monitor APIs, note that the API constant `SQLM_IDENT_SZ` is used to define the length of this element. Only the first 8 characters are currently used.

## Rows Deleted

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
<b>API Element Name</b>	rows_deleted	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “When Counters are Initialized” on page 20</li><li>• “Internal Rows Deleted” on page 157</li></ul>	

**Description:** This is the number of row deletions attempted.

**Usage:** You can use this element to gain insight into the current level of activity within the database manager.

This count does not include the attempts counted in *Internal Rows Deleted*.

## Rows Inserted

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
<b>API Element Name</b>	rows_inserted	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “When Counters are Initialized” on page 20</li></ul>	

**Description:** This is the number of row insertions attempted.

**Usage:** You can use this element to gain insight into the current level of activity within the database manager.

## Rows Updated

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
<b>API Element Name</b>	rows_updated	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “When Counters are Initialized” on page 20</li><li>• “Internal Rows Updated” on page 158</li></ul>	

**Description:** This is the number of row updates attempted.

**Usage:** You can use this element to gain insight into the current level of activity within the database manager.

This value does not include updates counted in *Internal Rows Updated*.. However, rows that are updated by more than one update statement are counted for each update.

## Rows Selected

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
<b>API Element Name</b>	rows_selected	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “When Counters are Initialized” on page 20</li><li>• “Select SQL Statements Executed” on page 168</li></ul>	

**Description:** This is the number of rows that have been selected and returned to the application.

**Usage:** You can use this element to gain insight into the current level of activity within the database manager.

This element does not include a count of rows read for actions such as COUNT(\*) or joins.

## Rows Written

Snapshot Information Level	API Structure(s)	Monitor Switch
Table	sqlm_table	Table
Application	sqlm_appl	Basic
	sqlm_stmt	Basic
	sqlm_subsection	Statement
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Connection	sqlm_conn_event	
Table	sqlm_table_event	
Statement	sqlm_stmt_event	
Transaction	sqlm_xaction_event	
API Element Name	rows_written	
Element Type	counter	
Related Information	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “When Counters are Initialized” on page 20</li> <li>• “Rows Read” on page 156</li> <li>• “Internal Rows Inserted” on page 159</li> <li>• “Internal Rows Deleted” on page 157</li> <li>• “Internal Rows Updated” on page 158</li> </ul>	

**Description:** This is the number of rows changed (inserted, deleted or updated) in the table.

**Usage:** A high value for table-level information indicates there is heavy usage of the table and you may want to use the Run Statistics (RUNSTATS) utility to maintain efficiency of the packages used for this table.

For application-connections and statements, this element includes the number of rows inserted, updated, and deleted in temporary tables.

At the application, transaction, and statement levels, this element can be useful for analyzing the relative activity levels, and for identifying candidates for tuning.

## Rows Read

Snapshot Information Level	API Structure(s)	Monitor Switch
Table	sqlm_table	Table
Application	sqlm_appl	Basic
	sqlm_stmt	Basic
	sqlm_subsection	Statement
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Connection	sqlm_conn_event	
Table	sqlm_table_event	
Statement	sqlm_stmt_event	
Transaction	sqlm_xaction_event	
API Element Name	rows_read	
Element Type	counter	
Related Information	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “When Counters are Initialized” on page 20</li><li>• “Rows Written” on page 155</li><li>• “Accesses to Overflowed Records” on page 157</li></ul>	

**Description:** This is the number of rows read from the table.

**Usage:** This element helps you identify tables with heavy usage for which you may want to create additional indexes. To avoid the maintenance of unnecessary indexes, you may use the SQL EXPLAIN statement, described in the *Administration Guide* to determine if the package uses an index.

This count is **not** the number of row that were returned to the calling application. Rather, it is the number of rows that had to be read in order to return the result set. For example, the following statement returns one row to the application, but many rows are read to determine the average salary:

```
SELECT AVG(SALARY) FROM USERID.EMPLOYEE
```

This count includes the value in *Accesses to Overflowed Records*.

## Accesses to Overflowed Records

Snapshot Information Level	API Structure(s)	Monitor Switch
Table	sqlm_table	Table
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Table	sqlm_table_event	
API Element Name	overflow_accesses	
Element Type	counter	
Related Information	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “When Counters are Initialized” on page 20</li> <li>• “Rows Read” on page 156</li> <li>• “Rows Written” on page 155</li> </ul>	

**Description:** The number of accesses (reads and writes) to overflowed rows of this table.

**Usage:** Overflowed rows indicate that data fragmentation has occurred. If this number is high, you may be able to improve table performance by reorganizing the table using the REORG utility, which cleans up this fragmentation.

A row overflows if it is updated and no longer fits in the data page where it was originally written. This usually happens as a result of an update of a VARCHAR or an ALTER TABLE statement.

## Internal Rows Deleted

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
	sqlm_stmt	Basic
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
Statement	sqlm_stmt_event	
API Element Name	int_rows_deleted	
Element Type	counter	
Related Information	<ul style="list-style-type: none"> <li>• “When Counters are Initialized” on page 20</li> <li>• “Rows Deleted” on page 153</li> </ul>	

**Description:** This is the number of rows deleted from the database as a result of internal activity.

**Usage:** This element can help you gain insight into internal activity within the database manager of which you might not be aware. If this activity is high, you may want to evaluate your table design to determine if the referential constraints or triggers that you have defined on your database are necessary.

Internal delete activity can be a result of:

- A cascading delete enforcing an ON CASCADE DELETE referential constraint
- A trigger being fired.

## Internal Rows Updated

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
Application	sqlm_appl sqlm_stmt	Basic Basic
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
Statement	sqlm_stmt_event	
API Element Name	int_rows_updated	
Element Type	counter	
Related Information	<ul style="list-style-type: none"> <li>• “When Counters are Initialized” on page 20</li> <li>• “Rows Updated” on page 154</li> </ul>	

**Description:** This is the number of rows updated from the database as a result of internal activity.

**Usage:** This element can help you gain insight into internal activity within the database manager of which you might not be aware. If this activity is high, you may want to evaluate your table design to determine if the referential constraints that you have defined on your database are necessary.

Internal update activity can be a result of:

- A *set null* row update enforcing a referential constraint defined with the ON DELETE SET NULL rule
- A trigger being fired.



## Internal Rows Inserted

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
	sqlm_stmt	Basic
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
Statement	sqlm_stmt_event	
<b>API Element Name</b>	int_rows_inserted	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “When Counters are Initialized” on page 20</li><li>• “Rows Inserted” on page 153</li></ul>	

**Description:** The number of rows inserted into the database as a result of internal activity caused by triggers.

**Usage:** This element can help you gain insight into the internal activity within the database manager. If this activity is high, you may want to evaluate your design to determine if you can alter it to reduce this activity.

## Table File ID

Snapshot Information Level	API Structure(s)	Monitor Switch
Application	sqlm_appl	Lock
Table	sqlm_table	Table
Lock	sqlm_appl_lock	Lock
	sqlm_lock	Lock
<b>Resettable</b>	No	
<b>API Element Name</b>	table_file_id	
<b>Element Type</b>	information	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “Table Name” on page 151</li><li>• “Table Schema Name” on page 152</li><li>• “Table Type” on page 150</li></ul>	

**Description:** This is the file ID (FID) for the table.

**Usage:** This element is provided for information purposes only. It is returned for compatibility with previous versions of the database system monitor, and it may **not** uniquely identify the table. Use *Table Name* and *Table Schema Name* to identify the table.

## SQL Cursors

The following elements provide information about the SQL cursors:

- “Open Remote Cursors”
- “Open Remote Cursors with Blocking” on page 161
- “Rejected Block Cursor Requests” on page 162
- “Accepted Block Cursor Requests” on page 162
- “Open Local Cursors” on page 163
- “Open Local Cursors with Blocking” on page 163

### Open Remote Cursors

<b>Snapshot Information Level</b>	<b>API Structure(s)</b>	<b>Monitor Switch</b>
Application	sqlm_appl	Basic
<b>Resettable</b>	No	
<b>API Element Name</b>	open_rem_curs	
<b>Element Type</b>	gauge	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Open Remote Cursors with Blocking” on page 161</li><li>• “Open Local Cursors” on page 163</li></ul>	

**Description:** The number of remote cursors currently open for this application, including those cursors counted by *Open Remote Cursors with Blocking*.

**Usage:** You may use this element in conjunction with *Open Remote Cursors with Blocking* to calculate the percentage of remote cursors that are blocking cursors. If the percentage is low, you may be able to improve performance by improving the row blocking in the application. See *Open Remote Cursors with Blocking* for more information.

For the number of open cursors used by applications connected to a local database, see *Open Local Cursors*.

## Open Remote Cursors with Blocking

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_appl	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	open_rem_curs_blk gauge	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Open Remote Cursors” on page 160</li><li>• “Rejected Block Cursor Requests” on page 162</li><li>• “Accepted Block Cursor Requests” on page 162</li><li>• “Open Local Cursors” on page 163</li><li>• “Open Local Cursors with Blocking” on page 163</li></ul>	

**Description:** The number of remote blocking cursors currently open for this application.

**Usage:** You can use this element in conjunction with *Open Remote Cursors* to calculate the percentage of remote cursors that are blocking cursors. If the percentage is low, you may be able to improve performance by improving the row blocking in the application:

- Check the pre-compile options for record blocking for treatment of ambiguous cursors
- Redefine cursors to allow for blocking (for example, if possible, specify FOR FETCH ONLY on your cursors).

*Rejected Block Cursor Requests* and *Accepted Block Cursor Requests* provide additional information that may help you tune your configuration parameters to improve row blocking in your application.

For the number of open blocking cursors used by applications connected to a local database see *Open Local Cursors with Blocking*.

## Rejected Block Cursor Requests

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_appl	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>Event Type</b> Connection	<b>Event Record(s)</b> sqlm_conn_event	
<b>API Element Name</b> <b>Element Type</b>	rej_curs_blk counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Accepted Block Cursor Requests” on page 162</li><li>• “Open Local Cursors” on page 163</li><li>• “Open Local Cursors with Blocking” on page 163</li></ul>	

**Description:** The number of times that a request for an I/O block at server was rejected and the request was converted to non-blocked I/O.

**Usage:** If there are many cursors blocking data, the communication heap may become full. When this heap is full, an error is not returned. Instead, no more I/O blocks are allocated for blocking cursors. If cursors are unable to block data, performance can be affected.

If a large number of cursors were unable to perform data blocking, you may be able to improve performance by:

- Increasing the size of the *query\_heap* database manager configuration parameter. For more information see the *Administration Guide*.

## Accepted Block Cursor Requests

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_appl	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>Event Type</b> Connection	<b>Event Record(s)</b> sqlm_conn_event	
<b>API Element Name</b> <b>Element Type</b>	acc_curs_blk counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Rejected Block Cursor Requests” on page 162</li><li>• “Open Local Cursors” on page 163</li><li>• “Open Local Cursors with Blocking” on page 163</li></ul>	

**Description:** The number of times that a request for an I/O block was accepted.

**Usage:** You can use this element in conjunction with *Rejected Block Cursor Requests* to calculate the percentage of blocking requests that are accepted and/or rejected.

See *Rejected Block Cursor Requests* for suggestions on how to use this information to tune your configuration parameters.

## Open Local Cursors

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_appl	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	open_loc_curs gauge	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Open Local Cursors with Blocking” on page 163</li><li>• “Open Remote Cursors” on page 160</li><li>• “Open Remote Cursors with Blocking” on page 161</li><li>• “Rejected Block Cursor Requests” on page 162</li><li>• “Accepted Block Cursor Requests” on page 162</li></ul>	

**Description:** The number of local cursors currently open for this application, including those cursors counted by *Open Local Cursors with Blocking*.

**Usage:** You may use this element in conjunction with *Open Local Cursors with Blocking* to calculate the percentage of local cursors that are blocking cursors. If the percentage is low, you may be able to improve performance by improving the row blocking in the application.

For cursors used by remote applications, see *Open Remote Cursors*.

## Open Local Cursors with Blocking

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_appl	<b>Monitor Switch</b> Basic
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	open_loc_curs_blk gauge	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Open Local Cursors” on page 163</li><li>• “Open Remote Cursors” on page 160</li><li>• “Open Remote Cursors with Blocking” on page 161</li><li>• “Rejected Block Cursor Requests” on page 162</li><li>• “Accepted Block Cursor Requests” on page 162</li></ul>	

**Description:** The number of local blocking cursors currently open for this application.

**Usage:** You may use this element in conjunction with *Open Local Cursors* to calculate the percentage of local cursors that are blocking cursors. If the percentage is low, you may be able to improve performance by improving the row blocking in the application:

- Check the pre-compile options for record blocking for treatment of ambiguous cursors
- Redefine cursors to allow for blocking (for example, if possible, specify FOR FETCH ONLY on your cursors).

*Rejected Block Cursor Requests* and *Accepted Block Cursor Requests* provide additional information that may help you tune your configuration parameters to improve row blocking in your application.

For blocking cursors used by remote applications, see *Open Remote Cursors with Blocking*.

## SQL Statement Activity

The following elements provide information about SQL statement activity:

- “Static SQL Statements Attempted”
- “Dynamic SQL Statements Attempted” on page 165
- “Failed Statement Operations” on page 165
- “Commit Statements Attempted” on page 166
- “Rollback Statements Attempted” on page 167
- “Select SQL Statements Executed” on page 168
- “Update/Insert/Delete SQL Statements Executed” on page 168
- “Data Definition Language (DDL) SQL Statements” on page 169
- “Internal Automatic Rebinds” on page 170
- “Internal Commits” on page 171
- “Internal Rollbacks” on page 172
- “Internal Rollbacks Due To Deadlock” on page 173
- “SQL Requests Since Last Commit” on page 173
- “Statement Node” on page 174
- “Binds/Precompiles Attempted” on page 174

### Static SQL Statements Attempted

<b>Snapshot Information Level</b>	<b>API Structure(s)</b>	<b>Monitor Switch</b>
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
<b>API Element Name</b>	static_sql_stmts	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “When Counters are Initialized” on page 20</li> <li>• “Failed Statement Operations” on page 165</li> </ul>	

**Description:** The number of static SQL statements that were attempted.

**Usage:** You can use this element to calculate the total number of successful SQL statements at the database or application level:

```

Dynamic SQL Statements Attempted
+ Static SQL Statements Attempted
- Failed Statement Operations
= throughput during monitoring period

```

## Dynamic SQL Statements Attempted

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
<b>API Element Name</b>	dynamic_sql_stmts	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “When Counters are Initialized” on page 20</li><li>• “Failed Statement Operations” on page 165</li></ul>	

**Description:** The number of dynamic SQL statements that were attempted.

**Usage:** You can use this element to calculate the total number of successful SQL statements at the database or application level:

Dynamic SQL Statements Attempted  
+ Static SQL Statements Attempted  
- Failed Statement Operations  
= throughput during monitoring period

## Failed Statement Operations

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
<b>API Element Name</b>	failed_sql_stmts	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “When Counters are Initialized” on page 20</li><li>• “Dynamic SQL Statements Attempted” on page 165</li><li>• “Static SQL Statements Attempted” on page 164</li></ul>	

**Description:** The number of SQL statements that were attempted, but failed.

**Usage:** You can use this element to calculate the total number of successful SQL statements at the database or application level:

Dynamic SQL Statements Attempted  
+ Static SQL Statements Attempted  
- Failed Statement Operations  
= throughput during monitoring period

This count includes all SQL statements that received a negative SQLCODE.

This element may also help you in determining reasons for poor performance, since failed statements mean time wasted by the database manager and as a result, lower throughput for the database.

## Commit Statements Attempted

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
API Element Name	commit_sql_stmts	
Element Type	counter	
Related Information	<ul style="list-style-type: none"> <li>• “When Counters are Initialized” on page 20</li> <li>• “Internal Commits” on page 171</li> <li>• “Rollback Statements Attempted” on page 167</li> <li>• “Internal Rollbacks” on page 172</li> <li>• “Internal Rollbacks Due To Deadlock” on page 173</li> </ul>	

**Description:** The total number of SQL COMMIT statements that have been attempted.

**Usage:** A small rate of change in this counter during the monitor period may indicate that applications are not doing frequent commits, which may lead to problems with logging and data concurrency.

You can also use this element to calculate the total number of units of work by calculating the sum of the following:

```

commit statements attempted
+ internal commits
+ rollback statements attempted
+ internal rollbacks

```

**Note:** The units of work calculated will only include those since the later of:

- The connection to the database (for database-level information, this is the time of the first connection)
- The last reset of the database monitor counters.

This calculation can be done at a database or application level.



## Rollback Statements Attempted

<b>Snapshot Information Level</b>	<b>API Structure(s)</b>	<b>Monitor Switch</b>
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
<b>API Element Name</b>	rollback_sql_stmts	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “When Counters are Initialized” on page 20</li><li>• “Statement Type” on page 175</li><li>• “Commit Statements Attempted” on page 166</li><li>• “Internal Commits” on page 171</li><li>• “Internal Rollbacks” on page 172</li><li>• “Internal Rollbacks Due To Deadlock” on page 173</li></ul>	

**Description:** The total number of SQL ROLLBACK statements that have been attempted.

**Usage:** A rollback can result from an application request, a deadlock, or an error situation. This element **only** counts the number of rollback statements issued from applications.

At the application level, this element can help you determine the level of database activity for the application and the amount of conflict with other applications. At the database level, it can help you determine the amount of activity in the database and the amount of conflict between applications on the database.

**Note:** You should try to minimize the number of rollbacks, since higher rollback activity results in lower throughput for the database.

It may also be used to calculate the total number of units of work, by calculating the sum of the following:

```
commit statements attempted
+ internal commits
+ rollback statements attempted
+ internal rollbacks
```

## Select SQL Statements Executed

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
Table Space	sqlm_tablespace	Basic
Application	sqlm_appl	Basic
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
API Element Name	select_sql_stmts	
Element Type	counter	
Related Information	<ul style="list-style-type: none"> <li>• “When Counters are Initialized” on page 20</li> <li>• “Static SQL Statements Attempted” on page 164</li> <li>• “Dynamic SQL Statements Attempted” on page 165</li> </ul>	

**Description:** The number of SQL SELECT statements that were executed.

**Usage:** You can use this element to determine the level of database activity at the application or database level.

You can also use the following formula to determine the ratio of SELECT statements to the total statements:

$$\frac{\text{select SQL statements executed}}{(\text{static SQL statements attempted} + \text{dynamic SQL statements attempted})}$$

This information can be useful for analyzing application activity and throughput.

## Update/Insert/Delete SQL Statements Executed

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
API Element Name	uid_sql_stmts	
Element Type	counter	
Related Information	<ul style="list-style-type: none"> <li>• “When Counters are Initialized” on page 20</li> <li>• “Static SQL Statements Attempted” on page 164</li> <li>• “Dynamic SQL Statements Attempted” on page 165</li> </ul>	

**Description:** The number of SQL UPDATE, INSERT, and DELETE statements that were executed.

**Usage:** You can use this element to determine the level of database activity at the application or database level.

You can also use the following formula to determine the ratio of UPDATE, INSERT and DELETE statements to the total number of statements:

$$\frac{\text{update/insert/delete SQL statements executed}}{\text{(static SQL statements attempted + dynamic SQL statements attempted)}}$$

This information can be useful for analyzing application activity and throughput.

## Data Definition Language (DDL) SQL Statements

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
API Element Name	ddl_sql_stmts	
Element Type	counter	
Related Information	<ul style="list-style-type: none"> <li>• “When Counters are Initialized” on page 20</li> </ul>	

**Description:** This element indicates the number of SQL Data Definition Language (DDL) statements that were executed.

**Usage:** You can use this element to determine the level of database activity at the application or database level. DDL statements are expensive to run due to their impact on the system catalog tables. As a result, if the value of this element is high, you should determine the cause, and possibly restrict this activity from being performed.

You can also use this element to determine the percentage of DDL activity using the following formula:

$$\frac{\text{data definition language (DDL) SQL statements}}{\text{total number of statements}}$$

This information can be useful for analyzing application activity and throughput. DDL statements can also impact the package cache, by invalidating sections that are stored there and causing additional system overhead due to section recompilation.

Examples of DDL statements are CREATE TABLE, CREATE VIEW, ALTER TABLE, and DROP INDEX.

## Internal Automatic Rebinds

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
<b>API Element Name</b>	int_auto_rebinds	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “When Counters are Initialized” on page 20</li><li>• “Binds/Precompiles Attempted” on page 174</li></ul>	

**Description:** The number of automatic rebinds (or recompiles) that have been attempted.

**Usage:** Automatic rebinds are the internal binds the system performs when an package has been invalidated. The rebind is performed the first time that the database manager needs to execute an SQL statement from the package. For example, packages are invalidated when you:

- Drop an object, such as a table, view, or index, on which the plan is dependent
- Add or drop a foreign key
- Revoke object privileges on which the plan is dependent.

You can use this element to determine the level of database activity at the application or database level. Since internal automatic rebinds can have a significant impact on performance, they should be minimized where possible.

You can also use this element to determine the percentage of rebind activity using the following formula:

$$\text{internal automatic rebinds} / \text{total number of statements}$$

This information can be useful for analyzing application activity and throughput.

## Internal Commits

<b>Snapshot Information Level</b>	<b>API Structure(s)</b>	<b>Monitor Switch</b>
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
<b>API Element Name</b>	int_commits	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “When Counters are Initialized” on page 20</li><li>• “Commit Statements Attempted” on page 166</li><li>• “Rollback Statements Attempted” on page 167</li><li>• “Internal Rollbacks” on page 172</li></ul>	

**Description:** The total number of commits initiated internally by the database manager.

**Usage:** An internal commit may occur during any of the following:

- A reorganization
- An import
- A bind or pre-compile
- An application ends without executing an explicit SQL COMMIT statement (on UNIX).

This value, which does not include explicit SQL COMMIT statements, represents the number of these internal commits since the later of:

- The connection to the database (for database-level information, this is the time of the first connection)
- The last reset of the database monitor counters.

You can use this element to calculate the total number of units of work by calculating the sum of the following:

```
commit statements attempted
+ internal commits
+ rollback statements attempted
+ internal rollbacks
```

**Note:** The units of work calculated will only include those since the later of:

- The connection to the database (for database-level information, this is the time of the first connection)
- The last reset of the database monitor counters.

This calculation can be done at the application or the database level.

## Internal Rollbacks

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
<b>Resettable</b>	Yes	
<b>Event Type</b>	<b>Event Record(s)</b>	
Database	sqlm_db_event	
Connection	sqlm_conn_event	
<b>API Element Name</b>	int_rollbacks	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “When Counters are Initialized” on page 20</li><li>• “Commit Statements Attempted” on page 166</li><li>• “Internal Commits” on page 171</li><li>• “Rollback Statements Attempted” on page 167</li><li>• “Internal Rollbacks Due To Deadlock” on page 173</li></ul>	

**Description:** The total number of rollbacks initiated internally by the database manager.

**Usage:** An internal rollback occurs when any of the following **cannot** complete successfully:

- A reorganization
- An import
- A bind or pre-compile
- An application ends as a result of a deadlock situation or lock timeout situation
- An application ends without executing an explicit commit or rollback statement (on Windows).

This value represents the number of these internal rollbacks since the later of:

- The connection to the database (for database-level information, this is the time of the first connection)
- The last reset of the database monitor counters.

While this value does not include explicit SQL ROLLBACK statements, the count from Internal Rollbacks Due To Deadlock is included.

You can use this element to calculate the total number of units of work by calculating the sum of the following:

```
commit statements attempted
+ internal commits
+ rollback statements attempted
+ internal rollbacks
```

**Note:** The units of work calculated will include those since the later of:

- The connection to the database (for database-level information, this is the time of the first connection)
- The last reset of the database monitor counters.

This calculation can be done at the application or the database level.

## Internal Rollbacks Due To Deadlock

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
<b>Resettable</b>	Yes	
Event Type	Event Record(s)	
Connection	sqlm_conn_event	
API Element Name	int_deadlock_rollbacks	
Element Type	counter	
Related Information	<ul style="list-style-type: none"> <li>• “When Counters are Initialized” on page 20</li> <li>• “Deadlocks Detected” on page 133</li> <li>• “Rollback Statements Attempted” on page 167</li> <li>• “Internal Rollbacks” on page 172</li> </ul>	

**Description:** The total number of forced rollbacks initiated by the database manager due to a deadlock. A rollback is performed on the current unit of work in an application selected by the database manager to resolve the deadlock.

**Usage:** This element shows the number of deadlocks that have been broken and can be used as an indicator of concurrency problems. It is important, since internal rollbacks due to deadlocks lower the throughput of the database.

This value is included in the value given by Internal Rollbacks.

## SQL Requests Since Last Commit

Snapshot Information Level	API Structure(s)	Monitor Switch
Application	sqlm_appl	Basic
<b>Resettable</b>	No	
API Element Name	sql_reqs_since_commit	
Element Type	information	
Related Information	<ul style="list-style-type: none"> <li>• None</li> </ul>	

**Description:** Number of SQL requests that have been submitted since the last commit.

**Usage:** You can use this element to monitor the progress of a transaction.

**Note:** This element is similar to the *cur\_reqs* field in the *sqlestat* output. See Appendix C, “DB2 Version 1 sqlestat Users” on page 271 for more information on sqlestat equivalent data elements.

## Statement Node

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_stmt	<b>Monitor Switch</b> Statement
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	stmt_node_number information	
<b>Related Information</b>	• None	

**Description:** Node where the statement was executed.

**Usage:** Used to correlate each statement with the node where it was executed.

## Binds/Precompiles Attempted

<b>Snapshot Information Level</b> Database Application	<b>API Structure(s)</b> sqlm_dbase sqlm_appl	<b>Monitor Switch</b> Basic Basic
<b>Resettable</b>	Yes	
<b>Event Type</b> Database Connection	<b>Event Record(s)</b> sqlm_db_event sqlm_conn_event	
<b>API Element Name</b> <b>Element Type</b>	binds_precompiles counter	
<b>Related Information</b>	• “When Counters are Initialized” on page 20 • “Internal Automatic Rebinds” on page 170	

**Description:** The number of binds and pre-compiles attempted.

**Usage:** You can use this element to gain insight into the current level of activity within the database manager.

This value does not include the count of *Internal Automatic Rebinds*, but it does include binds that occur as a result of the REBIND PACKAGE command.

## SQL Statement Details

The following elements provide details about the SQL statements:

- “Statement Type” on page 175
- “Statement Operation” on page 176
- “Package Name” on page 177
- “Section Number” on page 177
- “Cursor Name” on page 178
- “Application Creator” on page 179
- “Statement Operation Start Timestamp” on page 179
- “Statement Operation Stop Timestamp” on page 180
- “SQL Dynamic Statement Text” on page 180
- “Statement Sorts” on page 181



- “Number of Successful Fetches” on page 182
- “SQL Communications Area (SQLCA)” on page 182
- “Query Number of Rows Estimate” on page 183
- “Query Cost Estimate” on page 183

## Statement Type

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_appl sqlm_stmt	<b>Monitor Switch</b> Statement Statement
<b>Resettable</b>	No	
<b>Event Type</b> Statement	<b>Event Record(s)</b> sqlm_stmt_event	
<b>API Element Name</b> <b>Element Type</b>	stmt_type information	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “SQL Dynamic Statement Text” on page 180</li> <li>• “Application Creator” on page 179</li> <li>• “Section Number” on page 177</li> <li>• “Package Name” on page 177</li> </ul>	

**Description:** The type of statement processed.

**Usage:** You can use this element to determine the type of statement that is executing. It can be one of the following:

- A static SQL statement
- A dynamic SQL statement
- An operation other than an SQL statement; for example, a bind or pre-compile operation.

For the snapshot monitor, this element describes the statement that is currently being processed or was most recently processed.

**Note:** API users should refer to the *sqlmon.h* header file containing definitions of database system monitor constants.

## Statement Operation

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_appl sqlm_stmt	<b>Monitor Switch</b> Basic Statement
<b>Resettable</b>	No	
<b>Event Type</b> Statement	<b>Event Record(s)</b> sqlm_stmt_event	
<b>API Element Name</b>	stmt_operation (Snapshot) operation (Event)	
<b>Element Type</b>	information	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “Statement Type” on page 175</li><li>• “SQL Dynamic Statement Text” on page 180</li><li>• “Application Creator” on page 179</li><li>• “Section Number” on page 177</li><li>• “Package Name” on page 177</li><li>• “Number of Successful Fetches” on page 182</li></ul>	

**Description:** The statement operation currently being processed or most recently processed (if none currently running).

**Usage:** You can use this element to determine the operation that is executing or recently finished.

It can be one of the following.

For SQL operations:

- SELECT
- PREPARE
- EXECUTE
- EXECUTE IMMEDIATE
- OPEN
- FETCH
- CLOSE
- DESCRIBE
- STATIC COMMIT
- STATIC ROLLBACK
- FREE LOCATOR

For non-SQL operations:

- RUN STATISTICS
- REORG
- REBIND
- REDISTRIBUTE
- GET TABLE AUTHORIZATION
- GET ADMINISTRATIVE AUTHORIZATION

**Note:** API users should refer to the *sqlmon.h* header file containing definitions of database system monitor constants.

## Package Name

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_appl sqlm_stmt	<b>Monitor Switch</b> Statement Statement
<b>Resettable</b>	No	
<b>Event Type</b> Statement	<b>Event Record(s)</b> sqlm_stmt_event	
<b>API Element Name</b> <b>Element Type</b>	package_name information	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “Application Creator” on page 179</li> <li>• “Section Number” on page 177</li> <li>• “SQL Dynamic Statement Text” on page 180</li> </ul>	

**Description:** The name of the package that contains the SQL statement currently executing.

**Usage:** You may use this element to help identify the application program and the SQL statement that is executing.

If you are using the database system monitor APIs, note that the API constant `SQLM_IDENT_SZ` is used to define the length of this element. Only the first 8 characters are currently used.

## Section Number

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_appl sqlm_stmt	<b>Monitor Switch</b> Statement Statement
<b>Resettable</b>	No	
<b>Event Type</b> Statement	<b>Event Record(s)</b> sqlm_stmt_event	
<b>API Element Name</b> <b>Element Type</b>	section_number information	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “SQL Dynamic Statement Text” on page 180</li> <li>• “Application Creator” on page 179</li> <li>• “Package Name” on page 177</li> </ul>	

**Description:** The internal section number in the package for the SQL statement currently processing or most recently processed.

**Usage:** For static SQL, you can use this element along with Application Creator and Package Name to query the SYSCAT.STATEMENTS system catalog table and obtain the static SQL statement text, using the sample query as follows:

```
SELECT SEQNO, SUBSTR(TEXT,1,120)
FROM SYSCAT.STATEMENTS
WHERE PKGNAME = 'package_name' AND
      PKGSCHEMA = 'creator' AND
      SECTNO = section_number
ORDER BY SEQNO
```

**Note:** Exercise caution in obtaining static statement text, because this query against the system catalog table could cause lock contentions. Whenever possible, only use this query when there is little other activity against the database.

## Cursor Name

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_appl sqlm_stmt	<b>Monitor Switch</b> Statement Statement
<b>Resettable</b>	No	
<b>Event Type</b> Statement	<b>Event Record(s)</b> sqlm_stmt_event	
<b>API Element Name</b> <b>Element Type</b>	cursor_name information	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “SQL Dynamic Statement Text” on page 180</li> <li>• “Statement Type” on page 175</li> <li>• “Number of Successful Fetches” on page 182</li> </ul>	

**Description:** The name of the cursor corresponding to this SQL statement.

**Usage:** You may use this element to identify the SQL statement that is processing. This name will be used on an OPEN, FETCH, CLOSE, and PREPARE of an SQL SELECT statement. If a cursor is not used, this field will be blank.

If you are using the database system monitor APIs, note that the API constant SQLM\_IDENT\_SZ is used to define the length of this element. Only the first 8 characters are currently used.

## Application Creator

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_appl sqlm_stmt	<b>Monitor Switch</b> Statement Statement
<b>Resettable</b>	No	
<b>Event Type</b> Statement	<b>Event Record(s)</b> sqlm_stmt_event	
<b>API Element Name</b> <b>Element Type</b>	creator information	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “Package Name” on page 177</li><li>• “Section Number” on page 177</li></ul>	

**Description:** The authorization ID of the user that pre-compiled the application.

**Usage:** You may use this element to help identify the SQL statement that is processing, in conjunction with the CREATOR column of the package section information in the catalogs.

If you are using the database system monitor APIs, note that the API constant SQLM\_IDENT\_SZ is used to define the length of this element. Only the first 8 characters are currently used.

## Statement Operation Start Timestamp

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_appl sqlm_stmt	<b>Monitor Switch</b> Statement Statement
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	stmt_start timestamp	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “Statement Operation Stop Timestamp” on page 180</li><li>• “Statement Operation” on page 176</li></ul>	

**Description:** The date and time when the Statement Operation started executing.

**Usage:** You can use this element with Statement Operation Stop Timestamp to calculate the elapsed statement operation execution time.

## Statement Operation Stop Timestamp

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_appl sqlm_stmt	<b>Monitor Switch</b> Statement Statement
<b>Resettable</b>	No	
<b>Event Type</b> Statement	<b>Event Record(s)</b> sqlm_stmt_event	
<b>API Element Name</b>	stmt_stop (Snapshot) stop_time (event)	
<b>Element Type</b>	Timestamp	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “Statement Operation Start Timestamp” on page 179</li><li>• “Statement Operation” on page 176</li></ul>	

**Description:** The date and time when the Statement Operation stopped executing.

**Usage:** You can use this element with Statement Operation Start Timestamp to calculate the elapsed statement operation execution time.

## SQL Dynamic Statement Text

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_stmt	<b>Monitor Switch</b> Statement
<b>Resettable</b>	No	
<b>Event Type</b> Statement	<b>Event Record(s)</b> sqlm_stmt_event	
<b>API Element Name</b>	Relative offsets are used to return the text. See data structures in sqlmon.h>	
<b>Element Type</b>	stmt_text (Event) information	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “Statement Operation” on page 176</li><li>• “Cursor Name” on page 178</li><li>• “Input Database Alias” on page 193</li><li>• “Application Creator” on page 179</li><li>• “Package Name” on page 177</li><li>• “Section Number” on page 177</li></ul>	

**Description:** This is the text of the dynamic SQL statement.

**Usage:** For snapshots, this statement text helps you identify what the application was executing when the snapshot was taken, or most recently processed if no statement was being processed right at the time the snapshot was taken.

For event monitors, it is returned in the Statement event record for all dynamic statements.

See Section Number for information on how to query the system catalog tables to obtain static SQL statement text that is not provided due to performance considerations.

## Statement Sorts

Snapshot Information Level	API Structure(s)	Monitor Switch
Application	sqlm_appl sqlm_stmt	Statement Statement
<b>Resettable</b>	No	
<b>API Element Name</b>	stmt_sorts	
<b>Element Type</b>	counter	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “Total Sorts” on page 82</li> </ul>	

**Description:** The total number of times that a set of data was sorted in order to process the statement operation.

**Usage:** You can use this element to help identify the need for an index, since indexes can reduce the need for sorting of data. Using the related elements in the above table you can identify the SQL statement for which this element is providing sort information, and then analyze this statement to determine index candidates by looking at columns that are being sorted (for example, columns used in ORDER BY and GROUP BY clauses and join columns). See **explain** in the *Administration Guide* for information on checking whether your indexes are used to optimize sort performance.

This count includes sorts of temporary tables that were generated internally by the database manager to execute the statement. The number of sorts is associated with the first FETCH operation of the SQL statement. This information is returned to you when the operation for the statement is the first FETCH. You should note that for blocked cursors several fetches may be performed when the cursor is opened. In these cases it can be difficult to use the snapshot monitor to obtain the number of sorts, since a snapshot would need to be taken while DB2 was internally issuing the first FETCH.

A more reliable way to determine the number of sorts performed when using a blocked cursor would be with an event monitor declared for statements. The total sorts counter, in the statement event for the CLOSE cursor, contains the total number of sorts that were performed while executing the statement for which the cursor was defined.

## Number of Successful Fetches

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_stmt	<b>Monitor Switch</b> Statement
<b>Resettable</b>	Yes	
<b>Event Type</b> Statement	<b>Event Record(s)</b> sqlm_stmt_event	
<b>API Element Name</b> <b>Element Type</b>	fetch_count counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Statement Type” on page 175</li><li>• “Statement Operation” on page 176</li><li>• “Cursor Name” on page 178</li><li>• “Statement Operation Start Timestamp” on page 179</li><li>• “Statement Operation Stop Timestamp” on page 180</li></ul>	

**Description:** The number of successful fetches performed on a specific cursor.

**Usage:** You can use this element to gain insight into the current level of activity within the database manager.

For performance reasons, a statement event monitor does not generate a statement event record for every FETCH statement. A record event is only generated when a FETCH returns a non-zero SQLCODE.

## SQL Communications Area (SQLCA)

<b>Event Type</b> Statement	<b>Event Record(s)</b> sqlm_stmt_event
<b>API Element Name</b> <b>Element Type</b>	sqlca information
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Statement Operation” on page 176</li></ul>

**Description:** The SQLCA data structure that was returned to the application at statement completion.

**Usage:** The SQLCA data structure can be used to determine if the statement completed successfully. See the *SQL Reference* or *API Reference* for information about the content of the SQLCA.



## Query Number of Rows Estimate

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_stmt	<b>Monitor Switch</b> Statement
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	query_card_estimate information	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “Query Cost Estimate” on page 183</li></ul>	

**Description:** An estimate of the number of rows that will be returned by a query.

**Usage:** This estimate by the SQL compiler can be compared with the run time actuals.

## Query Cost Estimate

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_stmt	<b>Monitor Switch</b> Statement
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	query_cost_estimate information	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• None</li></ul>	

**Description:** Estimated cost, in timerons, for a query, as determined by the SQL compiler.

**Usage:** This allows correlation of actual run-time with the compile-time estimates.

## Subsection Details

When a statement is executed against a partitioned database, it is divided into subsections that may be executed on different nodes. An application may have several subsections simultaneously executing on a node. See “Monitoring Subsections” on page 26 and the *Administration Guide* for more information on subsections.

For problem determination, you may have to locate the problem subsection. For example, a subsection may be waiting on a tablequeue, because one of the writers to this tablequeue is in lock wait on another node. To get the overall picture for an application, you may have to issue an application snapshot on each node where the application is running.

The following database system monitor elements provide information about Subsections:

- “Subsection Number” on page 184
- “Subsection Node Number” on page 184
- “Subsection Status” on page 185
- “Execution Elapsed Time” on page 185
- “Number of Agents Working on a Subsection” on page 186

- “Waiting for Any Node to Send on a Tablequeue” on page 186
- “Waited for Node on a Tablequeue” on page 187
- “Total Number of Tablequeue Buffers Overflowed” on page 187
- “Current Number of Tablequeue Buffers Overflowed” on page 188
- “Number of Rows Read from Tablequeues” on page 188
- “Number of Rows Written to Tablequeues” on page 189

## Subsection Number

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_subsection	<b>Monitor Switch</b> Statement
<b>Resettable</b>	No	
<b>Event Type</b> Statement	<b>Event Record(s)</b> sqlm_subsection_event	
<b>API Element Name</b> <b>Element Type</b>	ss_number information	
<b>Related Information</b>	• None	

**Description:** Identifies the subsection associated with the returned information.

**Usage:** This number relates to the subsection number in the access plan that can be obtained with db2expln (see *Administration Guide*).

## Subsection Node Number

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_subsection	<b>Monitor Switch</b> Statement
<b>Resettable</b>	No	
<b>Event Type</b> Statement	<b>Event Record(s)</b> sqlm_subsection_event	
<b>API Element Name</b> <b>Element Type</b>	ss_node_number information	
<b>Related Information</b>	• None	

**Description:** Node where the subsection was executed.

**Usage:** Use to correlate each subsection with the database partition where it was executed.

## Subsection Status

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_subsection	<b>Monitor Switch</b> Statement
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	ss_status information	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Waited for Node on a Tablequeue” on page 187</li><li>• “Waiting for Any Node to Send on a Tablequeue” on page 186</li></ul>	

**Description:** The current status of an executing subsection.

**Usage:** The current status values can be:

- executing
- waiting for a lock
- waiting to receive data on a tablequeue
- waiting to send data on a tablequeue

## Execution Elapsed Time

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_subsection	<b>Monitor Switch</b> Statement
<b>Resettable</b>	No	
<b>Event Type</b> Statement	<b>Event Record(s)</b> sqlm_subsection_event	
<b>API Element Name</b> <b>Element Type</b>	ss_exec_time counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• None</li></ul>	

**Description:** The time in seconds that it took a subsection to execute.

**Usage:** Allows you to track the progress of a subsection.

## Number of Agents Working on a Subsection

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_subsection	<b>Monitor Switch</b> Statement
<b>Resettable</b>	No	
<b>Event Type</b> Statement	<b>Event Record(s)</b> sqlm_subsection_event	
<b>API Element Name</b> <b>Element Type</b>	num_subagents gauge	
<b>Related Information</b>	• None	

**Description:** Total number of subagents currently working on a subsection.

**Usage:** Indicates the current degree of parallelism. Helps you track how execution is progressing.

## Waiting for Any Node to Send on a Tablequeue

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_subsection	<b>Monitor Switch</b> Statement
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	tq_wait_for_any information	
<b>Related Information</b>	• "Subsection Status" on page 185 • "Waited for Node on a Tablequeue" on page 187	

**Description:** This flag is used to indicate that the subsection is blocked because it is waiting to receive rows from any node.

**Usage:** If Subsection Status indicates *waiting to receive data on a tablequeue* and this flag is TRUE, then the subsection is waiting to receive rows from any node. This generally indicates that the SQL statement has not processed to the point it can pass data to the waiting agent. For example, the writing agent may be performing a sort and will not write rows until the sort has completed. From the db2expln output, determine the subsection number associated with the tablequeue that the agent is waiting to receive rows from. You can then examine the status of that subsection by taking a snapshot on each node where it is executing.

## Waited for Node on a Tablequeue

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_subsection	<b>Monitor Switch</b> Statement
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	tq_node_waited_for information	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Subsection Status” on page 185</li><li>• “Waiting for Any Node to Send on a Tablequeue” on page 186</li></ul>	

**Description:** If the subsection status Subsection Status is *waiting to receive* or *waiting to send* and Waiting for Any Node to Send on a Tablequeue is FALSE, then this is the number of the node that this agent is waiting for.

**Usage:** This can be used for troubleshooting. You may want to take an application snapshot on the node that the subsection is waiting for. For example, the application could be in a lock wait on that node.

## Total Number of Tablequeue Buffers Overflowed

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_subsection	<b>Monitor Switch</b> Statement
<b>Resettable</b>	No	
<b>Event Type</b> Statement	<b>Event Record(s)</b> sqlm_subsection_event	
<b>API Element Name</b> <b>Element Type</b>	tq_tot_send_spills counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Subsection Status” on page 185</li><li>• “Current Number of Tablequeue Buffers Overflowed” on page 188</li></ul>	

**Description:** Total number of tablequeue buffers overflowed to a temporary table.

**Usage:** Indicates the total number of tablequeue buffers that have been written to a temporary table. See “Current Number of Tablequeue Buffers Overflowed” on page 188 for more information.

## Current Number of Tablequeue Buffers Overflowed

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_subsection	<b>Monitor Switch</b> Statement
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	tq_cur_send_spills gauge	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Subsection Status” on page 185</li><li>• “Total Number of Tablequeue Buffers Overflowed” on page 187</li></ul>	

**Description:** Current number of tablequeue buffers residing in a temporary table.

**Usage:** An agent writing to a tablequeue may be sending rows to several readers. The writing agent will overflow buffers to a temporary table when the agent that it is currently sending rows to is not accepting rows and another agent requires rows in order to proceed. Overflowing to temporary table allows both the writer and the other readers to continue processing.

Rows that have been overflowed will be sent to the reading agent when it is ready to accept more rows.

If this number is high, and queries fail with sqlcode -968, and there are messages in *db2diad.log* indicating that you ran out of temporary space in the TEMP table space, then tablequeue overflows may be the cause. This could indicate a problem on another node (such as locking). You would investigate by taking snapshots on all the partitions for this query.

There are also cases, perhaps because of the way data is partitioned, where many buffers need to be overflowed for the query. In these cases you will need to add more disk to the temporary table space.

## Number of Rows Read from Tablequeues

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_subsection	<b>Monitor Switch</b> Statement
<b>Resettable</b>	No	
<b>Event Type</b> Statement	<b>Event Record(s)</b> sqlm_subsection_event	
<b>API Element Name</b> <b>Element Type</b>	tq_rows_read counter	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• None</li></ul>	

**Description:** Total number of rows read from tablequeues.

**Usage:** If monitoring does not indicate that this number is increasing, then processing progress is not taking place.

If there is significant differences in this number between nodes, then some nodes may be over utilized while others are being under utilized.

If this number is large, then there is a lot of data being shipped between nodes, suggest that optimization might improve the access plan.

### Number of Rows Written to Tablequeues

<b>Snapshot Information Level</b> Application	<b>API Structure(s)</b> sqlm_subsection	<b>Monitor Switch</b> Statement
<b>Resettable</b>	No	
<b>Event Type</b> Statement	<b>Event Record(s)</b> sqlm_subsection_event	
<b>API Element Name</b> <b>Element Type</b>	tq_rows_written counter	
<b>Related Information</b>	• None	

**Description:** Total number of rows written to tablequeues.

**Usage:** If monitoring does not indicate that this number is increasing, then processing progress is not taking place.

If there is significant differences in this number between nodes, then some nodes may be over utilized while others are being under utilized.

If this number is large, then there is a lot of data being shipped between nodes, suggest that optimization might improve the access plan.

### Intra-query Parallelism

The following database system monitor elements provide information about queries for which the degree of parallelism is greater than 1:

- “Number of Agents Working on a Statement” on page 190
- “Number of Agents Created” on page 190
- “Degree of Parallelism” on page 191

## Number of Agents Working on a Statement

Snapshot Information Level	API Structure(s)	Monitor Switch
Statement	sqlm_stmt sqlm_subsection	Statement Statement
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	num_agents gauge	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Number of Agents Created” on page 190</li><li>• “Degree of Parallelism” on page 191</li></ul>	

**Description:** Number of concurrent agents currently executing a statement or subsection.

**Usage:** An indicator how well the query is parallelized. This is useful for tracking the progress of query execution, by taking successive snapshots.

## Number of Agents Created

Snapshot Information Level	API Structure(s)	Monitor Switch
Database Application	sqlm_dbase sqlm_stmt	Statement Statement
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	agents_top water mark	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Number of Agents Working on a Statement” on page 190</li><li>• “Degree of Parallelism” on page 191</li></ul>	

**Description:** This is the maximum number of agents that were used when executing the statement.

**Usage:** An indicator how well intra-query parallelism was realized.



## Degree of Parallelism

<b>Snapshot Information Level</b> Statement	<b>API Structure(s)</b> sqlm_stmt	<b>Monitor Switch</b> Statement
<b>Resettable</b>	No	
<b>API Element Name</b> <b>Element Type</b>	degree_parallelism information	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Number of Agents Working on a Statement” on page 190</li> <li>• “Number of Agents Created” on page 190</li> </ul>	

**Description:** The degree of parallelism requested when the query was bound.

**Usage:** Use with “Number of Agents Created” on page 190, to determine if the query achieved maximum level of parallelism.

## CPU Usage

The CPU usage for an application is broken down into **user CPU**, which is the CPU consumed while executing application code, and **system CPU**, which is the CPU consumed executing system calls.

CPU consumption is available at the application, transaction, statement, and subsection levels.

- “CPU Time Used”

## CPU Time Used

<b>Snapshot Information Level</b> Application Statement Subsection	<b>API Structure(s)</b> sqlm_appl sqlm_stmt sqlm_subsection	<b>Monitor Switch</b> Basic Statement Statement
<b>Resettable</b>	Yes, at the application level No, at other levels	
<b>Event Type</b> Connection Transaction Statement Subsection	<b>Event Record(s)</b> sqlm_conn_event sqlm_xaction_event sqlm_stmt_event sqlm_subsection_event	
<b>API Element Name</b>	agent_usr_cpu_time agent_sys_cpu_time stmt_usr_cpu_time stmt_sys_cpu_time ss_usr_cpu_time ss_sys_cpu_time user_cpu_time system_cpu_time	
<b>Element Type</b>	time	
<b>Related Elements</b>	<ul style="list-style-type: none"> <li>• None</li> </ul>	

**Description:** The total CPU time (in seconds and microseconds) used by database manager agents, while working on behalf of the application, transaction, statement, or subsection.

System CPU represents the time spent in system calls. User CPU represents time spent executing database manager code.

These counters include time spent on both SQL and non-SQL statements, as well as any fenced user defined functions (UDF) or stored procedures executed by the application.

**Usage:** These elements can help you identify applications or queries that consume large amounts of CPU.

**Note:** If this information is not available for your operating system, these elements will be returned as 0. For example, they are not available on OS/2.

## Snapshot Monitoring Elements

The following elements provide information about monitoring applications. They are returned as output for every snapshot:

- “Last Reset Timestamp”
- “Input Database Alias” on page 193
- “Snapshot Time” on page 193

### Last Reset Timestamp

Snapshot Information Level	API Structure(s)	Monitor Switch
Database Manager	sqlm_db2	Basic
Database	sqlm_dbase	Basic
Application	sqlm_appl	Basic
Table Space	sqlm_tablespace_header	Buffer Pool
Table	sqlm_table_header	Table
<b>Resettable</b>	No	
<b>API Element Name</b>	last_reset	
<b>Element Type</b>	timestamp	
<b>Related Information</b>	<ul style="list-style-type: none"> <li>• “Resetting Monitor Data” on page 21</li> <li>• “Input Database Alias” on page 193</li> </ul>	

**Description:** Indicates the date and time that the monitor counters were reset for the application issuing the GET SNAPSHOT.

**Usage:** You can use this element to help you determine the scope of information returned by the database system monitor.

If the counters have never been reset, this element will be zero.

The database manager counters will only be reset if you reset all active databases.

## Input Database Alias

Snapshot Information Level	API Structure(s)	Monitor Switch
Database	sqlm_dbase	Basic
Application	sqlm_appl_id_info	Basic
Table Space	sqlm_tablespace_header	Buffer Pool
Buffer Pool	sqlm_bufferpool	Buffer Pool
Table	sqlm_table_header	Table
Lock	sqlm_dbase_lock	Basic
<b>Resettable</b>	No	
<b>API Element Name</b>	input_db_alias	
<b>Element Type</b>	information	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• “Resetting Monitor Data” on page 21</li><li>• “Last Reset Timestamp” on page 192</li><li>• “Database Alias Used by Application” on page 54</li></ul>	

**Description:** The alias of the database provided when calling the snapshot function.

**Usage:** This element can be used to identify the specific database to which the monitor data applies. It contains blanks unless you requested monitor information related to a specific database.

The value of this field may be different than the value of the *Database Alias Used by Application* monitor element since a database can have many different aliases. Different applications and users can use different aliases to connect to the same database.

If you are using the database system monitor APIs, note that the API constant SQLM\_IDENT\_SZ is used to define the length of this element. Only the first 8 characters are currently used.

## Snapshot Time

Snapshot Information Level	API Structure(s)	Monitor Switch
Database Manager	sqlm_collected	Basic
<b>Resettable</b>	No	
<b>API Element Name</b>	time_stamp	
<b>Element Type</b>	timestamp	
<b>Related Information</b>	<ul style="list-style-type: none"><li>• None</li></ul>	

**Description:** The date and time when the database system monitor information was collected.

**Usage:** You can use this element to help relate data chronologically if you are saving the results in a file or database for ongoing analysis.



---

## Chapter 4. Event Monitor Output

This chapter explains the contents and format of the trace produced by an event monitor and different options that can be specified on the CREATE EVENT MONITOR statement that can influence the trace. It shows how to program for reading this trace, through the use of code samples.

---

### Output Stream Format

The output of an event monitor is a binary stream of data structures that are exactly the same for both pipe and file event monitors. You can format this trace using the db2evmon productivity tool.

Event Monitor records are defined in the **sqlmon.h** header file. You can look at the comments included in that file to see exactly which data elements are returned for each event type.

The following table illustrates the order in which records may appear in the event monitor stream. See "Information Available from Event Monitors" on page 17 for a list of events that trigger the writing of event records. Records in a trace are logically divided into three sections:

1. Prologue records - generated when an event monitor is activated.
2. Actual content records - generated as events occur.
3. Epilogue records - generated when a database is deactivated.

Record type	Record name	Information returned
<b>Prologue</b>		
Event Log Header	sqlm_event_log_header	Characteristics of the trace, for example server type and memory layout.
Database Header	sqlm_dbheader_event	Database name, path and activation time.
Event Monitor Start	sqlm_evmon_start_event	Time when the monitor was started or restarted.
Connection Header	sqlm_connheader_event	One for each current connection, includes connection time and application name.
<b>Actual Contents</b> (may appear mixed in with other connections).		
Connection Header	sqlm_connheader_event	One for each connection after activation, includes connection time and application identification.
Statement Event	sqlm_stmt_event	Statement level data, including text for dynamic statements.
Transaction Event	sqlm_xaction_event	Transaction level data.
Connection Event	sqlm_conn_event	Connection level data.
Deadlock Event	sqlm_deadlock_event	Deadlock level data.

Record type	Record name	Information returned
Deadlocked Connection Event	sqlm_dlconn_event	One for each connection involved, includes applications involved and locks in contention.
Overflow	sqlm_overflow_event	Number of records lost - generated when reader cannot keep up with a (non-blocked) event monitor.
<b>Epilogue</b>		
Database Event	sqlm_db_event	Database level data.
Buffer Pool Event	sqlm_bufferpool_event	Buffer pool level data.
Table Space Event	sqlm_tablespace_event	Table space level data.
Table Event	sqlm_table_event	Table level data.

**Note:** Event records may be generated for any connection and may therefore appear in mixed order in the stream. This means that you may get a transaction event for Connection 1, immediately followed by a connection event for Connection 2. However, records belonging to a single connection or a single event, will appear in their logical order. For example, a statement record (end of statement) always precedes a transaction record (end of UOW), if any. Similarly, a deadlock event record always precedes the deadlocked connection event records for each connection involved in the deadlock. The **application id** or **application handle (agent\_id)** can be used to match records with a connection.

For example, using the following event monitor,

```
db2 connect to sample
db2 "create event monitor ALL for
statements, transactions, connections,
deadlocks, database, bufferpools,
tablespaces, tables, write to
file '/tmp/all'"
mkdir /tmp/all
db2 connect reset
```

the following workload,

**Application 1**

```
db2 set event monitor ALL state 1
db2 select evmonname from
syscat.eventmonitors
db2 connect reset
```

**Application 2**

```
db2 connect to sample
db2 +c connect reset
```

the following trace might be generated. Listed in this sample are some of the fields in each event record to give a flavor of the type of information contained in a trace. See “Event Monitors” on page 10 for an example of deadlock events. Note, the numbers in this sample are used to illustrate the order in which records have been written.

## PROLOGUE

The Prologue information is generated when set event monitor all state 1 is executed. If this event monitor had been AUTOSTART, it would have been generated when the database was activated.

```
1) sqlm_event_log_header
   version:          SQLM_DBMON_VERSION5 - Trace was produced by UDB V5
   num_nodes_in_db2_instance: 1         - for a standalone system,
   byte_order:       SQLM_BIG_ENDIAN    - on a UNIX or AIX box,
   event_monitor_name: ALL              - by event monitor: 'ALL'

2) sqlm_dbheader_event
   db_name:          SAMPLE              - for database 'SAMPLE'

3) sqlm_connheader_event
   agent_id: 14                - Application 1 - handle
   appl_id: *LOCAL.bourbon.970602180712 - Application 1 - id with timestamp
```

## CONTENTS

Generated when Application 1 issues select name from syscat.eventmonitors. At the time that the event monitor is turned on, Application 2 has not yet connected.

```

4) sqlm_stmt_event
   agent_id: 14
   appl_id: *LOCAL.bourbon.970602180712
   operation:      SQLM_PREPARE
   package_name:   SQLC2BA4
   cursor:         SQLCUR201
   @stmt_text_offset: SELECT EVMONNAME FROM SYSCAT.EVENTMONITORS

5) sqlm_stmt_event
   agent_id: 14
   appl_id: *LOCAL.bourbon.970602180712
   operation:      SQLM_OPEN
   package_name:   SQLC2BA4
   cursor:         SQLCUR201
   @stmt_text_offset: SELECT EVMONNAME FROM SYSCAT.EVENTMONITORS

6) sqlm_stmt_event
   agent_id: 14
   appl_id: *LOCAL.bourbon.970602180712
   operation:      SQLM_FETCH
   package_name:   SQLC2BA4
   cursor:         SQLCUR201
   @stmt_text_offset: SELECT EVMONNAME FROM SYSCAT.EVENTMONITORS
   fetch_count:    2
   sqlca.sqlcode:  100 - (all rows in the SYSCAT.EVENTMONITORS table)
   SQL0100W No row was found for FETCH, UPDATE or DELETE; or the result of a
   query is an empty table.  SQLSTATE=02000
NOTE - A fetch event is generated only if the fetch fails or encounters end of table

7) sqlm_stmt_event
   agent_id: 14
   appl_id: *LOCAL.bourbon.970602180712
   operation:      SQLM_DESCRIBE
   package_name:   SQLC2BA4
   cursor:         SQLCUR201
   @stmt_text_offset: SELECT EVMONNAME FROM SYSCAT.EVENTMONITORS

8) sqlm_stmt_event
   agent_id: 14
   appl_id: *LOCAL.bourbon.970602180712
   operation:      SQLM_CLOSE
   package_name:   SQLC2BA4
   cursor:         SQLCUR201
   @stmt_text_offset: SELECT EVMONNAME FROM SYSCAT.EVENTMONITORS
   fetch_count:    2

9) sqlm_stmt_event
   agent_id: 14
   appl_id: *LOCAL.bourbon.970602180712
   operation:      SQLM_STATIC_COMMIT - generated by CLP after the SELECT
   package_name:   SQLC2BA4

10) sqlm_xaction_event
   agent_id: 14
   appl_id: *LOCAL.bourbon.970602180712
   status:        SQLM_UOWCOMMIT
   rows_read:     7

```



Application 2 is connecting to the database. Output is interleaved, as the DB2 agents are executing simultaneously:

```
11) sqlm_connheader_event
    agent_id: 15
    appl_id: *LOCAL.bourbon.970602180714 - Application 2 - handle
                                         - Application 2 - id with timestamp

12) sqlm_stmt_event
    agent_id: 15
    appl_id: *LOCAL.bourbon.970602180714
    operation:  SQLM_STATIC_COMMIT - generated by CLP on CONNECT

13) sqlm_xaction_event
    agent_id: 15
    appl_id: *LOCAL.bourbon.970602180714
    status:     SQLM_UOWCOMMIT

14) sqlm_stmt_event
    agent_id: 15
    appl_id: *LOCAL.bourbon.970602180714
    operation:  SQLM_STATIC_COMMIT - generated on CONNECT RESET

15) sqlm_xaction_event
    agent_id: 15
    appl_id: *LOCAL.bourbon.970602180714
    status:     SQLM_UOWCOMMIT

16) sqlm_conn_event
    agent_id: 15
    appl_id: *LOCAL.bourbon.970602180714
    commit_sql_stmts: 2

17) sqlm_stmt_event
    agent_id: 14
    appl_id: *LOCAL.bourbon.970602180712
    operation:  SQLM_STATIC_COMMIT - generated on CONNECT RESET
    package_name: SQLC2BA4

18) sqlm_xaction_event
    agent_id: 14
    appl_id: *LOCAL.bourbon.970602180712
    status:     SQLM_UOWCOMMIT
    rows_read: 2
    locks_held_top: 7

19) sqlm_conn_event
    agent_id: 14
    appl_id: *LOCAL.bourbon.970602180712
    select_sql_stmts: 1
    rows_selected: 2
```

## Epilogue

The Epilogue information is generated during database deactivation (last application finished disconnecting):

```
20) sqlm_table_event
    table_schema: SYSIBM
    table_name: SYSTABLES
    table_type: SQLM_CATALOG_TABLE
    rows_read: 2

21) sqlm_table_event
    table_schema: SYSIBM
    table_name: SYSDBAUTH
    table_type: SQLM_CATALOG_TABLE
    rows_read: 3

22) sqlm_tablespace_event
    tablespace_name: SYSCATSPACE

23) sqlm_tablespace_event
    tablespace_name: TEMPSPACE1

24) sqlm_tablespace_event
    tablespace_name: USERSPACE1

25) sqlm_bufferpool_event
    bp_name: IBMDEFAULTBP

26) sqlm_db_event
    connections_top: 2
```

**Note:** A WHERE clause on the CREATE EVENT MONITOR SQL statement can be used to restrict the applications that will generate events; see Appendix A, “Database System Monitor Interfaces” on page 213 for details.

---

## Matching Event Records with Their Application

Each record includes the application handle and application ID. These allow you to correlate each record with the application for which the record was generated.

The application handle (**agent\_id**) is unique system-wide for the duration of the application. However, it will eventually be reused (a 16 bit counter is used to generate this identifier). In most cases, this reuse is not a problem, since an application reading records from the trace is able to detect a connection that was terminated. For example, encountering (in the trace) a connection header with a known agent\_ID implies that the previous connection with this agent\_ID was terminated.

The application ID is a string identifier that includes a timestamp and is guaranteed to remain unique, even after stopping and restarting the database manager.

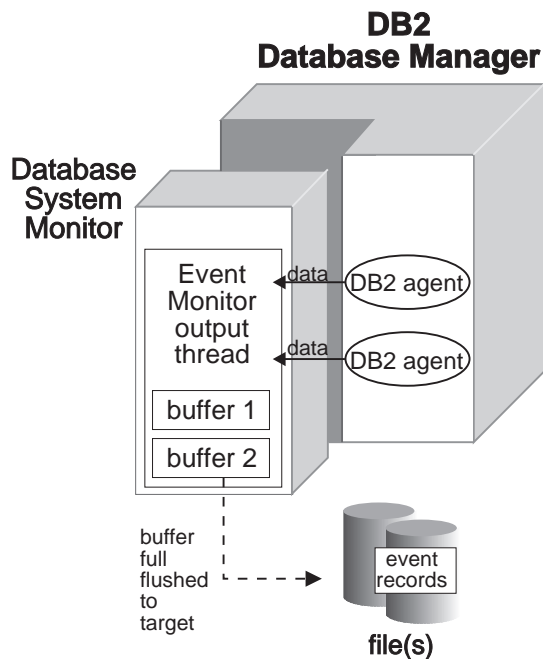
---

## File Event Monitor Buffering

The event monitor output thread buffers records, using two internal buffers, before writing them to disk. Records are written to the trace only when a buffer is full. To force an event monitor to flush its buffers you must turn it off. The size of these buffers can be specified on the CREATE EVENT MONITOR statement with the BUFFERSIZE argument. Specifying larger buffers reduces the number of disk accesses.

Figure 5 illustrates how event records are generated for a FILE statement event monitor: 2 applications are connected to a database, each having a single agent working on its behalf.

---



---

Figure 5. Event Monitor Buffers

In this example, each application agent has just finished executing a statement and is reporting the monitor data it has collected for its statement to the event monitor output thread. The output thread formats the records and writes them into one of its two buffers. The buffer gets written to a file when it is full. Having two buffers allows the output thread to continue receiving data from database agents, while a buffer is being written.

## Blocked Event Monitors

A blocked event monitor will suspend the agent(s) sending monitor data, when both of its buffers are full, until a buffer has been written. This can introduce a significant performance overhead, depending on the type of workload and the speed of the I/O device. But, a blocked event monitor never discards event records, as long as it is running. This is the default.

## Non-Blocked Event Monitors

A non-blocked event monitor will simply discard monitor data coming from the agents when the data is coming faster than it can write it. The following is an example of creating a non-blocked event monitor:

```
db2 "create event monitor STMT for
statements write to file '/tmp/all'
NONBLOCKED"
```

## Overflows

An event monitor that has discarded event records generates an **overflow event**. It specifies the start and stop time during which the monitor was discarding events, and the number of events that were discarded during that period.

**Unwritten Overflow Data:** It is possible for an event monitor to terminate or be deactivated with a pending overflow to report. If this occurs, the following message is written to the db2diag.log:

```
DIA1603I Event Monitor monitor-name had a pending overflow
record when it was deactivated.
```

## File Event Monitor Target

All the output of the event monitor goes in the directory supplied to the FILE argument on the CREATE EVENT MONITOR statement.

When a file event monitor is first activated, a control file is created in this directory. This binary file contains control information that is used to prevent two event monitors from simultaneously writing to the same target, and to keep track of the file and file location where the event monitor is supposed to write its next record. It is named *db2event.ctf*; do not remove or modify this file.

## Limiting Trace Size

By default, an event monitor writes its trace to a single file, called *00000000.evt*. This file will keep growing as long as there is space on the file system. You can limit the maximum size of a trace using the MAXFILESIZE and MAXFILES arguments of the create event monitor statement.

**Number of Files:** The trace produced by an event monitor can be quite large, and you may want to break it down into several files of a fixed size. This also allows you to remove files after processing them, while the event monitor is still running.

Files are numbered sequentially, beginning with `00000000.evt`. If you are using several files, then when a file is full, output is automatically directed to the next file. For example, the following event monitor will break down its trace into 4MB files. It keeps creating files as long as there is space on the file system.

```
db2 "create event monitor BIGONE
for statements, transactions, connections,
deadlocks write to file '/tmp/bigevmon'
MAXFILESIZE 1000
MAXFILES NONE"
```

This might result in the following files in its target directory.

File	size (bytes)
/tmp/bigevmon/db2evmon.ctl	300
/tmp/bigevmon/00000000.evt	4079766
/tmp/bigevmon/00000001.evt	4095128
/tmp/bigevmon/00000002.evt	4095602

The highest numbered file is always the active file. When the number of files reaches the maximum defined by `MAXFILES`, the event monitor deactivates itself and the following message is written to the `DB2DIAG.LOG`.

```
DIA1601I Event Monitor monitor-name was deactivated when
it reached its preset MAXFILES and MAXFILESIZE limit
```

You can avoid this situation by removing full files (see “Processing Data While Monitor is Active”). Any event file except the active file can be removed while the event monitor is still running.

## Running out of Disk Space

When a File event monitor runs out of disk space, it shuts itself down, after logging a system-error-level message in the error logs, `db2diag.log` and `db2err.log`.

## Processing Data While Monitor is Active

You may want an event monitor to collect data continuously so that no events are ever missed. For example, if you have a usage account system that uses an event monitor to collect data, you may want to process the data each night beginning at 2:00 AM, at which point you delete the files that have been processed.

An event monitor cannot be forced to switch to the next file unless you stop and restart it. It must also be in `APPEND` mode. In order to keep track of which events have been

processed in the active file, you can create an application that simply keeps track of the file number and location of the last record processed. When processing the trace the next time around, the application can simply seek to that file location.

Using Pipe event monitors is an easy way to read data produced by an active event monitor (see “Using Pipe Event Monitors” on page 18).

### Restarting a File Event Monitor

When a File event monitor is restarted, it can either erase any existing data, or append to it.

An APPEND event monitor starts writing at the end of the file it was last using (the file number is indicated in the *db2evmon.ctl* control file). If you have removed that file, then the next file number in sequence is used. For example, in the example above, if you remove all *.evt* files, and restart the event monitor, then event records will be written into *00000003.evt*. If you had not removed the files, then they would go into or append to *00000002.evt*. When an append event monitor is restarted, only the *start\_event* is generated. The event log header and database header are only generated for the first activation.

A REPLACE event monitor always deletes existing event files, and starts writing at *00000000.evt*.

---

## Programming to Read an Event Monitor Trace

Each record in the binary event monitor stream, except for the log header, starts with the size and type of the record. While reading the trace, it is extremely important that the size of a record is used for skipping a record in the trace, both to ensure that your application will be able to handle the traces produced by future releases of DB2, and because byte padding is sometimes used in the output stream. **WARNING: You should never use `sizeof()` on event monitor records.** Similarly, the type of every record should be checked. Skipping unknown or unwanted records will allow your application to handle any event monitor trace.

The log header describes the characteristics of the trace, containing information such as the memory model (for example big endian) of the server where the trace was collected, and the codepage of the database. You may have to do byte swapping on numerical values, if the system where you read the trace has a different memory model than the server (for example, Windows NT to UNIX). Codepage translation may also need to be done, if the database is configured in a different language than the machine from which you read the trace.

The following annotated fragments from the code samples in *sqlib/samples/c\_AIX* illustrates the most important considerations in programming to read an event monitor trace. Some error handling is omitted, for simplicity. This code should run on all platforms, except for PIPE I/O routines (however, more platforms are addressed in the actual code samples).

## Reading the Data Stream

The following routines illustrate how you can open, read, or skip bytes from a PIPE or FILE on a UNIX platform.

```
//-----  
// File functions - Using the ANSI C library  
//-----  
#include <stdio.h>  
#include <stdlib.h>  
#include <errno.h>  
//-----  
FILE* openFile(char *file_name) {  
    return fopen(file_name,"rb"); /* returns NULL if file cannot be opened */  
}  
//-----  
int closeFile(FILE* handle) {  
    return fclose(handle);  
}  
//-----  
int readFromFile(char* buffer, int size, FILE* fp) {  
    int rc=0; /* returns 0 (success); EOF; or errno */  
    int records_read = fread(buffer, size, 1, fp);  
    if (records_read != 1) {  
        if (feof(fp))  
            rc = EOF;  
        else rc = errno;  
    } /* end if no data was returned */  
    return rc;  
} /* end readFromFile */  
  
//-----  
// Pipe functions - for AIX  
//-----  
#include <unistd.h> /* for pipe functions on AIX */  
#include <fcntl.h> /* for definition of O_RDONLY and open() */  
//-----  
int openNamedPipe (char *pipe_name) {  
    return open(pipe_name, O_RDONLY);  
}  
//-----  
int closeNamedPipe (int handle) {  
    return close(handle);  
}  
//-----  
int readFromPipe(int handle, char* buffer, int size) {  
    int rc=0;  
    int num_bytes;  
    num_bytes = read(handle, buffer, size);  
    if (num_bytes != size) {  
        if (num_bytes==0)  
            rc=EOF;  
        else rc = num_bytes;  
    } /* end did not get the expected number of bytes back from read() */  
    return rc;  
} /* end readFromPipe */  
  
//-----  
// Read data from Event Monitor trace (FILE or PIPE) returns 0 (success) or EOF  
//-----  
int read_data(EventLog* evtlog,  
             char* buffer,  
             int size) {  
    int rc=0;  
    if (evtlog->type == EVMFile) {  
        rc = readFromFile(buffer, size, evtlog->current_fp);  
        if (rc && rc!=EOF) {
```

```

        fprintf(stderr, "ERROR: Could not read from: %s\n",
                evtlog->current_fn);
        exit(1);
    } /* end cannot read the log header from the file */
} /* end if the Event Monitor Log is read from a file */
else {
    rc = readFromPipe(evtlog->handle, buffer, size);
    if (rc && rc!=EOF) {
        fprintf(stderr, "ERROR: Could not read a data from: %s\n",
                evtlog->target);
        exit(2);
    } /* end cannot read from the pipe */
} /* end else the Event Log is read from a pipe */
return rc;
} /* end of read_data */

//-----
// Skip n bytes from current position in the trace
//-----
void skip_data(EventLog* evtlog, int n) {
    if (evtlog->type == EVMFile)
        fseek(evtlog->current_fp, n, SEEK_CUR);
    else if (evtlog->type == EVMPipe) {
        lseek(evtlog->handle, n, SEEK_CUR);
    } /* end else pipe event monitor */
} /* end skip_data */

```

## Swapping Bytes in Numerical Values

This code is required when transferring data between systems using different conventions for storing numerical values (for example, UNIX to Windows NT).

```

#include <sqlmon.h> // DB2 Database Monitor interface
//-----
// Byte conversion macros
//-----
#define SWAP2(s) (((s) >> 8) & 0xFF) | (((s) << 8) & 0xFF00)

#define SWAP4(l) (((l) >> 24) & 0xFF) | (((l) & 0xFF0000) >> 8) & 0xFF00 \
| (((l) & 0xFF00) << 8) | ((l) << 24)

//-----
void swapBytes_sqlm_event_log_header(sqlm_event_log_header* r) {
    r->size = SWAP4(r->size);
    r->version = SWAP4(r->version);
    r->codepage_id = SWAP2(r->codepage_id);
    r->country_code = SWAP2(r->country_code);
} // end of swapBytes_sqlm_event_log_header

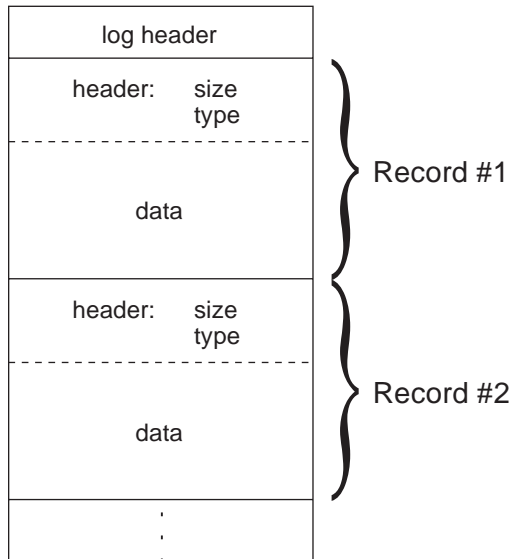
// . . .

```

## Reading the Event Records

After the log header has been read, all remaining records in the trace follow the following layout. Each record starts and terminates on a 4 byte boundary; and records are never split across 2 files.





The **size** accounts for both header and data.

```
//-----
// Read an event record - returns: 0 (success) or EOF
//-----
int read_event_record(EventLog *evtlog, char *buffer, int *event_type) {

    sqlm_event_rec_header* pHeader = (sqlm_event_rec_header*) buffer;

    //-----
    // Read the record header
    //-----
    int rc;
    rc=read_data(evtlog, (char *) pHeader, pHeader->size);
    if (rc)
        return rc; /* could be at EOF */

    *event_type = pHeader->type; // The event type is specified in the header

    if (evtlog->needByteReversal)
        swapBytes_sqlm_event_rec_header(pHeader);

    //-----
    // Read the rest of the data
    //-----
    rc=read_data(evtlog, buffer + pHeader->size,
                pHeader->size - pHeader->size);

    if (rc==0 && evtlog->needByteReversal)
        swapBytes(pHeader->type, buffer);

    return rc;
} /* end of read_event_record */
```

## Reading the Log Header

You must take care of byte reversal, and possibly code page conversion. This example only handles byte reversal.

```
// From sqlmon.h, the DB2 system monitor header file:
// LOG HEADER
typedef struct sqlm_event_log_header
{
    int             byte_order;           /* Big Endian or Little Endian */
    unsigned long  size;                 /* Size of this record */
    unsigned long  version;              /* Event Monitor Version */
    char           event_monitor_name[SQLM_IDENT_SZ]; /* Name of the Event Mon */
    unsigned short codepage_id;          /* Code page of Database */
    unsigned short country_code;         /* Country Code of Database */
    char           server_prdid[SQLM_IDENT_SZ]; /* Server Product Id */
    char           server_instance_name[SQLM_IDENT_SZ]; /*instance name of DB2 */
    unsigned long  number_nodes_in_system;
}sqlm_event_log_header;

//-----
// Attributes of an Event Log
// This structure is used to keep all information that is required to
// process the output of any Event Monitor (file or pipe).
//-----
typedef enum      EventMonitorType { EVMPipe, EVMFile } EventMonitorType;
typedef struct EventLog {
    sqlm_event_log_header  header;
    char*                  target;
    Boolean                 needByteReversal; // True if running on a machine
                                // with a different memory model
    EventMonitorType       type;             // type: file or pipe
    int                    handle;          // For Named pipe only
    FILE*                  current_fp;      // File currently open
    char                   current_fn[512]; // It's name.
} EventLog;

//-----
// Read the log header - store its attributes into an 'EventLog' structure.
//-----
#ifdef _AIX // Compiler pre-defined macro on AIX
#define BIG_ENDIAN_MEMORY
#else
// Assume an INTEL platform
#define LITTLE_ENDIAN_MEMORY
#endif
void read_log_header( /* output */ EventLog* evtlog) {

    // Read the Event Trace header
    read_data(evtlog, // input: event trace (or log)
              (char*) &evtlog->header, // output: log header
              sizeof(sqlm_event_log_header)); // input: number of bytes to read
                                              // is what we can handle

    // Check if the memory model is different and we need byte-reversal
    switch (evtlog->header.byte_order) {
        case SQLM_BIG_ENDIAN:
            #ifdef LITTLE_ENDIAN_MEMORY
                evtlog->needByteReversal = 1;
            #else
                evtlog->needByteReversal = 0;
            #endif
            break;
        case SQLM_LITTLE_ENDIAN:
            #ifndef LITTLE_ENDIAN_MEMORY
                evtlog->needByteReversal = 1;
            #else
                evtlog->needByteReversal = 0;
            #endif
    }
}
```

```

        #endif
        break;
    } // end switch

    // Convert the header, if the server had a different memory model than ours
    if (evtlog->needByteReversal)
        swapBytes_sqlm_event_log_header(&evtlog->header);

    // Skip extra bytes, if the record is bigger than what we can handle
    // (which may become the case in subsequent releases of the product)
    if (evtlog->header.size > sizeof(sqlm_event_log_header)) {
        skip_data(evtlog, evtlog->header.size - sizeof(sqlm_event_log_header));
    } // end if more bytes than expected for this record in the log file
} // end read_log_header

```

## Printing Event Records

All timestamps in event monitor records are **GMT time** since January 1, 1970.

All strings in event monitor records are padded with blanks, up to their maximum size. **Strings are NEVER NULL terminated.**

The following routines illustrate one method of handling blank-padded strings and converting GMT time into local time. It also shows how to print any event monitor.

```

#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <ctype.h> /* To use cprint() function */
#include <time.h>
#include <sqlmon.h> // DB2 Database Monitor interface

//-----
// Convert GMT time into local time, and format it into a printable string.
//-----
char* time_STRING(const sqlm_timestamp timestamp, char *timeString) {

    // Event Monitor returns GMT time, adjust it to local time
    struct tm *pstTm;
    pstTm = localtime((signed long*) &timestamp.seconds);
    if (timestamp.seconds == 0)
        strcpy(timeString, "");
    else {
        if (timestamp.microsec < 0) {
            sprintf(timeString, "%02d-%02d-%04d %02d:%02d:%02d",
                pstTm->tm_mon + 1, pstTm->tm_mday, pstTm->tm_year + 1900,
                pstTm->tm_hour, pstTm->tm_min, pstTm->tm_sec);
        } else {
            sprintf(timeString, "%02d-%02d-%04d %02d:%02d:%02d.%06.61d",
                pstTm->tm_mon + 1, pstTm->tm_mday, pstTm->tm_year + 1900,
                pstTm->tm_hour, pstTm->tm_min, pstTm->tm_sec, timestamp.microsec);
        } /* end else micro seconds are not null */
    } /* end if the timestamp is non-zero */
    return timeString;
} // end of time_STRING

```

All strings in an event monitor traces are BLANK PADDED up to their maximum size. They \*are not\* null terminated.

```

//-----
// Print a Blank Padded String of maximum length SZ
//-----
// note: strings returned by DB2 are NOT NULL-TERMINATED, they are all
//       blank padded up to some maximum length.
//-----

```

```

// For example, given:
//          char str[20] = "Contents of 1      ";
// the following invocation:
//          fpBPSTR(stdout, " String 1", str, sizeof(str));
// will print:
//          " String 1: Contents of 1"
//-----
#define fpBPSTR(fp, txt, str, SZ) \
{ \
    char newstr[SZ]; \
    int k=0; \
    while (str[k]!='\0' && k<SZ) { newstr[k]=str[k]; k++;} \
    if (k<SZ) newstr[k]='\0'; \
    fprintf(fp, txt ": %0.*s\n", SZ, newstr); \
}

//-----
char* byte_order_STRING(int val) {
    switch (val) {
        case SQLM_LITTLE_ENDIAN:    return "SQLM_LITTLE_ENDIAN";
        case SQLM_BIG_ENDIAN:       return "SQLM_BIG_ENDIAN";
    }
    return "";
} // end of byte_order_STRING

//-----
// Print the log header record
//-----
void print_sqlm_event_log_header(FILE* fp,
                                const sqlm_event_log_header *event_header){
    fprintf(fp,
"-----\n"
"          EVENT LOG HEADER\n");
    fpBPSTR(fp, " Event Monitor name",
            event_header->event_monitor_name, SQLM_IDENT_SZ);
    fpBPSTR(fp, " Server Product ID",
            event_header->server_prdid, SQLM_IDENT_SZ);
    fprintf(fp, " Version of event monitor data: %ld\n", event_header->version);
    fprintf(fp, " Byte order: %s\n",
            byte_order_STRING(event_header->byte_order));
    fprintf(fp, " Size of record: %ld\n", event_header->size);
    fprintf(fp, " Codepage of database: %d\n", event_header->codepage_id);
    fprintf(fp, " Country code of database: %d\n", event_header->country_code);
    fpBPSTR(fp, " Server instance name",
            event_header->server_instance_name, SQLM_IDENT_SZ);
    fprintf(fp,
"-----\n");
    fflush(fp);
} /* end of print_sqlm_event_log_header */

//-----
// Print an event record
//-----
void print_event_record(FILE* fp, int event_type, char* rec, int rec_no) {
    switch (event_type) {
        case SQLM_EVENT_DB:
            // print_sqlm_db_event(fp, (const sqlm_db_event*) rec, rec_no);
            break;
        case SQLM_EVENT_CONN:
            // print_sqlm_conn_event(fp, (sqlm_conn_event*) rec, rec_no);
            break;
        case SQLM_EVENT_TABLE:
            // print_sqlm_table_event(fp, (sqlm_table_event*) rec, rec_no);
            break;
        case SQLM_EVENT_STMT:
            // print_sqlm_stmt_event(fp, (sqlm_stmt_event*) rec, rec_no);
            break;
    }
}

```

```

    case SQLM_EVENT_STMTTEXT:
//      print_sqlm_stmttext_event(fp, (sqlm_stmttext_event*) rec, rec_no);
        break;
    case SQLM_EVENT_XACT:
//      print_sqlm_xaction_event(fp, (sqlm_xaction_event*) rec, rec_no);
        break;
    case SQLM_EVENT_DEADLOCK:
//      print_sqlm_deadlock_event(fp, (sqlm_deadlock_event*) rec, rec_no);
        break;
    case SQLM_EVENT_DLCONN:
//      print_sqlm_dlconn_event(fp, (sqlm_dlconn_event*) rec, rec_no);
        break;
    case SQLM_EVENT_TABLESPACE:
//      print_sqlm_tablespace_event(fp, (sqlm_tablespace_event*) rec, rec_no);
        break;
    case SQLM_EVENT_DBHEADER:
//      print_sqlm_dbheader_event(fp, (sqlm_dbheader_event*) rec, rec_no);
        break;
    case SQLM_EVENT_CONNHEADER:
//      print_sqlm_connheader_event(fp, (sqlm_connheader_event*) rec, rec_no);
        break;
    case SQLM_EVENT_OVERFLOW:
//      print_sqlm_overflow_event(fp, (sqlm_overflow_event*) rec, rec_no);
        break;
    case SQLM_EVENT_START:
//      print_sqlm_evmon_start_event(fp, (sqlm_evmon_start_event*) rec, rec_no);
        break;
    default:
//      print_unknown_event(fp, rec, rec_no);
        break;
} /* end switch on event type */
fflush (fp);
} /* end of print_event_record */

```

## Reading Events from a FILE Trace

This sample illustrates how to handle multiple files.

A file event monitor writes to files that are created in the directory identified by its target. When initially turned on, it starts writing to file: 00000000.evt, when this file is full (as specified by the MAXFILESIZE parameter on create event monitor), it moves on to file: 00000001.evt, and so on.

```

//-----
// Build a fully qualified Event Monitor file name, given a file number
//-----
#ifdef _AIX
#define PATH_SEP '/'
#define PATH_SEP_STR "/"
#else /* Assume Intel platform */
#define PATH_SEP '\\\'
#define PATH_SEP_STR "\\\"
#endif
void build_event_monitor_file_name(const char* target, int fnum, char *fn) {

    // Build the full filename (path + filename)
    int len = strlen(target);
    if (target[len-1] == PATH_SEP) {
        sprintf(fn, "%s%0.8d.evt", target, fnum);
    } else {
        sprintf(fn, "%s%s%0.8d.evt", target, PATH_SEP_STR, fnum);
    } // end else need to append path delimiter to directory name
} /* end of build_event_monitor_file_name */

```

```

//-----
// Read Events from files
//-----
void read_events_from_file(char* target) {
    EventLog      evtlog;      /* Attributes for this event log */
    int record_no;           /* current record number */
    int file_no;           /* current file number */
    int rc=0;
    int end_of_trace=0;      /* True when no more files to read */
    char buffer[4096];       /* buffer for reading Event records */

    //-----
    // Initialize the attributes of this Event Log
    //-----
    evtlog.type = EVMFile;    /* A File log
    evtlog.target = target;   /* Directory where the files reside

    //-----
    // Open the first file and read the log header
    //-----
    file_no=0;               /* First file to open is 00000000.evt */
    build_event_monitor_file_name(evtlog.target, file_no, evtlog.current_fn);
    read_log_header(&evtlog);
    print_sqlm_event_log_header(stdout, &evtlog.header); // Print it

    //-----
    // Read/print events from the trace file(s)
    //-----
    record_no=0;             /* Number of event records processed */
    while (!end_of_trace) {
        int event_type;      /* Type of event read from the trace */
        rc=read_event_record(&evtlog, buffer, &event_type);
        if (rc == EOF) {

            /* Try to open the next trace file, if any */
            closeFile(evtlog.current_fp);
            build_event_monitor_file_name(evtlog.target,
                ++file_no, evtlog.current_fn);
            if ((evtlog.current_fp = openFile(evtlog.current_fn))==NULL)
                end_of_trace=true; /* No more files to read */
            else {
                rc=read_event_record(&evtlog, buffer, &event_type);
                if (rc==EOF)
                    end_of_trace = true;
            } /* else read the event from the next file */
        } else if (rc) end_of_trace = true;

        if (rc==0) {
            // Process the event
            print_event_record(stdout, event_type, buffer, ++record_no);
        } /* end if we got an event record */
    } /* end while there are more files in the Event Monitor trace */
} // end of read_events_from_file

```

See evm.c sample application in engn/samples/c\_AIX for the complete source.

## Appendix A. Database System Monitor Interfaces

This appendix contains reference material for the commands, SQL statements, and APIs associated with the database system monitor. The following tables list the commands or APIs to use for a given task.

*Table 1. Snapshot Monitor Commands and APIs for a Given Task*

Snapshot Monitoring Task	API	Command
Taking a snapshot	"sqlmonss - Get Snapshot API" on page 248 "sqlmonsz - Estimate Size Required for sqlmonss() Output Buffer API" on page 260	"GET SNAPSHOT Command" on page 232 "LIST ACTIVE DATABASES Command" on page 235 "LIST APPLICATIONS - Command" on page 237 "LIST DCS APPLICATIONS - Command" on page 239 "GET DATABASE MANAGER MONITOR SWITCHES Command" on page 228
Getting monitor switch status for an application or session	"sqlmon - Get/Update Monitor Switches API" on page 244	"GET MONITOR SWITCHES Command" on page 230
Updating monitor switch status for an application or session	"sqlmon - Get/Update Monitor Switches API" on page 244	"UPDATE MONITOR SWITCHES Command" on page 266
Getting monitor switch status for the database manager level	"sqlmonss - Get Snapshot API" on page 248	"GET DATABASE MANAGER MONITOR SWITCHES Command" on page 228
Resetting counters for an application or session	"sqlmrset - Reset Monitor API" on page 263	"RESET MONITOR Command" on page 241

*Table 2. Event Monitor Commands (or SQL) for a Given Task*

Event Monitor Task	Command
Creating an event monitor	"CREATE EVENT MONITOR Command and SQL" on page 214
Determining if an event monitor is active	"EVENT_MON_STATE SQL Function" on page 227
Activating/Deactivating an event monitor	"SET EVENT MONITOR STATE Command and SQL" on page 242
Deleting an event monitor	"DROP EVENT MONITOR Command and SQL" on page 226
Formatting trace to stdout	"db2evmon - Event Monitor Trace Formatter Command" on page 224
Reading trace with the GUI	"db2eva - Event Analyzer Command" on page 222

## CREATE EVENT MONITOR Command and SQL

---

### CREATE EVENT MONITOR Command and SQL

#### Purpose

Stores an Event Monitor definition in the database catalogs. When activated (see “SET EVENT MONITOR STATE Command and SQL” on page 242), an event monitor will log monitor data when certain events occur while using the database. You should read “Event Monitors” on page 10 and Chapter 4, “Event Monitor Output” on page 195 before using this command.

#### Context

This statement can be embedded in an application program or issued interactively. It is an executable statement that can be dynamically prepared.

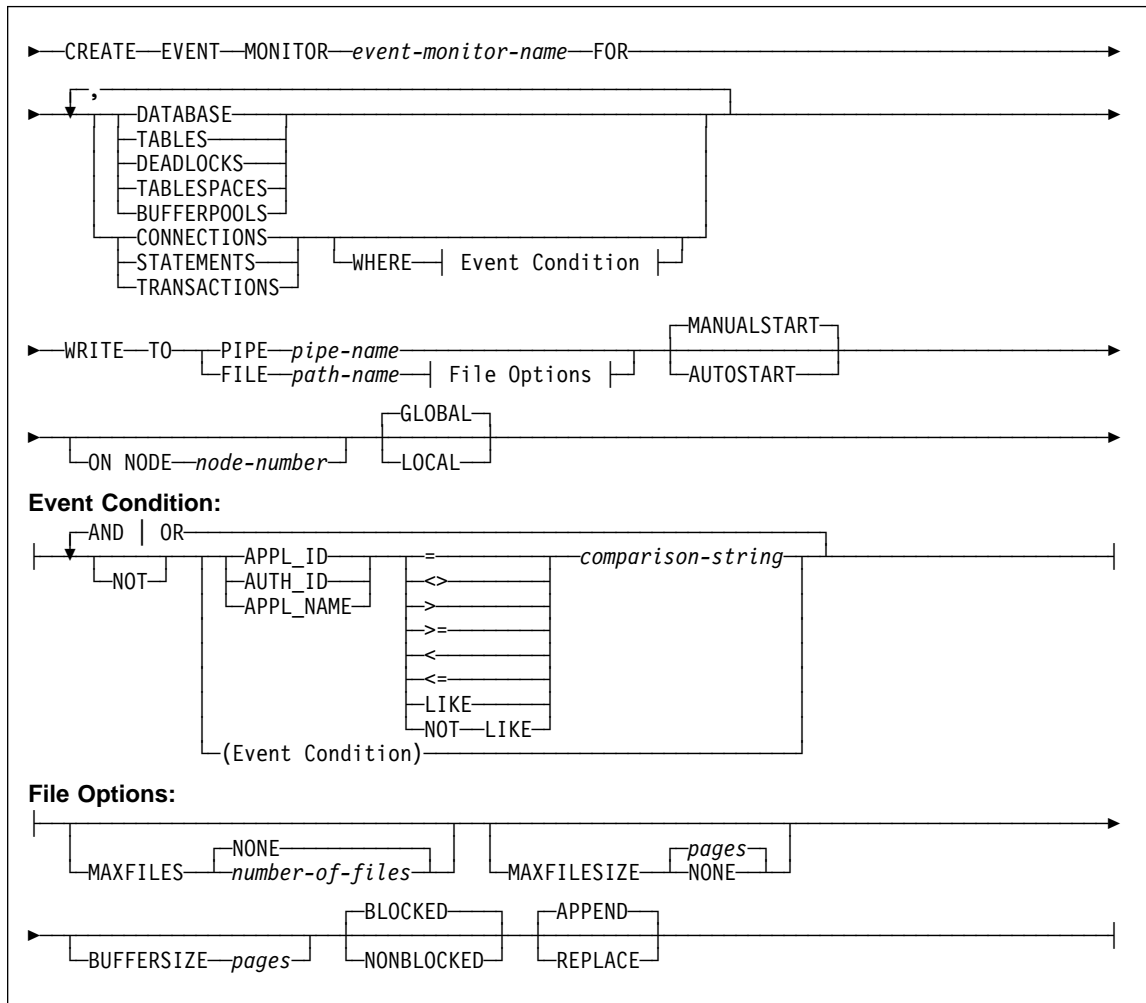
#### Authorization

The privileges held by the authorization ID must include either SYSADM or DBADM authority (SQLSTATE 42502).

#### Format



## CREATE EVENT MONITOR Command and SQL



### Parameters

*event-monitor-name*

Names the event monitor. This is a one-part name. It is an SQL identifier (either ordinary or delimited). The *event-monitor-name* must not identify an event monitor that already exists in the catalog (SQLSTATE 42710).

### FOR

Introduces the type of events to record. See “Event Types” on page 18 for a list of the records produced for each event type and the events that trigger writing them.

### DATABASE

Specifies that the event monitor writes a database record when the database is deactivated.

## CREATE EVENT MONITOR Command and SQL

### TABLES

Specifies that the event monitor writes a table record for each table that has been accessed since database activation, when the database is deactivated. An active table is a table that has changed since the first connection to the database.

### DEADLOCKS

Specifies that the event monitor writes a deadlock record whenever a deadlock occurs.

### TABLESPACES

Specifies that the event monitor writes a table space record for each table space when the database is deactivated.

### BUFFERPOOLS

Specifies that the event monitor writes a buffer pool record when the database is deactivated.

### CONNECTIONS

Specifies that the event monitor writes a connection record when an application disconnects from the database.

### STATEMENTS

Specifies that the event monitor writes a statement record whenever a SQL statement finishes executing.

### TRANSACTIONS

Specifies that the event monitor writes a transaction record whenever a transaction completes (that is, whenever there is a commit or rollback operation).

### WHERE *event condition*

Defines a filter that determines which will be monitored. If the result of the event condition is TRUE for a particular connection, then the event monitor will generate the requested event records for that connection.

This clause is a special form of the WHERE clause that should not be confused with a standard search condition.

If no WHERE clause is specified then all connections will be monitored.

### APPL\_ID

Specifies that the comparison string is an application ID of a connection.

### AUTH\_ID

Specifies that the comparison string is the authorization ID of a connection.

### APPL\_NAME

Specifies that the comparison string is the application program name of the connection.

### *comparison-string*

A string to be compared with the APPL\_ID, AUTH\_ID, or APPL\_NAME of each application that connects to the database. *comparison-string* must be

## CREATE EVENT MONITOR Command and SQL

a string constant (that is, host variables and other string expressions are not permitted).

### WRITE TO

Introduces the target for the data.

### PIPE

Specifies that the target for the event monitor data is a named pipe. The event monitor writes the data to the pipe in a single stream (that is, as if it were a single, infinitely long file). When writing the data to a pipe, an event monitor does not perform blocked writes. If there is no room in the pipe buffer, then the event monitor will discard the data. It is the monitoring application's responsibility to read the data promptly if it wishes to ensure no data loss. See "Using Pipe Event Monitors" on page 18 for more details and examples.

#### *pipe-name*

The name of the pipe (FIFO on AIX) to which the event monitor will write the data.

The naming rules for pipes are platform specific. On UNIX operating systems pipe names are treated like file names. As a result, relative pipe names are permitted, and are treated like relative path-names (see *path-name* below). However, on OS/2, Windows 95 and Windows NT, there is a special syntax for a pipe name. As a result, on OS/2, Windows 95 and Windows NT absolute pipe names are required.

The existence of the pipe will not be checked at event monitor creation time. It is the responsibility of the user to have created and opened the pipe for reading at the time that the event monitor is activated. If the pipe is not available at this time, then the event monitor will turn itself off, and will log an error. (That is, if the event monitor was activated at database start time as a result of the AUTOSTART option, then the event monitor will log an error in the system error log.) If the event monitor is activated via the SET EVENT MONITOR STATE SQL statement, then that statement will fail (SQLSTATE 58030).

### FILE

Indicates that the target for the event monitor data is a file (or set of files). The event monitor writes out the stream of data as a series of 8 character numbered files, with the extension "evt". (for example, 00000000.evt, 00000001.evt, 00000002.evt, etc). The data should be considered to be one logical file even though the data is broken up into smaller pieces (that is, the start of the data stream is the first byte in the file 00000000.evt; the end of the data stream is the last byte in the file nnnnnnnn.evt).

The maximum size of each file can be defined as well as the maximum number of files. An event monitor will never split a single event record across two files. However, an event monitor may write related records in two different files. It is the responsibility of the application that uses this data to keep track of such related information when processing the event files. See "File Event Monitor Target" on page 202 for more information.

## CREATE EVENT MONITOR Command and SQL

### *path-name*

The name of the directory in which the event monitor should write the event files data. The path must be known at the server, however, the path itself could reside on another partition (or node) (for example, in a UNIX-based system, this might be an NFS mounted file). A string constant must be used when specifying the *path-name*.

The directory does not have to exist at CREATE EVENT MONITOR time. However, a check is made for the existence of the target path when the event monitor is activated. At that time, if the target path does not exist, an error (SQLSTATE 428A3) is raised.

If an absolute path (a path that starts with the root directory on AIX, or a disk identifier on OS/2, Windows 95 and Windows NT) is specified, then the specified path will be the one used. If a relative path (a path that does not start with the root) is specified, then the path relative to the DB2EVENT directory in the database directory will be used.

When a relative path is specified, the DB2EVENT directory is used to convert it into an absolute path. Thereafter, no distinction is made between absolute and relative paths. The absolute path is stored in the SYSCAT.EVENTMONITORS catalog view.

It is possible to specify two or more event monitors that have the same target path. However, once one of the event monitors has been activated for the first time, and as long as the target directory is not empty, it will be impossible to activate any of the other event monitors.

### **File Options**

Specifies the options for the file format.

#### **MAXFILES NONE**

Specifies that there is no limit to the number of event files that the event monitor will create. This is the default.

#### **MAXFILES *number-of-files***

Specifies that there is a limit on the number of event monitor files that will exist for a particular event monitor at any time. Whenever an event monitor has to create another file, it will check to make sure that the number of .evt files in the directory is less than *number-of-files*. If this limit has already been reached, then the event monitor will turn itself off.

If an application removes the event files from the directory after they have been written, then the total number of files that an event monitor can produce can exceed *number-of-files*. This option has been provided to allow a user to guarantee that the event data will not consume more than a specified amount of disk space.

## CREATE EVENT MONITOR Command and SQL

### **MAXFILESIZE** *pages*

Specifies that there is a limit to the size of each event monitor file. Whenever an event monitor writes a new event record to a file, it checks that the file will not grow to be greater than *pages* (in units of 4K pages). If the resulting file would be too large, then the event monitor switches to the next file. The default for this option is:

- OS/2, Windows 95 and Windows NT - 200 4K pages
- UNIX - 1000 4K pages

The number of pages must be greater than at least the size of the event buffer in pages. If this requirement is not met, then an error (SQLSTATE 428A4) is raised.

### **MAXFILESIZE NONE**

Specifies that there is no set limit on a file's size. If MAXFILESIZE NONE is specified, then MAXFILES 1 must also be specified. This option means that one file will contain all of the event data for a particular event monitor. In this case the only event file will be 00000000.evt. This is the default.

### **BUFFERSIZE** *pages*

Specifies the size of the event monitor buffers (in units of 4K pages). All FILE event monitor I/O is buffered to improve the performance of the event monitors. The larger the buffers, the less I/O will be performed by the event monitor. Highly active event monitors should have larger buffers than relatively inactive event monitors. When the monitor is started, two buffers of the specified size are allocated. Event monitors use double buffering to permit asynchronous I/O. See "File Event Monitor Buffering" on page 201 for more information.

The minimum and default size of each buffer (if this option is not specified) is 1 page (that is, 2 buffers, each 4 K in size). The maximum size of the buffers is limited by the size of the database manager monitor heap (MON\_HEAP\_SZ) since the buffers are allocated from this heap.

### **BLOCKED**

Specifies that each agent that generates an event should wait for the event monitor to finish writing a buffer out to disk when both are full. BLOCKED should be selected to guarantee no event data loss. This is the default option. See "Blocked Event Monitors" on page 202 and "Non-Blocked Event Monitors" on page 202 for more information.

### **NONBLOCKED**

Specifies that each agent that generates an event should not wait when the event monitor buffers are full. NONBLOCKED event monitors do not slow down database operations to the extent of

## CREATE EVENT MONITOR Command and SQL

BLOCKED event monitors. However, NONBLOCKED event monitors are subject to data loss on highly active systems. See “Blocked Event Monitors” on page 202 and “Non-Blocked Event Monitors” on page 202 for more information.

### APPEND

Specifies that if event data files already exist when the event monitor is turned on, then the event monitor will append the new event data to the existing stream of data files. When the event monitor is re-activated, it will resume writing to the event files as if it had never been turned off. APPEND is the default option. See “Restarting a File Event Monitor” on page 204 for more information.

The APPEND option does not apply at CREATE EVENT MONITOR time, if there is existing event data in the directory where the newly created event monitor is to write its event data.

### REPLACE

Specifies that if event data files already exist when the event monitor is turned on, then the event monitor will erase all of the event files and start writing data to file 00000000.evt.

### MANUALSTART

Specifies that the event monitor not be started automatically each time the database is started. Event monitors with the MANUALSTART option must be activated manually using the SET EVENT MONITOR STATE statement. This is the default option.

### AUTOSTART

Specifies that the event monitor be started automatically each time the database is started.

### ON NODE

*node-number*

Specifies a node number where the event monitor output thread or process runs. If defined as global, the event monitor output data, from all nodes, will be directed to that node. The default is the node where the command was issued.

### LOCAL

Event monitor reports activity only for the node where it is running (the monitor node). This is the default.

### GLOBAL

Event monitor reports activity from all nodes. In DB2 Version 5, for a partitioned database only deadlock event monitors can be defined as GLOBAL. The GLOBAL event monitor will report deadlocks involving applications running on any node in the system.

## CREATE EVENT MONITOR Command and SQL

### Usage

- Creating an event monitor **does not** activate it or create the target directory.

### Comments

- Event monitor definitions are recorded in the SYSCAT.EVENTMONITORS catalog view. The events themselves are recorded in the SYSCAT.EVENTS catalog view.

### Related Information

“SET EVENT MONITOR STATE Command and SQL” on page 242

“DROP EVENT MONITOR Command and SQL” on page 226

## db2eva - Event Analyzer Command

---

### db2eva - Event Analyzer Command

#### Purpose

Starts the event analyzer GUI from the command line, allowing the user to view traces produced an event monitor.

#### Authorization


None, unless connecting to the database (-evm and -db,), then one of the following is required:

*sysadm*  
*sysctrl*  
*sysmaint*  
*dbadm*

#### Required Connection

None

#### Format



#### Parameters

Two methods of operation are provided for reading traces with db2eva, you can:

1. Specify the directory where the trace files reside (using -path) This mode allow users to move trace files from a server and analyze them locally, even if the Event Monitor has been dropped.
2. Specify the database and event monitor names; db2eva then automatically locates the trace files. When this mode is used, db2eva connects to the database, and issues a *select target from syscat.eventmonitors* to locate the directory where the Event Monitor writes the event records. The connection is then released, unless -conn is specified. This method cannot be used if the event monitor has been dropped.

**-path** *evmon-target*

Specifies the absolute path of the event monitor target, which can either be a directory or a named pipe.



## db2eva - Event Analyzer Command

### **-evm** *evmon-name* **-db** *database-alias*

When **-evm** and **-db** are supplied, db2eva connects to the database, and obtains the directory or pipe to which the event monitor is writing, by issuing an SQL select from the syscat.eventmonitor catalog. For FILE event monitors, this means that you cannot move the trace files to a different directory. Specify the database for which the event monitor is defined, as catalogued on the machine where the trace is analyzed. Using *database-alias* overrides the database name specified in the trace.

### **-conn**

Requests db2eva to maintain a connection to the database specified by the **-db** option. Or if **-db** is not supplied, then to the database specified in the trace file header. Keeping a connection allows the event analyzer to obtain information that is not contained in the trace files, for example the text for static SQL statements. (The statement text events for static SQL contain package creator, package number and section number. When the **-conn** option is specified, db2eva connects to the database and retrieves the text from the database system catalog using these fields). The default is not to keep a connection.

## Comments

This tool is only available on OS/2 and Windows platforms.

It does not display database, table space, buffer pool, and table event records, but properly reads traces containing them.

## db2evmon - Event Monitor Trace Formatter Command

---

### db2evmon - Event Monitor Trace Formatter Command

#### Purpose

Formats the trace produced by file or pipe event monitors, and writes it to standard output. This tool is located in the misc subdirectory of the sql1lib directory of the instance.

#### Authorization

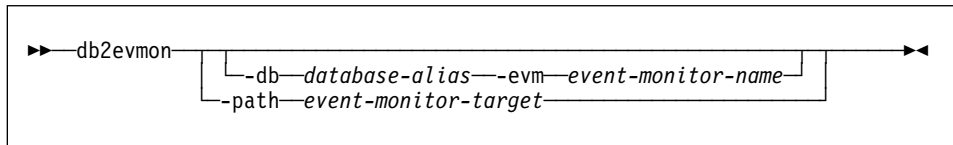
None, unless connecting to the database (-evm and -db are specified), then one of the following is required:

*sysadm*  
*sysctrl*  
*sysmaint*  
*dbadm*

#### Required Connection

None

#### Format



#### Parameters

Two methods of operation are provided for formatting a trace:

1. Specify where the trace resides, using the -path option.
2. Let **db2evmon** locate the trace by selecting the event monitor target from the SYSCAT.EVENTMONITORS database catalog. You need to use -evm and -db.

**-path** *event-monitor-target*

Specifies the name of the directory containing the event monitor trace files, or the name of the pipe where the event monitor is writing its records.

**-db** *database-alias* **-evm** *event-monitor-name*

Specifies the database where the event monitor is defined, event-monitor-name one-part name of the event monitor. An ordinary or delimited SQL identifier.

#### Comments

If the data is being written to a pipe, the tool formats the output for display using standard output as event records are written occur. In this case, the tool must be started **before** the monitor is turned on.

## db2evmon - Event Monitor Trace Formatter Command

### See Also

"Programming to Read an Event Monitor Trace" on page 204

## DROP EVENT MONITOR Command and SQL

---

### DROP EVENT MONITOR Command and SQL

#### Purpose

Removes an event monitor definition from the Database catalogs. Whenever an object is deleted, its description is deleted from the catalog and any packages that reference the object are invalidated.

#### Context

This statement can be embedded in an application program or issued interactively. It is an executable statement that can be dynamically prepared.

#### Authorization

The authorization ID of the DROP statement when dropping an event monitor must have SYSADM or DBADM authority

#### Format

```
►►—DROP—EVENT—MONITOR—event-monitor-name—◄◄
```

#### Parameters

**EVENT MONITOR** *event-monitor-name*

Identifies the event monitor that is to be dropped. The *event-monitor-name* must identify an event monitor that is described in the catalog (SQLSTATE 42704).

If the identified event monitor is ON, an error (SQLSTATE 55034) is raised. Otherwise, the event monitor is deleted.

If there are event files in the target path of the event monitor when the event monitor is dropped, the event files are not deleted.

#### Usage

An event monitor must be stopped or OFF before it can be deleted. Dropping an event monitor does not erase the target directory.

---

## EVENT\_MON\_STATE SQL Function

### Purpose

The EVENT\_MON\_STATE function returns the current state of an event monitor.

### Context

This statement can be embedded in an application program or issued interactively. It is an executable statement that can be dynamically prepared.

### Format

▶▶—EVENT_MON_STATE—( <i>string-expression</i> )—————▶▶
--

### Parameters

*string-expression*

The argument is a string expression with a resulting type of CHAR or VARCHAR and a value that is the name of an event monitor. If the named event monitor does not exist in the SYSCAT.EVENTMONITORS catalog table, SQLSTATE 42704 will be returned.

### Usage

The schema is SYSIBM.

The result is an integer with one of the following values:

- 0 The event monitor is inactive.
- 1 The event monitor is active.

If the argument can be null, the result can be null; if the argument is null, the result is the null value.

Example:

- The following example selects all of the defined event monitors, and indicates whether each is active or inactive:

```

SELECT EVMONNAME,
       CASE
         WHEN EVENT_MON_STATE(EVMONNAME) = 0 THEN 'Inactive'
         WHEN EVENT_MON_STATE(EVMONNAME) = 1 THEN 'Active'
       END
FROM SYSCAT.EVENTMONITORS

```

## GET DATABASE MANAGER MONITOR SWITCHES Command

---

### GET DATABASE MANAGER MONITOR SWITCHES Command

#### Purpose

Displays the status of the database manager monitor switches. Monitor switches instruct the database system manager to collect statistics about its operation and performance and that of the applications using it. A database manager-level switch is on when any monitoring application has turned it on, an active event monitor requires this switch to be on, or it has been set in the database manager configuration. This command is used to determine if the database manager is currently collecting data.

#### Authorization


One of the following:

*sysadm*  
*sysctrl*  
*sysmaint*

#### Required Connection

Instance. To display the settings for a remote instance, or for a different local instance, it is necessary to first attach to that instance.

#### Format



#### Parameters

None

#### Example

The following is sample output from GET DATABASE MANAGER MONITOR SWITCHES:

```
DBM System Monitor Information Collected
Buffer Pool Activity Information (BUFFERPOOL) = ON 06-11-1997 10:11:01.738377
Lock Information (LOCK) = OFF
Sorting Information (SORT) = ON 06-11-1997 10:11:01.738400
SQL Statement Information (STATEMENT) = OFF
Table Activity Information (TABLE) = OFF
Unit of Work Information (UOW) = ON 06-11-1997 10:11:01.738353
```

## GET DATABASE MANAGER MONITOR SWITCHES Command

### Comments

This command returns the settings for the database manager, indicating whether the database manager is currently collecting monitor data. To see the switch settings for your session issue the GET MONITOR SWITCHES command.

Default switch settings can be set in the database manager configuration file. If switches are set ON in the database manager configuration file, then the database manager will always be collecting monitor data, even if all monitoring applications have turned off their switches.

### See Also

“GET MONITOR SWITCHES Command” on page 230

“GET SNAPSHOT Command” on page 232

“RESET MONITOR Command” on page 241

“UPDATE MONITOR SWITCHES Command” on page 266

“sqlmonss - Get Snapshot API” on page 248.

## GET MONITOR SWITCHES Command

---

### GET MONITOR SWITCHES Command

#### Purpose

Displays the status of the database system monitor switches for the current session. Monitor switches instruct the database system manager to collect database activity information. Each application using the database system monitor interface has its own set of monitor switches. This command displays them. To display the database manager-level switches, use “GET DATABASE MANAGER MONITOR SWITCHES Command” on page 228. If a particular switch is on, this command also displays the time stamp for when the switch was turned on.

#### Authorization

One of the following:

*sysadm*  
*sysctrl*  
*sysmaint*

#### Required Connection

Instance.

#### Format

▶▶ GET MONITOR SWITCHES ◀◀

#### Parameters

None

#### Example

The following is sample output from GET MONITOR SWITCHES:

```
Monitor Recording Switches
Buffer Pool Activity Information (BUFFERPOOL) = ON 02-20-1997 16:04:30.070073
Lock Information (LOCK) = OFF
Sorting Information (SORT) = OFF
SQL Statement Information (STATEMENT) = ON 02-20-1997 16:04:30.070073
Table Activity Information (TABLE) = OFF
Unit of Work Information (UOW) = ON 02-20-1997 16:04:30.070073
```



## GET MONITOR SWITCHES Command

### Comments

When a database system monitor command is first issued the session inherits the switch settings of the database manager configuration. The settings can be overridden using "UPDATE MONITOR SWITCHES Command" on page 266.

### See Also

"GET DATABASE MANAGER MONITOR SWITCHES Command" on page 228

"GET SNAPSHOT Command" on page 232

"RESET MONITOR Command" on page 241

"UPDATE MONITOR SWITCHES Command" on page 266

"sqlmon - Get/Update Monitor Switches API" on page 244.

# GET SNAPSHOT Command

---

## GET SNAPSHOT Command

### Purpose

Collects some of the data that the database manager maintains about its operation and performance. The information returned represents a snapshot of this data at the time the command is issued.

### Scope

Returns data only for the node to which the session is attached.

### Authorization

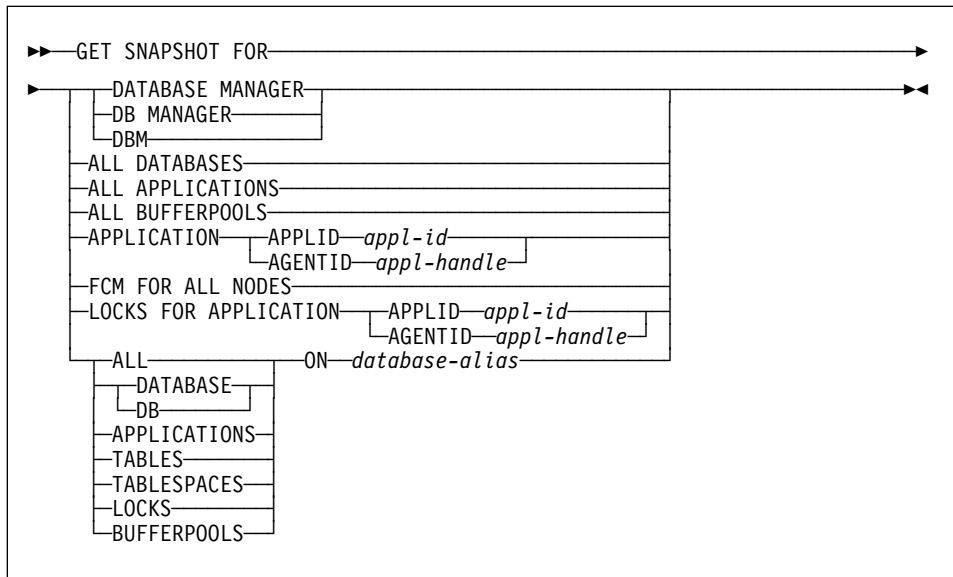
One of the following:

- sysadm*
- sysctrl*
- sysmaint*

### Required Connection

Instance. To obtain a snapshot of a remote instance, it is necessary to first attach to that instance.

### Format



### Parameters

Parameters are group by the way you would use them.

**Note:** The appl-id and appl-handle can be obtained by issuing a LIST APPLICATIONS command (see “LIST APPLICATIONS - Command” on page 237).

#### **DATABASE MANAGER**

Database manager level information, including internal monitor switch settings. On multi-node systems FCM information is also returned.

#### **FCM FOR ALL NODES**

Provides Fast Communication Manager (FCM) statistics with for this node with respect to all other nodes.

#### **DATABASE ON *database-alias***

Database level information and counters for a database. Information is returned only if there is at least one application connected to the database.

#### **ALL DATABASES**

Same information for each database active on the node.

#### **ALL ON *database-alias***

For the specified database, returns: database snapshot, lock snapshot, buffer pool snapshot, and application snapshot(for each connection to the database).

#### **APPLICATION APPLID *appl-id***

Application level information, includes cumulative counters, status information. If the statement switch is ON, it also returns statistics about each cursor currently open and the most recent SQL statement executed.

#### **APPLICATION AGENTID *appl-handle***

Same as APPLICATION APPLID.

#### **APPLICATIONS ON *database-alias***

Same information as APPLICATION APPLID, for each application that is connected to the database on this node.

#### **ALL APPLICATIONS**

Same information as APPLICATION APPLID, for each application that is connected to a database on the current node.

#### **TABLES ON *database-alias***

Returns Table activity information at the database and application level for each application connected to the database, and at the table level for each table that \*was accessed\* by an application connected to the database. Requires table switch to be ON.

#### **LOCKS FOR APPLICATION APPLID *appl-id***

Provides information about each lock held by the application. Provides lock wait information, if application is waiting for a lock. Requires lock switch to be ON.

#### **LOCKS FOR APPLICATION AGENTID *appl-handle***

Same information as LOCKS FOR APPLICATION APPLID.

## GET SNAPSHOT Command

### **LOCKS ON** *database-alias*

Same information as LOCKS FOR APPLICATION APPLID for each application connected to the mentioned database. Plus a database level summary. Requires lock switch

### **TABLESPACES ON** *database-alias*

Information about tablespace activity at the database level; the application level for each application connected to the database; and the tablespace level for each tablespace that has been accessed by an application connected to the database. Requires buffer pool switch.

### **ALL BUFFERPOOLS**

Provides buffer pool activity counters. Requires buffer pool switch to be ON.

### **BUFFERPOOLS ON** *database-alias*

Same information as ALL BUFFERPOOLS, but only for the specified database.

## Comments

**INSTANCE CONNECTION:** If not attached to an instance, issuing this command will automatically attach your session to the instance specified by the DB2INSTANCE environment variable.

To obtain a snapshot from a remote instance (or a different local instance), it is necessary to first attach to that instance. If an alias for a database residing at a different instance is specified, an error message is returned.

**DATA COLLECTED UNDER SWITCH CONTROL:** Data elements that are collected by the DBMS only if a monitor switch is ON are either not returned, or returned as 'Not Collected'. Check Chapter 3, "Database System Monitor Data Elements" on page 31 to determine if a switch needs to be turned on for a data element.

## See Also

"GET MONITOR SWITCHES Command" on page 230  
"RESET MONITOR Command" on page 241  
"UPDATE MONITOR SWITCHES Command" on page 266  
"LIST ACTIVE DATABASES Command" on page 235  
"LIST APPLICATIONS - Command" on page 237  
"LIST DCS APPLICATIONS - Command" on page 239  
"sqlmonss - Get Snapshot API" on page 248

**LIST ACTIVE DATABASES Command**

**Purpose**

Displays the list of databases that are active on this instance. This is a subset of the information listed by the GET SNAPSHOT FOR ALL DATABASES command (see “GET SNAPSHOT Command” on page 232). For each active database, this command displays the following:

- Database name
- Number of applications currently connected to the database
- Database path.

**Scope**

Returns data only for the node to which the session is attached.


**Authorization**

None

**Required Connection**

Instance. To obtain a snapshot of a remote instance, it is necessary to first attach to that instance.

**Format**



**Parameters**

None

**Examples**

Following is sample output from the LIST ACTIVE DATABASES command:

```

Active Databases

Database name           = TEST
Applications connected currently = 0
Database path          = /home/smith/smith/NODE00000/SQL00002/

Database name           = SAMPLE
Applications connected currently = 1
Database path          = /home/smith/smith/NODE00000/SQL00001/

```

## LIST ACTIVE DATABASES

### See Also

- “GET SNAPSHOT Command” on page 232.
- “LIST APPLICATIONS - Command” on page 237
- “sqlmonss - Get Snapshot API” on page 248

---

## LIST APPLICATIONS - Command

### Purpose

Displays the list of applications connected to a database on the instance, including secondary connections established to access a partitioned database.

### Scope

This command only returns information for the node on which it is issued.

### Authorization


One of the following:

*sysadm*  
*sysctrl*  
*sysmaint*

### Required Connection

Instance. To list applications for a remote instance, it is necessary to first attach to that instance.

### Format



### Parameters

#### FOR DATABASE *database-alias*

Information for each application that is connected to the specified database is to be displayed. Database name information is not displayed. If this option is not specified, the command displays the information for each application that is currently connected to any database at the node to which the user is currently attached.

The default application information is comprised of the following:

- Authorization ID
- Application program name
- Application handle
- Application ID
- Database name.

## LIST APPLICATIONS - Command

### SHOW DETAIL

Output will include the following additional information:

- Sequence #
- Application status
- Status change time
- Database path.

**Note:** If this option is specified, it is recommended that the output be redirected to a file, and that the report be viewed with the help of an editor. The output lines may wrap around when displayed on the screen.

### Example

The following is sample output from LIST APPLICATIONS:

Auth Id	Application Name	Appl. Handle	Application Id	DB Name	# of Agents
smith	db2bp_32	12	*LOCAL.smith.970220191502	TEST	1
smith	db2bp_32	11	*LOCAL.smith.970220191453	SAMPLE	1

**Note:** For more information about these fields, see Chapter 3, “Database System Monitor Data Elements” on page 31.

### Comments

The database administrator can use the output from this command as an aid to problem determination. In addition, this information is required if the database administrator wants to use “GET SNAPSHOT Command” on page 232 or FORCE an application.

To list applications at a remote instance (or a different local instance), it is necessary to first attach to that instance. If FOR DATABASE is specified when an attachment exists, and the database resides at an instance which differs from the current attachment, the command will fail.

### See Also

“sqlmonss - Get Snapshot API” on page 248

“GET SNAPSHOT Command” on page 232



---

## LIST DCS APPLICATIONS - Command

### Purpose

Displays the contents of the Database Connection Services (DCS) directory to standard output.

### Authorization


One of the following:

*sysadm*  
*sysctrl*  
*sysmaint*

### Required Connection

Instance. To list the DCS applications at a remote instance, it is necessary to first attach to that instance.

### Format



### Parameters

#### LIST DCS APPLICATIONS

The default application information includes:

- Host authorization ID (*username*)
- Application program name
- Agent ID
- Outbound application ID (*luwid*).

#### SHOW DETAIL

Specifies that output include the following additional information:

- Application ID
- Application sequence number
- Client database alias
- Client node name (*nname*)
- Client product ID
- Code page ID
- Code page
- Outbound sequence number
- Host database name
- Host product ID.

## LIST DCS APPLICATIONS - Command

### Example

The following is sample output from LIST DCS APPLICATIONS:

Auth Id	Application Name	Agent Id	Outbound Application Id
DDCSUS1	db2bp	89330	CAIBMOML.OMXT4H08.A79EAA3C6E29

**Note:** For more information about these fields, see Chapter 3, “Database System Monitor Data Elements” on page 31.

### Comments

The database administrator can use this command to match client application connections *to* the gateway with corresponding host connections *from* the gateway.

The database administrator can also use agent ID information to force specified applications off a DDCS server.

---

## RESET MONITOR Command

### Purpose

Resets the counters maintained by the database system monitor for the current session for specified database, or for all active databases, to zero.

### Authorization


One of the following:

*sysadm*  
*sysctrl*  
*sysmaint*

### Required Connection

Instance.

### Format



### Parameters

#### ALL

This option indicates that the internal counters should be reset for all databases. Some database manager information is also reset.

#### FOR DATABASE *database-alias*

This option indicates that only the counters for the database with alias *database-alias* should be reset.

### Comments

Each session (instance) has its own private view of the monitor data. If one user resets, other users are not affected.

This resets the data for all monitor switches. To reset data for a single switch turn it OFF, then ON, using "UPDATE MONITOR SWITCHES Command" on page 266.

### See Also

- "GET SNAPSHOT Command" on page 232
- "GET MONITOR SWITCHES Command" on page 230
- "UPDATE MONITOR SWITCHES Command" on page 266
- "sqlmrset - Reset Monitor API" on page 263

## SET EVENT MONITOR STATE Command and SQL

---

### SET EVENT MONITOR STATE Command and SQL

#### Purpose

The SET EVENT MONITOR STATE statement activates or deactivates an event monitor. The SET EVENT MONITOR STATE statement is not under transaction control.


#### Context

This statement can be embedded in an application program or issued through the use of dynamic SQL statements. It is an executable statement that can be dynamically prepared.

#### Authorization

The authorization ID of the statement must hold either SYSADM or DBADM authority (SQLSTATE 42815).

#### Format



#### Parameters

##### *event-monitor-name*

Identifies the event monitor to activate or deactivate. The name must identify an event monitor that exists in the catalog (SQLSTATE 42704).

##### *new-state*

*new-state* can be specified either as an integer constant or as the name of a host variable that will contain the appropriate value at run time. The following may be specified:

- |          |   |
|----------|---|
| <b>0</b> | Indicates that the specified event monitor should be deactivated.   |
| <b>1</b> | Indicates that the specified event monitor should be activated. The event monitor should not already be active; otherwise a warning (SQLSTATE 01598) is issued. |

##### *host-variable*

The data type is INTEGER. The value specified must be 0 or 1 (SQLSTATE 42815). If *host-variable* has an associated indicator variable, the value of that indicator variable must not indicate a null value (SQLSTATE 42815).

## SET EVENT MONITOR STATE Command and SQL

### Usage

- Although an unlimited number of event monitors may be defined, there is a limit of 32 event monitors that can be simultaneously active (SQLSTATE 54030).
- In order to activate an event monitor, the transaction in which the event monitor was created must have been committed (SQLSTATE 55033). This rule prevents (in one unit of work) creating an event monitor, activating the monitor, then rolling back the transaction.
- If the number or size of the event monitor files exceeds the values specified for MAXFILES or MAXFILESIZE on the CREATE EVENT MONITOR statement, an error (SQLSTATE 54031) is raised.
- If the target path of the event monitor (that was specified on the CREATE EVENT MONITOR statement) is already in use by another event monitor, an error (SQLSTATE 51026) is raised.

### Comments

- Activating an event monitor performs a reset of any counters associated with it.

The following example activates an event monitor called SMITHPAY.

```
SET EVENT MONITOR SMITHPAY STATE = 1
```

The current state of an event monitor (active or inactive) is determined by using the EVENT\_MON\_STATE built-in function.

## sqlmon - Get/Update Monitor Switches API

---

### sqlmon - Get/Update Monitor Switches API

#### Purpose

Selectively turns on or off switches for groups of monitor data to be collected by the database manager. Returns the current state of these switches for the application issuing the call.

To get the current state of the switches at the database manager level use “sqlmonss - Get Snapshot API” on page 248.

#### Scope

This API only affects the application making the call.

#### Authorization

One of the following:

*sysadm*  
*sysctrl*  
*sysmaint*

#### Required Connection

Instance.

#### API Include File

*sqlmon.h*

#### C API Syntax

```
/* File: sqlmon.h */
/* API: Get/Update Monitor Switches */
/* ... */
int SQL_API_FN
sqlmon (
    unsigned long        version,
    _SQLOLDCHAR         *reserved,
    sqlm_recording_group group_states[],
    struct sqlca         *sqlca);
/* ... */
```

#### API Parameters

*sqlca*

Output. A pointer to the *sqlca* structure, that returns error information.

## sqlmon - Get/Update Monitor Switches API

### *group\_states*

Input/Output. Specifies the monitor switches to update and returns their values. It is an array of `sqlm_recording_group`, one for each monitor switch:

<code>input_state</code>	To request the setting for a switch use <code>SQLM_HOLD</code> . To turn a switch OFF use <code>SQLM_OFF</code> . To turn a switch ON use <code>SQLM_ON</code> .
<code>output_state</code>	The current setting for the switch, either <code>SQLM_OFF</code> or <code>SQLM_ON</code> .
<code>start_time</code>	A timestamp indicating the time a switch was turned ON. A value of 0 (zero) is returned if the switch is OFF.

### *reserved*

Reserved for future use. Users should set this value to NULL.

### *version*

Input. Version ID of the database monitor data to collect. The database monitor only returns data that was available for the requested version. Set this parameter to one of the following symbolic constants:

- `SQLM_DBMON_VERSION1`
- `SQLM_DBMON_VERSION2`
- `SQLM_DBMON_VERSION5`

If requesting data for a version higher than the current server, the database monitor only returns data for its level.

**Note:** If `SQLM_DBMON_VERSION1` is specified as the version, the APIs cannot be run remotely.

## Comments

To obtain the status of the switches at the database manager level, call “`sqlmonss - Get Snapshot API`” on page 248, specifying `SQMA_DB2` for `OBJ_TYPE` (get snapshot for database manager).

## See Also

“`sqlmonss - Get Snapshot API`” on page 248

“`sqlmonsz - Estimate Size Required for sqlmonss() Output Buffer API`” on page 260

“`sqlmrset - Reset Monitor API`” on page 263

“`UPDATE MONITOR SWITCHES Command`” on page 266

“`GET MONITOR SWITCHES Command`” on page 230

## Code Sample

The following example illustrates how to update the monitor switches and print their current settings.

## sqlmon - Get/Update Monitor Switches API

```
/*
   Database Monitor Switch API
*/
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <time.h>
#include "sqlca.h"
#include "sqlutil.h"          // for using sqlaintp
#include "sqlmon.h"
//-----
// Database Monitor Switch API Sample
//-----
char* sw_status_string(int val) {
    switch (val) {
        case SQLM_OFF: return "OFF";
        case SQLM_ON:  return "ON";
    }
    return "";
}

void print_sws(sqlm_recording_group group_states[SQLM_NUM_GROUPS]);
void print_sw_set_times(sqlm_recording_group group_states[SQLM_NUM_GROUPS]);
int main() {
    //-----
    // Set Table switch ON, UOW switch OFF, and query the current (default)
    // values for the other switches. (see Database Manager Configuration)
    //-----
    sqlm_recording_group group_states[SQLM_NUM_GROUPS];
    struct sqlca sqlca;
    group_states[SQLM_TABLE_SW].input_state = SQLM_ON;
    group_states[SQLM_UOW_SW].input_state = SQLM_OFF;
    group_states[SQLM_STATEMENT_SW].input_state = SQLM_HOLD;
    group_states[SQLM_BUFFER_POOL_SW].input_state = SQLM_HOLD;
    group_states[SQLM_LOCK_SW].input_state = SQLM_HOLD;
    group_states[SQLM_SORT_SW].input_state = SQLM_HOLD;
    //-----
    // Perform the call
    //-----
    sqlmon(SQLM_DBMON_VERSION5, NULL, group_states, &sqlca);
    if (sqlca.sqlcode<0) {
        // get and display a printable error message
        char msg[1024];
        sqlaintp (msg, sizeof(msg), 60, &sqlca);
        printf("%s", msg);
    }
    //-----
    // Print the output
    //-----
    print_sws(group_states);          // Print the switch values
    print_sw_set_times(group_states); // Print their switch set time (if ON)
} // end of Database Monitor Switch API sample
//-----
// print switch values
//-----
void print_sws(sqlm_recording_group group_states[SQLM_NUM_GROUPS]) {
    printf("SQLM_UOW_SW:      %s\n",
           sw_status_string(group_states[SQLM_UOW_SW].output_state));
    printf("SQLM_STATEMENT_SW:   %s\n",
           sw_status_string(group_states[SQLM_STATEMENT_SW].output_state));
    printf("SQLM_TABLE_SW:         %s\n",
           sw_status_string(group_states[SQLM_TABLE_SW].output_state));
    printf("SQLM_BUFFER_POOL_SW:   %s\n",
           sw_status_string(group_states[SQLM_BUFFER_POOL_SW].output_state));
    printf("SQLM_LOCK_SW:          %s\n",
           sw_status_string(group_states[SQLM_LOCK_SW].output_state));
    printf("SQLM_SORT_SW:         %s\n",
           sw_status_string(group_states[SQLM_SORT_SW].output_state));
}
```



## sqlmon - Get/Update Monitor Switches API

```
        sw_status_string(group_states[SQLM_SORT_SW].output_state));
} // end print_sws
//-----
// print switch set times (if ON)
//-----
void print_sw_set_times(
    sqlm_recording_group group_states[SQLM_NUM_GROUPS]) {
    if (group_states[SQLM_UOW_SW].start_time.seconds) {
        printf("SQLM_UOW_SW start_time:      %s\n", ctime((time_t *)
            &group_states[SQLM_UOW_SW].start_time.seconds));
    }
    if (group_states[SQLM_STATEMENT_SW].start_time.seconds) {
        printf("SQLM_STATEMENT_SW start_time:  %s\n", ctime((time_t *)
            &group_states[SQLM_STATEMENT_SW].start_time.seconds));
    }
    if (group_states[SQLM_TABLE_SW].start_time.seconds) {
        printf("SQLM_TABLE_SW start_time:      %s\n", ctime((time_t *)
            &group_states[SQLM_TABLE_SW].start_time.seconds));
    }
    if (group_states[SQLM_BUFFER_POOL_SW].start_time.seconds) {
        printf("SQLM_BUFFER_POOL_SW start time: %s\n", ctime((time_t *)
            &group_states[SQLM_BUFFER_POOL_SW].start_time.seconds));
    }
    if (group_states[SQLM_LOCK_SW].start_time.seconds) {
        printf("SQLM_LOCK_SW start time:      %s\n", ctime((time_t *)
            &group_states[SQLM_LOCK_SW].start_time.seconds));
    }
    if (group_states[SQLM_SORT_SW].start_time.seconds) {
        printf("SQLM_SORT_SW start_time:      %s\n", ctime((time_t *)
            &group_states[SQLM_SORT_SW].start_time.seconds));
    }
} // end print_sw_set_times (if ON)
```

## sqlmonss - Get Snapshot API

---

### sqlmonss - Get Snapshot API

#### Purpose

Collects database manager monitor information and returns it to a user-allocated data buffer. The information returned represents a *snapshot* of the database manager operational status at the time the API was called.

#### Scope

This API returns information only for the instance.

#### Authorization

One of the following:

*sysadm*  
*sysctrl*  
*sysmaint*

#### Required Connection

Instance. To obtain a snapshot from a remote instance (or a different local instance), it is necessary to first attach to that instance.

#### API Include File

*sqlmon.h*

#### C API Syntax

```
/* File: sqlmon.h */
/* API: Get Snapshot */
/* ... */
int SQL_API_FN
sqlmonss (
    unsigned long    version,
    _SQLOLDCHAR     *reserved,
    sqlma           *sqlma_ptr,
    unsigned long    buffer_length,
    void            *buffer_area,
    sqlm_collected *collected,
    struct sqlca    *sqlca);
/* ... */
```

#### API Parameters

## sqlmonss - Get Snapshot API

### *version*

Input. Version ID of the database monitor data to collect. The database monitor only returns data that was available for the requested version. Set this parameter to one of the following symbolic constants:

- `SQLM_DBMON_VERSION1`
- `SQLM_DBMON_VERSION2`
- `SQLM_DBMON_VERSION5`

If requesting data for a version higher than the current server, the database monitor only returns data for its level.

**Note:** If `SQLM_DBMON_VERSION1` is specified as the version, the APIs cannot be run remotely.

### *reserved*

Input. Reserved for future use. Must be set to `NULL`.

### *sqlma\_ptr*

Input. Pointer to the user-allocated *sqlma* (monitor area) structure. This structure specifies the **snapshot requests** to be returned by this invocation of `sqlmonss()`.

### *buffer\_length*

Input. The length of the data buffer. You may want to first call `sqlmonsz()` to estimate which size would be required for a given `sqlmonss()` invocation. However, if you will be issuing frequent `sqlmonss()` calls, and especially if peak activity is predictable on your system, then you will get better performance by allocating a fixed size buffer in your application.

### *buffer\_area*

Output. Pointer to the user-defined data buffer into which the snapshot information will be returned.

### *collected*

Output. A pointer to the *sqlm\_collected* structure which provides information about the server and the number of top-level structures returned in the output buffer area.

### *sqlca*

Output. A pointer to the *sqlca* structure where error information is returned. For more information see the *API Reference*.

## Snapshot Requests Supported

The following table lists the Snapshot request types that are supported.

API request type	Data structures that may be returned (Record type)	Information returned
<code>SQLMA_APPLINFO_ALL</code>	<code>sqlm_applinfo</code> ( <code>SQLM_APPLINFO_SS</code> )	Application identification information for all applications currently connected to a database that is managed by the DB2 instance on the node where snapshot is taken.

## sqlmonss - Get Snapshot API

API request type	Data structures that may be returned (Record type)	Information returned
SQLMA_DBASE_APPLINFO	sqlm_applinfo (SQLM_APPLINFO_SS)	Same as SQLMA_APPLINFO_ALL for each application currently connected to the specified database.
SQLMA_DCS_APPLINFO_ALL	sqlm_dcs_applinfo (SQLM_DCS_APPLINFO_SS)	Application identification info for all DCS applications currently connected to a database that is managed by the DB2 instance on the node where snapshot is taken.
SQLMA_DB2	sqlm_db2 (SQLM_DB2_SS)	Database manager level information, including internal monitor switch settings.
	sqlm_fcm (SQLM_FCM_SS) sqlm_fcm_node (SQLM_NODE_SS)	Fast communication manager information for each node in a partitioned database that this node is communicating with.
SQLMA_DBASE	sqlm_dbase (SQLM_DBASE_SS)	Database level information and counters for a database. Information is returned only if there is at least one application connected to the database..
	sqlm_rollfwd_info (SQLM_DBASE_ROLLFWD_SS) sqlm_rollfwd_ts_info (SQLM_DBASE_ROLLFWD_TS_SS)	Rollforward information if a rollforward is in progress.
SQLMA_DBASE_ALL		Same as SQLMA_DBASE for each database active on the node.
SQLMA_APPL	sqlm_appl (SQLM_APPL_SS) sqlm_lock_wait (SQLM_LOCK_WAIT_SS)	Application level information, includes cumulative counters, status information. Lock information for every agent working for this application that is waiting for a lock.
	sqlm_stmt (SQLM_STMT_SS)	Statement information for each open cursor and the last statement executed.
	sqlm_subsection (SQLM_SUBSECTION_SS)	Subsection information that immediately follows its parent sqlm_stmt.
	sqlm_subagent (SQLM_SUBAGENT_SS)	In partitioned databases, an sqlm_subagent always follows its parent sqlm_subsection or sqlm_stmt. An sqlm_subagent is not returned for a coordinator agent.
SQLMA_AGENT_ID		Same as SQLMA_APPL.
SQLMA_DBASE_APPLS		Same as SQLMA_APPL, for each application that is connected to the database on the node.
SQLMA_APPL_ALL		Same as SQLMA_APPL. for each application that is active on the node.

API request type	Data structures that may be returned (Record type)	Information returned
SQLMA_DBASE_TABLES	sqlm_table_header (SQLM_TABLE_HEADER_SS) sqlm_table (SQLM_TABLE_SS)	Table activity information for each table that was accessed, sqlm_table_header followed by an sqlm_table for each table.
SQLMA_APPL_LOCKS	sqlm_appl_lock (SQLM_APPL_LOCK_SS) sqlm_lock_wait (SQLM_LOCK_WAIT_SS) sqlm_lock (SQLM_LOCK_SS)	List of locks held by the application, and any lock wait information, sqlm_appl_lock is followed by sqlm_lock_wait if an application is in a lock wait, followed by an sqlm_lock for each lock held.
SQLMA_APPL_LOCKS_AGENT_ID		Same as SQLMA_APPL_LOCKS.
SQLMA_DBASE_LOCKS	sqlm_dbase_lock (SQLM_DBASE_LOCK_SS) (SQLMA_APPL_LOCKS)	Lock information at the database level, and application level for each application connected to the database, followed by an SQLMA_APPL_LOCKS for each application.
SQLMA_DBASE_TABLESPACES	sqlm_tablespace_header (SQLM_TABLESPACE_HEADER_SS) sqlm_tablespace (SQLM_TABLESPACE_SS)	Information about table space activity the database level, the application level for each application connected to the database, and the table space level for each table space that has been accessed by an application connected to the database, sqlm_tablespace_header followed by an sqlm_tablespace for each table space.
SQLMA_BUFFERPOOLS_ALL	sqlm_bufferpool (SQLM_BUFFERPOOL_SS)	Bufferpool activity counters.
SQLMA_DBASE_BUFFERPOOLS		Same as SQLMA_BUFFERPOOLS_ALL, but for specified database only.

### Specifying the Snapshot Requests

An invocation of sqlmonss() can specify several requests.

The *sqlma* supplied as input argument to sqlmonss() contains an array of **sqlm\_obj\_struct**. Each sqlm\_obj\_struct is an individual snapshot request.

sqlm\_obj\_struct is defined as follows:

```
typedef struct sqlm_obj_struct      /* SNAPSHOT REQUEST */
{
    unsigned long agent_id;         /* used for requests based on agentid */
    unsigned long obj_type;        /* Snapshot Request Type (SQLMA_XXXX) */
    char          object[SQLM_OBJECT_SZ]; /* used for requests based on object */
                                           /* name, such as 'get snapshot for database' */
}sqlm_obj_struct;
```

## sqlmonss - Get Snapshot API

Where *agent\_id* and *object* are only required if applicable for the request type, and are mutually exclusive. For example: a database name is required when the type is SQLMA\_DBASE\_LOCKS (get snapshot for locks on database), whereas an *agent\_id* is required when the type is SQLMA\_APPL\_LOCKS\_AGENT\_ID. Both *agent\_id* and *object* are ignored for requests such as SQLMA\_APPLINFO\_ALL (list applications).

Note that *agent\_id* is the **application handle** for an application. It does not correspond to any Operating System process Id (it is named this way for source compatibility with older releases of DB2).

### Setting up the sqlma and issuing the snapshot call

The following example sets up the *sqlma* for a call to *sqlmonss()* that requests two different snapshots. The first request requires an object name, the database alias, the second request requires an *agent\_id*, the application handle:

```
#include "string.h"
#include "stdlib.h"
#include "stdio.h"
#include "sqlutil.h"
#include "sqlmon.h" // System Monitor interface
main() {
    struct sqlca sqlca;
    int rc;

    #define BUFFER_SZ 4096 // Use a fixed size output buffer
    char snap_buffer[BUFFER_SZ]; // Snapshot output buffer
    sqlm_collected collected;

    //-----
    // Request SQLMA_DBASE, and SQLMA_APPL_LOCKS_AGENT_ID in the sqlma
    //-----
    unsigned long agent_id = 0; // STUB: Obtain by issuing 'list application'

    // Allocate the variable size sqlma structure
    struct sqlma* sqlma = (struct sqlma *) malloc(SQLMASIZE(2));

    // Request 2 different snapshots in same call
    sqlma->obj_num = 2;
    sqlma->obj_var[0].obj_type = SQLMA_DBASE;
    strcpy(sqlma->obj_var[0].object, "SAMPLE");

    sqlma->obj_var[1].obj_type = SQLMA_APPL_LOCKS_AGENT_ID;
    sqlma->obj_var[1].agent_id = agent_id;

    //-----
    // Perform the snapshot
    //-----
    rc = sqlmonss(SQLM_DBMON_VERSION5, NULL, sqlma,
                 BUFFER_SZ, snap_buffer,
                 &collected,
                 &sqlca);
    if (sqlca.sqlcode < 0) {
        // get and display a printable error message
        char msg[1024];
        sqlaintp(msg, sizeof(msg), 60, &sqlca);
        printf("%s", msg);
    }
    free(sqlma);
    return rc;
}
```

Application handles can be retrieved by issuing an SQLMA\_APPLINFO\_ALL request (list applications). An application connecting to the database can also retrieve its application handle (agent\_id) from the sqlca of the CONNECT request (See "Obtaining application handle (AGENT\_ID) from the CONNECT request").

### Reading the Snapshot Output Buffer

The sqlmonss() routine returns data as contiguous data structures in the user supplied buffer.

The sqlmon.h header file contains the definitions for all records returned by the sqlmonss() routine. It is the first place you should look for record information. It contains comments that explain how records are laid out in the output buffer. You may want to print a copy of this file, and reference it as you read this section.

The data structures returned in the snapshot output buffer are arranged in a two-level hierarchy:

- top-level structures
- secondary-level structures. These always come in the buffer following their parent top-level structure.

Each record contains a field that specifies its **type** and a **size** field that specifies its total size in bytes. The size **must be used** to read and skip this record in the output buffer. For example, following an SQLMA\_DB2 snapshot request (GET SNAPSHOT FOR DATABASE MANAGER) on a parallel system. The buffer could contain:

buffer_ptr →			
record 1 sqlm_db2	size: 204 type: SQLM_DB2_SS	top- level structure	
	num_sec_dbm_structs: 3		
record 2 sqlm_fcm	size: 44 type: SQLM_FCM_SS	secondary- level structure	
record 3 sqlm_fcm_node	size: 20 type: SQLM_NODE_SS	secondary- level structure	
record 4 sqlm_fcm_node	size: 20 type: SQLM_NODE_SS	secondary- level structure	

The **sqlm\_collected** output structure indicates the number of top-level structures returned in the buffer. Each top-level structure indicates the number of secondary-level structures that may follow it.

**Attention:** WARNING: It is imperative that you always use the size field for skipping a record in the output buffer. **Never use sizeof() on a snapshot record.**

Your application should also always read the record type. For some snapshot requests, the order in which records are returned is not guaranteed, and some record may not be

## sqlmonss - Get Snapshot API

returned when a monitor switch is OFF (see “sqlmrset - Reset Monitor API” on page 263 and “sqlmon - Get/Update Monitor Switches API” on page 244).

### Loop for reading snapshot output buffer

The following code illustrates how an application should be reading and skipping the records returned in the snapshot output buffer.

```
#include "stdlib.h"
#include "stdio.h"
#include "sqlutil.h"
#include "string.h"
#include "sqlmon.h" // System Monitor interface
//-----
// PROCESS EACH RECORD THAT MAY BE RETURNED IN THE SNAPSHOT OUTPUT BUFFER
//-----
while (collected.num_top_level_structs--) {

    // Check the record type, (5th byte of any top-level structure)
    switch ((unsigned char) *(snap+sizeof(unsigned long))) {
    case SQLM_DB2_SS: {
        sqlm_db2 *db2_snap;
        db2_snap = (sqlm_db2*) snap;
        // Process the database manager snapshot
        printf("Processing database manager snapshot\n");
        // ...

        // Skip all its records in the output buffer
        snap += db2_snap->size;

        // Skip the secondary level entries
        while (db2_snap->num_sec_dbm_structs--) snap+=*(unsigned long*)snap;
        } break;
    case SQLM_DBASE_SS: {
        sqlm_dbase *db_snap = (sqlm_dbase*) snap;
        // Process the snapshot ...
        printf("Processing database snapshot\n");

        // Skip the database snapshot and any secondary structures
        snap += db_snap->size;
        while (db_snap->num_sec_dbase_structs--) snap+=*(unsigned long*)snap;
        } break;
    case SQLM_APPL_SS: {
        sqlm_appl *appl_snap = (sqlm_appl*) snap;
        printf("Processing application snapshot\n");
        while (appl_snap->num_sec_appl_structs--) snap+=*(unsigned long*)snap;
        } break;
    case SQLM_APPLINFO_SS: {
        sqlm_applinfo *appinfo_snap = (sqlm_applinfo*) snap;
        printf("Processing list application\n");
        snap+=appinfo_snap->size;
        } break;
    case SQLM_DCS_APPLINFO_SS: {
        sqlm_dcs_applinfo *dcs_snap = (sqlm_dcs_applinfo*) snap;
        printf("Processing list dcs application\n");
        snap+=dcs_snap->size;
        } break;
    case SQLM_TABLE_HEADER_SS: {
        sqlm_table_header *tabh_snap = (sqlm_table_header*) snap;
        int numtabs = tabh_snap->num_tables;
        printf("Processing list tables\n");
        // ...
        // Skip it in the output buffer
        snap += tabh_snap->size;
        while (numtabs--) {
```



## sqlmonss - Get Snapshot API

```
        sqlm_table *tab_snap= (sqlm_table*) snap;
        snap += tab_snap->size;
    }
    } break;
case SQLM_DBASE_LOCK_SS: {
    sqlm_dbase_lock *dbase_lock_snap = (sqlm_dbase_lock*) snap;
    printf("Processing snapshot for locks on database\n");
    dump_dbase_lock(stdout,, dbase_lock_snap);
    // this routine provided the following section Printing Snapshot Output Records

    // Skip it in the snapshot output buffer
    snap = skip_db_lock_snap(dbase_lock_snap);
    } break;
case SQLM_APPL_LOCK_SS: {
    sqlm_appl_lock* appl_lock_snap = (sqlm_appl_lock*) snap;
    printf("Processing snapshot for locks for application\n");

    // Skip it in the snapshot output buffer
    snap = skip_appl_lock_snap(appl_lock_snap);
    } break;
case SQLM_TABLESPACE_HEADER_SS: {
    sqlm_tablespace_header *tspace_snap = (sqlm_tablespace_header*) snap;
    printf("Processing snapshot for tablespaces\n");

    // Skip it in the snapshot output buffer
    snap = skip_tspace_snap(tspace_snap);
    } break;
default:
    printf("%s:%d: Unexpected record type %d in snapshot output buffer!\n",
        __FILE__, __LINE__, (unsigned char) *(snap+sizeof(unsigned long)));
    } // end check the current snapshot buffer structure
} // end while there are top-level structures in the snapshot output buffer
```

Note an anomaly to the interface: some top-level data structures do not have a field that returns the number of all generic secondary structures that follows. Instead, they specify the number of structures returned for a specific secondary-level type. Special treatment is required for these snapshot requests. For example, the following routines, referenced in the main loop above, skip lock and table space snapshot requests in the snapshot output buffer.

### Routines to skip specific snapshot requests in the snapshot output buffer

```
//-----
// Skip a get snapshot for locks for application (SQLM_APPL_LOCK_SS)
//-----
char * skip_appl_lock_snap(sqlm_appl_lock * appl_lock_snap) {
    char *snap= (char*) appl_lock_snap;

    int numlocks=appl_lock_snap->num_locks;
    snap+=appl_lock_snap->size;

    // Skip the lock entries for this application
    while (numlocks--) snap+=*(unsigned long*)snap;
    return snap;
} // end skip_appl_lock_snap

//-----
// Skip a get snapshot for locks for database (SQLM_DBASE_LOCK_SS)
//-----
char * skip_db_lock_snap(sqlm_dbase_lock * dbase_lock_snap) {
    char *snap= (char*) dbase_lock_snap;

    int numaplocks=dbase_lock_snap->num_appls;
    snap+=*(unsigned long*)snap;
```

## sqlmonss - Get Snapshot API

```
    // Skip the appl lock entries
    while (numaplocks--) {
        snap = skip_appl_lock_snap((sqlm_appl_lock*) snap);
    }
    return snap;
} // end skip_db_lock_snap

//-----
// Skip a get snapshot for table spaces (SQLM_DBASE_LOCK_SS)
//-----
char * skip_tspace_snap(sqlm_tablespace_header * ts_header) {
    char *snap= (char*) ts_header;

    int numtspace = ts_header->num_tablespace;
    snap+=ts_header->size; // Skip the header

    // Skip the sqlm_tablespace entries
    while (numtspace--) snap+=*(unsigned long*)snap;
    return snap;
} // end skip_db_lock_snap
```

## Printing Snapshot Output Records

**Attention:** No string is NULL-terminated.

All strings returned in snapshot output records are blank-padded up to their maximum length.

The following example illustrates how a GET SNAPSHOT FOR LOCKS on database can be printed.

```
//-----
// Print a Blank Padded String of maximum length SZ
//-----
// note: strings returned by sqlmonss are NOT NULL-TERMINATED, they are all
//       blank padded up to some maximum length.
//-----
#define dump_BPSTRING(fp, str, SZ) \
{ \
    int k=0; \
    while (str[k]!='&&k<SZ) k++; \
    if (k<SZ) str[k]='\0'; \
    fprintf(fp, #str": %0.*s\n", SZ, str); \
} \
//-----
// Map Application Status to a printable string
//-----
// note: These #define may be found in sqlmon.h
//-----
char* appl_status_string(int val) {
    switch (val) {
        case SQLM_CONNECTPEND:    return "SQLM_CONNECTPEND";
        case SQLM_CONNECTED:      return "SQLM_CONNECTED";
        case SQLM_UOWEXEC:        return "SQLM_UOWEXEC";
        case SQLM_UOWWAIT:        return "SQLM_UOWWAIT";
        case SQLM_LOCKWAIT:       return "SQLM_LOCKWAIT";
        case SQLM_COMMIT_ACT:     return "SQLM_COMMIT_ACT";
        case SQLM_ROLLBACK_ACT:   return "SQLM_ROLLBACK_ACT";
        case SQLM_RECOMP:         return "SQLM_RECOMP";
        case SQLM_COMP:           return "SQLM_COMP";
        case SQLM_INTR:           return "SQLM_INTR";
        case SQLM_DISCONNECTPEND: return "SQLM_DISCONNECTPEND";
        case SQLM_TPREP:          return "SQLM_TPREP";
        case SQLM_THCOMT:         return "SQLM_THCOMT";
    }
```

## sqlmonss - Get Snapshot API

```
    case SQLM_THABRT:      return "SQLM_THABRT";
    case SQLM_TEND:       return "SQLM_TEND";
    case SQLM_CREATE_DB:  return "SQLM_CREATE_DB";
    case SQLM_RESTART:    return "SQLM_RESTART";
    case SQLM_RESTORE:    return "SQLM_RESTORE";
    case SQLM_BACKUP:     return "SQLM_BACKUP";
    case SQLM_LOAD:       return "SQLM_LOAD";
    case SQLM_UNLOAD:     return "SQLM_UNLOAD";
    case SQLM_IOERROR_WAIT: return "SQLM_IOERROR_WAIT";
    case SQLM_QUIESCE_TABLESPACE: return "SQLM_QUIESCE_TABLESPACE";
  }
  return "";
} // end of appl_status_string
//-----
// Map lock_object_type to a printable string
//-----
char* lock_object_type_string(int val) {
  char *result="";
  switch (val) {
    case SQLM_TABLE_LOCK:      result = "SQLM_TABLE_LOCK";
    case SQLM_ROW_LOCK:        result = "SQLM_ROW_LOCK";
    case SQLM_INTERNAL_LOCK:   result = "SQLM_INTERNAL_LOCK";
    case SQLM_TABLESPACE_LOCK: result = "SQLM_TABLESPACE_LOCK";
    case 0:                    result = "No Lock wait";
  }
  return result;
} // end of lock_object_type_string
//-----
// Map lock_mode to a printable string
//-----
char* lock_mode_string(int val) {
  char result="";
  switch (val) {
    case SQLM_LNON: result = "NO";
    case SQLM_LOIS: result = "IS";
    case SQLM_LOIX: result = "IX";
    case SQLM_LOOS: result = "S";
    case SQLM_LSIX: result = "SIX";
    case SQLM_LOOX: result = "X";
    case SQLM_LOIN: result = "IN";
    case SQLM_LOOZ: result = "Z";
    case SQLM_LOOU: result = "U";
    case SQLM_LONS: result = "NS";
    case SQLM_LONX: result = "NX";
  }
  return result;
}
//-----
// Map lock_status to a printable string
//-----
inline char* lock_status_string(int val) {
  switch (val) {
    case SQLM_LRBGRNT: result = "Granted";
    case SQLM_LRBCONV: result = "Converting";
  }
  return result;
}
//-----
// Print an sqlm_dbase_lock
//-----
void dump_sqlm_dbase_lock(FILE* fp, sqlm_dbase_lock* db_lock) {
  fprintf(fp, "\nsqlm_dbase_lock contains: \n");
  fprintf(fp, "db_lock->info_type: %s\n",          "SQLM_DBASE_LOCK_SS");
  fprintf(fp, "db_lock->locks_held: %ld\n",       db_lock->locks_held);
  fprintf(fp, "db_lock->appls_cur_cons: %ld\n",   db_lock->appls_cur_cons);
  fprintf(fp, "db_lock->num_appls: %ld\n",        db_lock->num_appls);
}
```

## sqlmonss - Get Snapshot API

```
fprintf(fp,"db_lock->locks_waiting: %ld\n", db_lock->locks_waiting);
dump_BPSTRING(fp, db_lock->input_db_alias, SQLM_DBPATH_SZ);
dump_BPSTRING(fp, db_lock->db_name, SQLM_IDENT_SZ);
dump_BPSTRING(fp, db_lock->db_path, SQLM_DBPATH_SZ);
} // end of dump_sqlm_dbase_lock
//-----
// Print an sqlm_appl_lock
// routine referred to in earlier section 'Loop for reading snapshot output buffer'
//-----
void dump_sqlm_appl_lock(FILE* fp, sqlm_appl_lock* appl_lock) {
    fprintf(fp, "\nsqlm_appl_lock contains: \n");
    fprintf(fp, "appl_lock->info_type: %s\n", "SQLM_APPL_LOCK_SS");
    fprintf(fp, "appl_lock->agent_id: %6.6ld\n", appl_lock->agent_id);
    fprintf(fp, "appl_lock->appl_status: %s\n",
            appl_status_string(appl_lock->appl_status));
    fprintf(fp, "appl_lock->codepage_id: %ld\n", appl_lock->codepage_id);
    fprintf(fp, "appl_lock->locks_held: %ld\n", appl_lock->locks_held);
    fprintf(fp, "appl_lock->num_locks: %ld\n", appl_lock->num_locks);
    // Print the status change time, only if non-zero
    if (appl_lock->status_change_time.seconds) {
        fprintf(fp, "appl_lock->status_change_time: %s\n",
                ctime((time_t*) &appl_lock->status_change_time.seconds));
    } // end if status changed
    dump_BPSTRING(fp, appl_lock->appl_id, SQLM_APPLID_SZ);
    dump_BPSTRING(fp, appl_lock->sequence_no, SQLM_SEQ_SZ);
    dump_BPSTRING(fp, appl_lock->appl_name, SQLM_IDENT_SZ);
    dump_BPSTRING(fp, appl_lock->auth_id, SQLM_IDENT_SZ);
    dump_BPSTRING(fp, appl_lock->client_db_alias, SQLM_IDENT_SZ);
    // The following information is returned only if the application is in
    // a lock wait (otherwise it is zeroed, or set to blank spaces if a string):
    if (appl_lock->appl_status==SQLM_LOCKWAIT) {
        fprintf(fp, "appl_lock->lock_object_name: %ld\n",
                appl_lock->lock_object_name);
        fprintf(fp, "appl_lock->agent_id_holding_lk: %6.6ld\n",
                appl_lock->agent_id_holding_lk);
        fprintf(fp, "appl_lock->lock_object_type: %s\n",
                lock_object_type_string(appl_lock->lock_object_type));
        fprintf(fp, "appl_lock->table_file_id: %ld\n", appl_lock->table_file_id);
        dump_BPSTRING(fp, appl_lock->appl_id_holding_lk, SQLM_APPLID_SZ);
        dump_BPSTRING(fp, appl_lock->sequence_no_holding_lk, SQLM_SEQ_SZ);
        dump_BPSTRING(fp, appl_lock->table_name, SQLM_IDENT_SZ);
        dump_BPSTRING(fp, appl_lock->table_schema, SQLM_IDENT_SZ);
        dump_BPSTRING(fp, appl_lock->tablespace_name, SQLM_IDENT_SZ);
    } // end if this application is in a lock_wait
} // end of dump_sqlm_appl_lock

//-----
// Print an sqlm_lock
//-----
void dump_sqlm_lock(FILE* fp, sqlm_lock* lock) {
    fprintf(fp, "\nsqlm_lock contains: \n");
    fprintf(fp, "lock->info_type: %s\n", "SQLM_LOCK_SS");
    fprintf(fp, "lock->lock_object_type: %s\n",
            lock_object_type_string(lock->lock_object_type));
    // Print the object of this lock
    switch (lock->lock_object_type) {
    case SQLM_ROW_LOCK:
    case SQLM_TABLE_LOCK:
        dump_BPSTRING(fp, lock->table_name, SQLM_IDENT_SZ);
        dump_BPSTRING(fp, lock->table_schema, SQLM_IDENT_SZ);
        fprintf(fp, "lock->table_file_id: %ld\n", lock->table_file_id);
        break;
    case SQLM_TABLESPACE_LOCK:
        dump_BPSTRING(fp, lock->tablespace_name, SQLM_IDENT_SZ);
        break;
    case SQLM_INTERNAL_LOCK:
        break;
    }
```

## sqlmonss - Get Snapshot API

```
    }  
    fprintf(fp,"lock->lock_mode: %s\n", lock_mode_string(lock->lock_mode));  
    fprintf(fp,"lock->lock_status: %s\n",lock_status_string(lock->lock_status));  
    fprintf(fp,"lock->lock_object_name: %ld\n", lock->lock_object_name);  
} // end of dump_sqlm_lock
```

### Sample Programs

**C**            \sqllib\samples\c\dbsnap.c

### See Also

“sqlmon - Get/Update Monitor Switches API” on page 244

“sqlmonsz - Estimate Size Required for sqlmonss() Output Buffer API” on page 260

“sqlmrset - Reset Monitor API” on page 263.

“GET SNAPSHOT Command” on page 232.

---

## sqlmonsz - Estimate Size Required for sqlmonss() Output Buffer API

### Purpose

Estimates the buffer size needed by “sqlmonss - Get Snapshot API” on page 248.

### Scope

This API only affects the instance to which the calling application is attached.

### Authorization

One of the following:

```
sysadm
sysctrl
sysmaint
```

### Required Connection

Instance. To obtain information from a remote instance (or a different local instance), it is necessary to first attach to that instance. If an attachment does not exist, an implicit instance attachment is made to the node specified by the **DB2INSTANCE** environment variable.

### API Include File

```
sqlmon.h
sqlca.h
```

### C API Syntax

```
/* File: sqlmon.h */
/* API: Estimate Database System Monitor Buffer Size */
/* ... */
int SQL_API_FN
sqlmonsz (
    unsigned long    version,
    _SQLOLDCHAR     *reserved,
    sqlma            *sqlma_ptr,
    unsigned long    *buff_size,
    struct sqlca     *sqlca);
/* ... */
```

### API Parameters

*sqlca*

Output. A pointer to the *sqlca* structure where the database manager returns error information. See *API Reference* for more information.

<i>buff_size</i>	Output. A pointer to the returned estimated buffer size needed by the GET SNAPSHOT API.
<i>sqlma_ptr</i>	Input. Pointer to the user-allocated <i>sqlma</i> data structure containing an array of <b>sqlm_obj_struct</b> , each one being a request for monitored data (see “sqlmonss - Get Snapshot API” on page 248).
<i>reserved</i>	Reserved for future use. Must be set to NULL.
<i>version</i>	Input. Version ID of the database monitor data to collect. The database monitor only returns data that was available for the requested version. Set this parameter to one of the following symbolic constants: <ul style="list-style-type: none"> <li>• SQLM_DBMON_VERSION1</li> <li>• SQLM_DBMON_VERSION2</li> <li>• SQLM_DBMON_VERSION5</li> </ul> <p>If requesting data for a version higher than the current server, the database monitor only returns data for its level.</p> <p><b>Note:</b> If SQLM_DBMON_VERSION1 is specified as the version, the APIs cannot be run remotely.</p>

## Comments

The size returned by *sqlmonsz* is always a little larger than what is actually required.

This function generates a significant amount of overhead (it basically takes a snapshot that does not return any data). Allocating and freeing memory dynamically for every **sqlmonss** call is also expensive. If calling **sqlmonss** repeatedly, for example, when sampling data over a period of time, it may be preferable to allocate a buffer of fixed size (for example 32K), possibly from the stack, and continue reusing that same buffer, rather than call **sqlmonsz**. If SQLM\_RC\_BUFFER\_FULL is returned then you can call *sqlmonsz* and allocate a new larger buffer.

If the database system monitor finds no active databases or applications, it may return a buffer size of zero (if, for example, lock information related to a database that is not active is requested). Verify that the estimated buffer size returned by this API is non-zero before calling “sqlmonss - Get Snapshot API” on page 248.

## See Also

- “sqlmon - Get/Update Monitor Switches API” on page 244
- “sqlmonss - Get Snapshot API” on page 248
- “sqlmrset - Reset Monitor API” on page 263.

## Code Sample

The following example estimates the buffer size required when issuing a snapshot for locks, tables, and database level information.

```
/*
 * Database Monitor - Estimate Buffer Size for Snapshot
 */
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include "sqlca.h"
#include "sqlutil.h"
#include "sqlmon.h"
main () {
    char *dbname = "SAMPLE"; // Name of the database to monitor
    int rc;
    struct sqlca sqlca;
    unsigned long buffer_sz;
    //-----
    // Request SQLMA_DBASE, SQLM_DBASE_TABLES, and SQLMA_DBASE_LOCKS in sqlma
    //-----
    struct sqlma* sqlma = (struct sqlma *) malloc(SQLMASIZE(3));
    sqlma->obj_num = 3;
    sqlma->obj_var[0].obj_type = SQLMA_DBASE;
    strcpy(sqlma->obj_var[0].object, dbname);
    sqlma->obj_var[1].obj_type = SQLMA_DBASE_LOCKS;
    strcpy(sqlma->obj_var[1].object, dbname);
    sqlma->obj_var[2].obj_type = SQLMA_DBASE_TABLES;
    strcpy(sqlma->obj_var[2].object, dbname);
    //-----
    // Estimate Buffer size required for this request
    //-----
    sqlmonsz(SQLM_DBMON_VERSION5, NULL, sqlma, &buffer_sz, &sqlca);
    if (sqlca.sqlcode) { // note: Positive return codes indicate a Warning
        // get and display a printable error message
        char msg[1024];
        sqlaintp (msg, sizeof(msg), 60, &sqlca);
        printf("%s", msg);
    }
    else printf ("\nBuffer size required for this snapshot is: %d\n",buffer_sz);
    //-----
    // ...Take the Snapshot...
    //-----
    {
        char* buffer_ptr;
        sqlm_collected collected;
        buffer_ptr = (char *) malloc(buffer_sz); // Allocate the buffer
        rc = sqlmonss(SQLM_DBMON_VERSION5, NULL, sqlma, buffer_sz, buffer_ptr,
            &collected, &sqlca);
        // .... Process snapshot output in buffer_ptr ...
        delete buffer_ptr; // Free the buffer
    }
}
```



---

## sqlmrset - Reset Monitor API

### Purpose

Resets the database system monitor data of a specified database, or of all active databases for the application issuing the call.

### Scope

This API only affects the application making the call.

### Authorization

One of the following:

```
sysadm
sysctrl
sysmaint
```

### Required Connection

Instance. To reset the monitor switches for a remote instance (or a different local instance), it is necessary to first attach to that instance.

### API Include File

*sqlmon.h*

### C API Syntax

```
/* File: sqlmon.h */
/* API: Reset Monitor */
/* ... */
int SQL_API_FN
sqlmrset (
    unsigned long version,
    _SQLOLDCHAR *reserved,
    unsigned long reset_all,
    _SQLOLDCHAR *db_alias,
    struct sqlca *sqlca);
/* ... */
```

### API Parameters

*reset\_all*

Input. An input value of 0 (zero) resets monitor counters for the database specified in **db\_alias**. An input value of 1 resets monitor counters for all databases and at the database manager level (**db\_alias** is then ignored).

## sqlmrset - Reset Monitor API

*db\_alias*

Input. A null terminated string that identifies the alias of the database being reset. It is ignored if **reset\_all** is set to 1.

*sqlca*

Output. A pointer to the *sqlca* structure where the database manager returns error information.

*version*

Input. Version ID of the database monitor data to reset. Use the same value that you would specify for the sqlmonss call.

## Sample Programs

C \sqllib\samples\c\monreset.c

## Comments

Each process (attachment) has its own private view of the monitor data. If one user resets, or turns off a monitor switch, other users are not affected. When an application first calls any database monitor function, it inherits the default switch settings from the database manager configuration file. These settings can be overridden with "sqlmon - Get/Update Monitor Switches API" on page 244.

If all active databases are reset, some database manager information is also reset to maintain the consistency of the data that is returned.

This API cannot be used to selectively reset specific data items or specific monitor groups. However a specific group can be reset by turning its switch OFF and then ON, using "sqlmon - Get/Update Monitor Switches API" on page 244.

## See Also

"sqlmon - Get/Update Monitor Switches API" on page 244

"sqlmonss - Get Snapshot API" on page 248

"sqlmonsz - Estimate Size Required for sqlmonss() Output Buffer API" on page 260.

"RESET MONITOR Command" on page 241.

## Code Sample

The following example resets the resettable data elements for a single database.

```
/*
 Database Monitor - Reset Data Elements
 */
#include <stdio.h>
#include "sqlca.h"
#include "sqlutil.h"
#include "sqlmon.h"
//-----
// Reset monitored data for a single database
//-----
struct sqlca sqlca;
sqlmrset(SQLM_DBMON_VERSION5, NULL, SQLM_OFF, "SAMPLE", &sqlca);
if (sqlca.sqlcode) { // note: Positive return codes indicate a Warning
 // get and display a printable error message
 char msg[1024];
 sqlaintp(msg, sizeof(msg), 60, &sqlca);
 printf("%s", msg);
}
```

## sqlmrset - Reset Monitor API

```
} // end if reset was successful  
else printf("Database Monitor Reset for '%s' was successful!\n", argv[1]);
```

## UPDATE MONITOR SWITCHES Command

---

### UPDATE MONITOR SWITCHES Command

#### Purpose

Turns one or more database monitor recording switches on or off.

The database manager records a base set of information at all times. Users who require more than this basic information can turn on the appropriate switches, but at a cost to system performance. Switches control information that is expensive to collect. See Chapter 3, “Database System Monitor Data Elements” on page 31 to determine if a switch is required for a particular data element.

#### Authorization


One of the following:

*sysadm*  
*sysctrl*  
*sysmaint*

#### Required Connection

Instance.

#### Format



#### Parameters

**USING** *switch-name*

Elements under switch control are grouped as follows:

<b>BUFFERPOOL</b>	Buffer pool activity information
<b>LOCK</b>	Lock information
<b>SORT</b>	Sorting information
<b>STATEMENT</b>	SQL statement information.
<b>TABLE</b>	Table activity information
<b>UOW</b>	Unit of work information.

#### Comments

To view the switch settings for your session, use “GET MONITOR SWITCHES Command” on page 230.

## UPDATE MONITOR SWITCHES Command

To clear the information related to a particular switch, set the switch off, then on.

When the application starts, it inherits the settings in the database manager configuration. See GET DATABASE MANAGER CONFIGURATION in the *Command Reference* and *dft\_mon\_xxx* configuration parameters.

### See Also

“RESET MONITOR Command” on page 241

“GET MONITOR SWITCHES Command” on page 230

“UPDATE MONITOR SWITCHES Command” on page 266

“GET SNAPSHOT Command” on page 232

“sqlmonss - Get Snapshot API” on page 248

## UPDATE MONITOR SWITCHES Command

---

## Appendix B. Parallel Edition Version 1.2 Users

In DB2 Version 5 the database system monitor interface has been simplified and is now the same for all database and system configurations. This harmonization of the interface means that some of the request types that were available with the Parallel Edition (PE) V1.2 system monitor are no longer supported.

The most significant change affects how you monitor an application. In DB2 Version 5 an application snapshot returns all the relevant application information, including a breakdown of the application statistics at the subsection or agent level (if applicable). For example, assuming an application is running a query composed of several subsections, a GET SNAPSHOT FOR APPLICATION will return:

- Lock wait information for each agent that is working for this application and is waiting for a lock.
- Tablequeue activity for each subsection executed by this application. This allows you to track progression of a query that is against a partitioned database.
- A list of process IDs or thread IDs for each agent associated with the application.

This information is available on both the coordinator and non-coordinator nodes. In PE V1.2, you would have to request information about individual agents or tablequeues and correlate the output obtained at these levels with the application.

**Note:** PE V1.2 applications are not compatible with DB2 Version 5.

PE V1.2 applications that are not using any of the requests that have are obsolete in DB2 Version 5 can be recompiled after changing the request type from SQLM\_DBMON\_PARALLEL1 to SQLM\_DBMON\_VERSION1. No other changes should be required. See the following tables for obsolete requests.

### agent\_id

You should note that **agent\_id** no longer corresponds to the process ID of the agent process. This field has not been renamed in the API to ensure source compatibility with previous versions, however it has become a globally unique identifier for the application.

**Agent ID and application handle are synonymous.** See “Partitioned Database Considerations” on page 23 for more information.

---

## API Changes

Obsolete sqlmonss() Request Type	Description	Replacement
SQLMA_AGENT_APPL SQLMA_AGENT_AGENTID	Get snapshot for <b>agent</b>	Replaced with SQLMA_APPL which will report a breakdown per agent, if and when applicable.
SQLMA_COORD_AGENTS	List all coordinator agents	Replaced with SQLMA_APPLINFO_ALL, which returns sqlm_applinfo for each application. It identifies the node where the coordinator agent runs and provides both its application handle and agent thread or process ID.
SQLMA_FCM_NODE_ALL SQLMA_FCM_NODE	Get Fast Communication Manager	Replaced with SQLMA_DB2, which gets <b>all</b> database manager information and returns FCM information (if applicable).
SQLMA_AGENT_ALL SQLMA_COORD_AGENTS	Get snapshot for all agents Get snapshot for coordinator agents	In PE V1.2, returned an SQLMA_AGENT_AGENTID snapshot for all agents, including the coordinators (or the coordinators only). Replaced with SQLMA_APPL_ALL (GET SNAPSHOT FOR APPLICATIONS). Note that information is not returned for agents that are not associated with any applications, as their counts would be zeroes.
SQLMA_DBASE_AGENTS	Get snapshot for all agents for a database	Replaced with SQLMA_DBASE_APPLS.

---

## Obsolete Commands

Obsolete PE V1.2 Command	Replacement
get snapshot for all agents	get snapshot for all applications
get snapshot for all coord agents	get snapshot for all applications
get snapshot for agents on dbname	get snapshot for applications on dbname
get snapshot for agents for application	get snapshot for application
get snapshot for coordinating agent	get snapshot for application
get snapshot for tablequeues	get snapshot for application

**Note:** GET SNAPSHOT FOR FCM is still supported, however the command processor maps it to a GET SNAPSHOT FOR DBM and extracts the FCM information from the returned output.



---

## Appendix C. DB2 Version 1 sqliestat Users

The following information previously available with the sqliestat API on OS/2 for DB2 Version 1 is now available through snapshot monitoring.

<b>sqliestat Name</b>	<b>Data Element</b>
component_id	"Product Identification" on page 38
corr_serv_lvl	"Service Level" on page 37
curr_reqs_lvl	"SQL Requests Since Last Commit" on page 173
db_type	"Server Operating System" on page 37
location	"Database Location" on page 43
node	"Catalog Node Network Name" on page 42
product_name	"Product Name" on page 38



---

## Appendix D. How the DB2 Library Is Structured

The DB2 Universal Database library consists of SmartGuides, online help, and books. This section describes the information that is provided, and how to access it.

To help you access product information online, DB2 provides the Information Center on OS/2, Windows 95, and the Windows NT operating systems. You can view task information, DB2 books, troubleshooting information, sample programs, and DB2 information on the Web. "About the Information Center" on page 280 has more details.

---

### SmartGuides

SmartGuides help you complete some administration tasks by taking you through each task one step at a time. SmartGuides are available on OS/2, Windows 95, and the Windows NT operating systems. The following table lists the SmartGuides.

---

<b>SmartGuide</b>	<b>Helps you to...</b>	<b>How to Access...</b>
<i>Add Database</i>	Catalog a database on a client workstation.	From the Client Configuration Assistant, click on <b>Add</b> .
<i>Create Database</i>	Create a database, and to perform some basic configuration tasks.	From the Control Center, click with the right mouse button on the <b>Databases</b> icon and select <b>Create-&gt;New</b> .
<i>Performance Configuration</i>	Tune the performance of a database by updating configuration parameters to match your business requirements.	From the Control Center, click with the right mouse button on the database you want to tune and select <b>Configure performance</b> .
<i>Backup Database</i>	Determine, create, and schedule a backup plan.	From the Control Center, click with the right mouse button on the database you want to backup and select <b>Backup-&gt;Database using SmartGuide</b> .
<i>Restore Database</i>	Recover a database after a failure. It helps you understand which backup to use, and which logs to replay.	From the Control Center, click with the right mouse button on the database you want to restore and select <b>Restore-&gt;Database using SmartGuide</b> .
<i>Create Table</i>	Select basic data types, and create a primary key for the table.	From the Control Center, click with the right mouse button on the <b>Tables</b> icon and select <b>Create-&gt;Table using SmartGuide</b> .
<i>Create Table Space</i>	Create a new table space.	From the Control Center, click with the right mouse button on the <b>Table spaces</b> icon and select <b>Create-&gt;Table space using SmartGuide</b> .

---

---

## Online Help

Online help is available with all DB2 components. The following table describes the various types of help.

Type of Help	Contents	How to Access...
<i>Command Help</i>	Explains the syntax of commands in the command line processor.	From the command line processor in interactive mode, enter:  <i>? command</i>  where <i>command</i> is a keyword or the entire command.  For example, <i>? catalog</i> displays help for all the CATALOG commands, whereas <i>? catalog database</i> displays help for the CATALOG DATABASE command.
<i>Control Center Help</i>	Explains the tasks you can perform in a window or notebook. The help includes prerequisite information you need to know, and describes how to use the window or notebook controls.	From a window or notebook, click on the <b>Help</b> push button or press the F1 key.
<i>Message Help</i>	Describes the cause of a message number, and any action you should take.	From the command line processor in interactive mode, enter:  <i>? message number</i>  where <i>message number</i> is a valid message number.  For example, <i>? SQL30081</i> displays help about the SQL30081 message.  To view message help one screen at a time, enter:  <i>? XXXnnnnn   more</i>  where <i>XXX</i> is the message prefix, such as SQL, and <i>nnnn</i> is the message number, such as 30081.  To save message help in a file, enter:  <i>? XXXnnnnn &gt; filename.ext</i>  where <i>filename.ext</i> is the file where you want to save the message help.  <b>Note:</b> On UNIX-based systems, enter:  <i>\? XXXnnnnn   more</i> or  <i>\? XXXnnnnn &gt; filename.ext</i>

---

Type of Help	Contents	How to Access...
<i>SQL Help</i>	Explains the syntax of SQL statements.	<p>From the command line processor in interactive mode, enter:</p> <p><b>help</b> <i>statement</i></p> <p>where <i>statement</i> is an SQL statement.</p> <p>For example, <b>help</b> <i>SELECT</i> displays help about the SELECT statement.</p>
<i>SQLSTATE Help</i>	Explains SQL states and class codes.	<p>From the command line processor in interactive mode, enter:</p> <p><b>?</b> <i>sqlstate</i> or <b>?</b> <i>class-code</i></p> <p>where <i>sqlstate</i> is a valid five digit SQL state and <i>class-code</i> is a valid two digit class code.</p> <p>For example, <b>?</b> <i>08003</i> displays help for the 08003 SQL state, whereas <b>?</b> <i>08</i> displays help for the 08 class code.</p>

---

## DB2 Books

The table in this section lists the DB2 books. They are divided into two groups:

- Cross-platform books: These books are for DB2 on any of the supported platforms.
- Platform-specific books: These books are for DB2 on a specific platform. For example, there is a separate *Quick Beginnings* book for DB2 on OS/2, Windows NT, and UNIX-based operating systems.

Most books are available in HTML and PostScript format, and in hardcopy that you can order from IBM. The exceptions are noted in the table.

You can obtain DB2 books and access information in a variety of different ways:

- View** To view an HTML book, you can do the following:
- If you are running DB2 administration tools on OS/2, Windows 95, or the Windows NT operating systems, you can use the Information Center. “About the Information Center” on page 280 has more details.
  - Use the open file function of the Web browser supplied by DB2 (or one of your own) to open the following page:  
`sqllib/doc/html/index.htm`  
The page contains descriptions of and links to the DB2 books. The path is located on the drive where DB2 is installed.  
You can also open the page by double-clicking on the **DB2 Online Books** icon. Depending on the system you are using, the icon is in the main product folder or the Windows Start menu.
- Search** To search for information in the HTML books, you can do the following:
- Click on **Search the DB2 Books** at the bottom of any page in the HTML books. Use the search form to find a specific topic.
  - Click on **Index** at the bottom of any page in an HTML book. Use the Index to find a specific topic in the book.
  - Display the Table of Contents or Index of the HTML book, and then use the find function of the Web browser to find a specific topic in the book.
  - Use the bookmark function of the Web browser to quickly return to a specific topic.
  - Use the search function of the Information Center to find specific topics. “About the Information Center” on page 280 has more details.
- Print** To print a book on a PostScript printer, look for the file name shown in the table.
- Order** To order a hardcopy book from IBM, use the form number.

<b>Book Name</b>	<b>Book Description</b>	<b>Form Number File Name</b>
<b>Cross-Platform Books</b>		
<i>Administration Getting Started</i>	Introduces basic DB2 database administration concepts and tasks, and walks you through the primary administrative tasks.	S10J-8154 db2k0x50
<i>Administration Guide</i>	Contains information required to design, implement, and maintain a database to be accessed either locally or in a client/server environment.	S10J-8157 db2d0x50
<i>API Reference</i>	Describes the DB2 application programming interfaces (APIs) and data structures you can use to manage your databases. Explains how to call APIs from your applications.	S10J-8167 db2b0x50
<i>CLI Guide and Reference</i>	Explains how to develop applications that access DB2 databases using the DB2 Call Level Interface, a callable SQL interface that is compatible with the Microsoft ODBC specification.	S10J-8159 db2l0x50
<i>Command Reference</i>	Explains how to use the command line processor, and describes the DB2 commands you can use to manage your database.	S10J-8166 db2n0x50
<i>DB2 Connect Enterprise Edition Quick Beginnings</i>	Provides planning, installing, configuring, and using information for DB2 Connect Enterprise Edition. Also contains installation and setup information for all supported clients.	S10J-7888 db2cyx50
<i>DB2 Connect Personal Edition Quick Beginnings</i>	Provides planning, installing, configuring, and using information for DB2 Connect Personal Edition.	S10J-8162 db2c1x50
<i>DB2 Connect User's Guide</i>	Provides concepts, programming and general using information about the DB2 Connect products.	S10J-8163 db2c0x50
<i>DB2 Connectivity Supplement</i>	Provides setup and reference information for customers who want to use DB2 for AS/400, DB2 for OS/390, DB2 for MVS, or DB2 for VM as DRDA Application Requesters with DB2 Universal Database servers, and customers who want to use DRDA Application Servers with DB2 Connect (formerly DDCS) application requesters.  <b>Note:</b> Available in HTML and PostScript formats only.	No form number db2h1x50
<i>Embedded SQL Programming Guide</i>	Explains how to develop applications that access DB2 databases using embedded SQL, and includes discussions about programming techniques and performance considerations.	S10J-8158 db2a0x50
<i>Glossary</i>	Provides a comprehensive list of all DB2 terms and definitions.  <b>Note:</b> Available in HTML format only.	No form number db2t0x50

<b>Book Name</b>	<b>Book Description</b>	<b>Form Number</b> <b>File Name</b>
<i>Installing and Configuring DB2 Clients</i>	Provides installation and setup information for all DB2 Client Application Enablers and DB2 Software Developer's Kits.  <b>Note:</b> Available in HTML and PostScript formats only.	No form number db2iyx50
<i>Master Index</i>	Contains a cross reference to the major topics covered in the DB2 library.  <b>Note:</b> Available in PostScript format and hardcopy only.	S10J-8170 db2w0x50
<i>Message Reference</i>	Lists messages and codes issued by DB2, and describes the actions you should take.	S10J-8168 db2m0x50
<i>Replication Guide and Reference</i>	Provides planning, configuring, administering, and using information for the IBM Replication tools supplied with DB2.	S95H-0999 db2e0x50
<i>Road Map to DB2 Programming</i>	Introduces the different ways your applications can access DB2, describes key DB2 features you can use in your applications, and points to detailed sources of information for DB2 programming.	S10J-8155 db2u0x50
<i>SQL Getting Started</i>	Introduces SQL concepts, and provides examples for many constructs and tasks.	S10J-8156 db2y0x50
<i>SQL Reference</i>	Describes SQL syntax, semantics, and the rules of the language. Also includes information about release-to-release incompatibilities, product limits, and catalog views.	S10J-8165 db2s0x50
<i>System Monitor Guide and Reference</i>	Describes how to collect different kinds of information about your database and the database manager. Explains how you can use the information to understand database activity, improve performance, and determine the cause of problems.	S10J-8164 db2f0x50
<i>Troubleshooting Guide</i>	Helps you determine the source of errors, recover from problems, and use diagnostic tools in consultation with DB2 Customer Service.	S10J-8169 db2p0x50
<i>What's New</i>	Describes the new features, functions, and enhancements in DB2 Universal Database.  <b>Note:</b> Available in HTML and PostScript formats only.	No form number db2q0x50
<b>Platform-Specific Books</b>		
<i>Building Applications for UNIX Environments</i>	Provides environment setup information and step-by-step instructions to compile, link, and run DB2 applications on a UNIX system.	S10J-8161 db2axx50
<i>Building Applications for Windows and OS/2 Environments</i>	Provides environment setup information and step-by-step instructions to compile, link, and run DB2 applications on a Windows or OS/2 system.	S10J-8160 db2a1x50



<b>Book Name</b>	<b>Book Description</b>	<b>Form Number</b> <b>File Name</b>
<i>DB2 Extended Enterprise Edition Quick Beginnings</i>	Provides planning, installing, configuring, and using information for DB2 Universal Database Extended Enterprise Edition for AIX.	S72H-9620 db2v3x50
<i>DB2 Personal Edition Quick Beginnings</i>	Provides planning, installing, configuring, and using information for DB2 Universal Database Personal Edition on OS/2, Windows 95, and the Windows NT operating systems.	S10J-8150 db2i1x50
<i>DB2 SDK for Macintosh Building Your Applications</i>	Provides environment setup information and step-by-step instructions to compile, link, and run DB2 applications on a Macintosh system.  <b>Note:</b> Available in PostScript format and hardcopy for DB2 Version 2.1.2 only.	S50H-0528 sqla7x02
<i>DB2 SDK for SCO OpenServer Building Your Applications</i>	Provides environment setup information and step-by-step instructions to compile, link, and run DB2 applications on a SCO OpenServer system.  <b>Note:</b> Available for DB2 Version 2.1.2 only.	S89H-3242 sqla9x02
<i>DB2 SDK for Silicon Graphics IRIX Building Your Applications</i>	Provides environment setup information and step-by-step instructions to compile, link, and run DB2 applications on a Silicon Graphics system.  <b>Note:</b> Available in PostScript format and hardcopy for DB2 Version 2.1.2 only.	S89H-4032 sqlaax02
<i>DB2 SDK for SINIX Building Your Applications</i>	Provides environment setup information and step-by-step instructions to compile, link, and run DB2 applications on a SINIX system.  <b>Note:</b> Available in PostScript format and hardcopy for DB2 Version 2.1.2 only.	S50H-0530 sqla8x00
<i>Quick Beginnings for OS/2</i>	Provides planning, installing, configuring, and using information for DB2 Universal Database on OS/2. Also contains installing and setup information for all supported clients.	S10J-8147 db2i2x50
<i>Quick Beginnings for UNIX</i>	Provides planning, installing, configuring, and using information for DB2 Universal Database on UNIX-based platforms. Also contains installing and setup information for all supported clients.	S10J-8148 db2ixx50
<i>Quick Beginnings for Windows NT</i>	Provides planning, installing, configuring, and using information for DB2 Universal Database on the Windows NT operating system. Also contains installing and setup information for all supported clients.	S10J-8149 db2i6x50

**Notes:**

1. The character in the sixth position of the file name indicates the language of a book. For example, the file name db2d0e50 indicates that the *Administration Guide* is in English. The following letters are used in the file names to indicate the language of a book:

<b>Language</b>	<b>Identifier</b>	<b>Language</b>	<b>Identifier</b>
Brazilian Portuguese	B	Hungarian	H
Bulgarian	U	Italian	I
Czech	X	Norwegian	N
Danish	D	Polish	P
English	E	Russian	R
Finnish	Y	Slovenian	L
French	F	Spanish	Z
German	G	Swedish	S

2. For late breaking information that could not be included in the DB2 books, see the README file. Each DB2 product includes a README file which you can find in the directory where the product is installed.

---

## About the Information Center

The Information Center provides quick access to DB2 product information. The Information Center is available on OS/2, Windows 95, and the Windows NT operating systems. You must install the DB2 administration tools to see the Information Center.

Depending on your system, you can access the Information Center from the:

- Main product folder
- Toolbar in the Control Center
- Windows Start menu.

The Information Center provides the following kinds of information. Click on the appropriate tab to look at the information:

<b>Tasks</b>	Lists tasks you can perform using DB2.
<b>Reference</b>	Lists DB2 reference information, such as keywords, commands, and APIs.
<b>Books</b>	Lists DB2 books.
<b>Troubleshooting</b>	Lists categories of error messages and their recovery actions.
<b>Sample Programs</b>	Lists sample programs that come with the DB2 Software Developer's Kit. If the Software Developer's Kit is not installed, this tab is not displayed.
<b>Web</b>	Lists DB2 information on the World Wide Web. To access this information, you must have a connection to the Web from your system.

When you select an item in one of the lists, the Information Center launches a viewer to display the information. The viewer might be the system help viewer, an editor, or a Web browser, depending on the kind of information you select.

The Information Center provides search capabilities so you can look for specific topics, and filter capabilities to limit the scope of your searches.



---

## Appendix E. Notices

Any reference to an IBM licensed program in this publication is not intended to state or imply that only IBM's licensed program may be used. Any functionally equivalent product, program or service that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, is the user's responsibility.

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.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Canada Limited  
Department 071  
1150 Eglinton Ave. East  
North York, Ontario  
M3C 1H7  
CANADA

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

This publication may contain examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

---

## Trademarks

The following terms are trademarks or registered trademarks of the IBM Corporation in the United States and/or other countries:

ACF/VTAM	MVS/ESA
ADSTAR	MVS/XA
AISPO	NetView
AIX	OS/400
AIXwindows	OS/390
AnyNet	OS/2
APPN	PowerPC
AS/400	QMF
CICS	RACF
C Set++	RISC System/6000
C/370	SAA
DATABASE 2	SP
DatagLANce	SQL/DS
DataHub	SQL/400
DataJoiner	S/370
DataPropagator	System/370
DataRefresher	System/390
DB2	SystemView
Distributed Relational Database Architecture	VisualAge
DRDA	VM/ESA
Extended Services	VSE/ESA
FFST	VTAM
First Failure Support Technology	WIN-OS/2
IBM	
IMS	
Lan Distance	

---

## Trademarks of Other Companies

The following terms are trademarks or registered trademarks of the companies listed:

C-bus is a trademark of Corollary, Inc.

HP-UX is a trademark of Hewlett-Packard.

Java and HotJava are trademarks of Sun Microsystems, Inc.

Microsoft, Windows, Windows NT, and the Windows 95 logo are trademarks or registered trademarks of Microsoft Corporation.

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

Solaris is a trademark of Sun Microsystems, Inc.

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

Other company, product, or service names, which may be denoted by a double asterisk (\*\*), may be trademarks or service marks of others.

---

## Index

### A

- acc\_curs\_blk element 162
- accepted block cursor requests, monitor element 162
- accesses to overflowed records, monitor element 157
- activating an event monitor 15
- active sorts, monitor element 84
- active\_sorts element 84
- agent and connection data elements
  - agents assigned from pool 75
  - agents created due to empty agent pool 76
  - agents registered 73
  - agents waiting for a token 73
  - applications connected currently 72
  - applications executing in the database currently 72
  - committed private memory 78
  - connects since first database connect 71
  - local connections 69
  - local connections executing in the database manager 70
  - local databases with current connects 71
  - maximum number of agents registered 74
  - maximum number of associated agents 77
  - maximum number of coordinating agents 76
  - number of idle agents 75
  - remote connections executing in the database manager 69
  - remote connections to database manager 68
  - secondary connections 78
  - stolen agents 77
- agent ID holding lock, monitor element 145
- agent pool 67
- agent\_id 200
- agent\_id element 45
- agent\_id\_holding\_lock element 145
- agent\_pid element 66
- agent\_usr\_cpu\_time element 191
- agents
  - associated 67
  - coordinator 67
  - idle 67
  - subagent 67
- agents assigned from pool, monitor element 75
- agents created due to empty agent pool, monitor element 76
- agents registered, monitor element 73
- agents waiting for a token, monitor element 73
- agents\_created\_empty\_pool element 76
- agents\_from\_pool element 75
- agents\_registered element 73
- agents\_registered\_top element 74
- agents\_stolen element 77
- agents\_top element 190
- agents\_waiting\_on\_token element 73
- agents\_waiting\_top element 74
- appl\_con\_time element 62
- appl\_id element 50
- appl\_id\_holding\_lk element 146
- appl\_idle\_time element 66
- appl\_name element 49
- appl\_priority element 59
- appl\_priority\_type element 60
- appl\_section\_inserts element 126
- appl\_section\_lookups element 125
- appl\_status element 46
- application agent priority, monitor element 59
- application creator, monitor element 179
- application handle
  - See agent\_id
- application handle (agent ID), monitor element 45
- application ID holding lock, monitor element 146
- application ID, monitor element 50
- application identification data elements
  - application agent priority 59
  - application handle (agent ID) 45
  - application ID 50
  - application idle time 66
  - application name 49
  - application priority type 60
  - application status 46
  - application status change time 48
  - authorization ID 52
  - client communication protocol 58
  - client operating platform 58
  - client process ID 57
  - client product/version ID 53
  - configuration NNAME of client 53
  - connection request completion timestamp 62
  - connection request start timestamp 62
  - coordinating node 61
  - database alias used by application 54
  - database country code 59

- application identification data elements (*continued*)
  - drda correlation token 57
  - host product/version ID 55
  - ID of code page used by application 48
  - outbound application ID 55
  - outbound sequence number 56
  - previous transaction stop time 65
  - previous unit of work completion timestamp 62
  - process or thread id 66
  - sequence number 52
  - unit of work completion status 65
  - unit of work start timestamp 63
  - unit of work stop timestamp 64
  - user authorization level 60
  - user login ID 56
- application idle time, monitor element 66
- application name, monitor element 49
- application priority type, monitor element 60
- application snapshot 7
- application status change time, monitor element 48
- application status data elements
  - See application identification data elements
- application status, monitor element 46
- applications connected currently, monitor element 72
- applications executing in the database currently, monitor element 72
- appls\_cur\_cons 72
- appls\_in\_db2 element 72
- associated agent 67
- associated\_agents\_top element 77
- auth\_id element 52
- authority required
  - for event monitors 15, 195
  - for snapshot monitoring 6
- authority\_lvl element 60
- authorization ID, monitor element 52
- autostarting an event monitor 15
- availability of data
  - snapshot monitoring 10

## B

- binds\_precompiles element 174
- binds/precompiles attempted, monitor element 174
- blocked event monitors 202
- bp\_info element 109
- buff\_free element 85
- buff\_free\_bottom element 85
- buffer overflows
  - pipe 20

- buffer pool 90
- buffer pool activity data elements
  - buffer pool asynchronous data reads 101
  - buffer pool asynchronous data writes 102
  - buffer pool asynchronous index reads 104
  - buffer pool asynchronous index writes 103
  - buffer pool asynchronous read requests 107
  - buffer pool asynchronous read time 105
  - buffer pool asynchronous write time 106
  - buffer pool data logical reads 93
  - buffer pool data physical reads 94
  - buffer pool data writes 95
  - buffer pool index logical reads 96
  - buffer pool index physical reads 97
  - buffer pool index writes 98
  - buffer pool information 109
  - buffer pool log space cleaners triggered 107
  - buffer pool threshold cleaners triggered 109
  - buffer pool victim page cleaners triggered 108
  - database files closed 100
  - time waited for prefetch 110
  - total buffer pool physical read time 99
  - total buffer pool physical write time 100
- buffer pool asynchronous data reads, monitor element 101
- buffer pool asynchronous data writes, monitor element 102
- buffer pool asynchronous index reads, monitor element 104
- buffer pool asynchronous index writes, monitor element 103
- buffer pool asynchronous read requests, monitor element 107
- buffer pool asynchronous read time, monitor element 105
- buffer pool asynchronous write time, monitor element 106
- buffer pool data logical reads, monitor element 93
- buffer pool data pages from extended storage, monitor element 113
- buffer pool data pages to extended storage, monitor element 112
- buffer pool data physical reads, monitor element 94
- buffer pool data writes, monitor element 95
- buffer pool event monitor 18
- buffer pool hit ratio 91
- buffer pool index logical reads, monitor element 96
- buffer pool index pages from extended storage, monitor element 114



- buffer pool index pages to extended storage, monitor element 112
- buffer pool index physical reads, monitor element 97
- buffer pool index writes, monitor element 98
- buffer pool information, monitor element 109
- buffer pool log space cleaners triggered, monitor element 107
- buffer pool snapshot 7
- buffer pool threshold cleaners triggered, monitor element 109
- buffer pool victim page cleaners triggered, monitor element 108
- buffering, event monitor 201

## C

- cache
  - catalog 119
  - package 122
- capabilities, monitoring
  - activity monitoring 1
  - performance analysis 2
  - problem determination 1
  - system configuration 2
- cat\_cache\_heap\_full element 122
- cat\_cache\_inserts element 120
- cat\_cache\_lookups element 119
- cat\_cache\_overflows element 121
- catalog cache
  - See cache
- catalog cache data elements
  - catalog cache heap full 122
  - catalog cache inserts 120
  - catalog cache lookups 119
  - catalog cache overflows 121
- catalog cache heap full, monitor element 122
- catalog cache inserts, monitor element 120
- catalog cache lookups, monitor element 119
- catalog cache overflows, monitor element 121
- catalog node network name, monitor element 42
- catalog node number, monitor element 43
- catalog\_node element 43
- catalog\_node\_name element 42
- CE\_free element 87
- CE\_free\_bottom element 87
- client communication protocol, monitor element 58
- client operating platform, monitor element 58
- client process ID, monitor element 57
- client product/version ID, monitor element 53
- client\_db\_alias element 54
- client\_nname element 53
- client\_pid element 57
- client\_platform element 58
- client\_prdid element 53
- client\_protocol element 58
- codepage\_id element 48
- comm\_private\_mem element 78
- commit statements attempted, monitor element 166
- commit\_sql\_stmts element 166
- committed private memory, monitor element 78
- component\_id element 38
- con\_local\_databases 71
- configuration NNAME of client, monitor element 53
- configuration NNAME at monitoring (server) node, monitor element 34
- conn\_complete\_time element 62
- connection data elements
  - See application status data elements
- connection entries currently free, monitor element 87
- connection for snapshot 9
- connection information 7
- connection request completion timestamp, monitor element 62
- connection request start timestamp, monitor element 62
- connection status, monitor element 89
- connection\_status element 89
- connections event monitor 18
- connections involved in deadlock, monitor element 140
- connects since first database connect, monitor element 71
- coord\_agents\_top element 76
- coord\_node element 61
- coordinating node, monitor element 61
- coordinator agent 27, 67
- corr\_token element 57
- counters 20, 33
- country\_code element 59
- CPU time used, monitor element 191
- cpu usage data elements
  - CPU time used 191
- CREATE EVENT MONITOR statement 214
- creating an event monitor 15
- creator element 179
- current agents waiting on locks, monitor element 144
- current number of tablequeue buffers overflowed, monitor element 188
- cursor name, monitor element 178

cursor\_name element 178

## D

data definition language (DDL) SQL statements, monitor element 169

data element

definition 28

types 33

data elements (by API element name) 57

acc\_curs\_blk 162

active\_sorts 84

agent\_id 45

agent\_id\_holding\_lock 145

agent\_pid 66

agent\_sys\_cpu\_time 191

agent\_usr\_cpu\_time 191

agents\_created\_empty\_pool 76

agents\_from\_pool 75

agents\_registered 73

agents\_registered\_top 74

agents\_stolen 77

agents\_top 190

agents\_waiting\_on\_token 73

agents\_waiting\_top 74

appl\_con\_time 62

appl\_id 50

appl\_idle\_time 66

appl\_name 49

appl\_priority 59

appl\_priority\_type 60

appl\_section\_inserts 126

appl\_section\_lookups 125

appl\_status 46

appli\_id\_holding\_lk 146

appls\_cur\_con 72

appls\_in\_db2 72

associated\_agents\_top 77

auth\_id 52

authority\_lvl 60

binds\_precompiles 174

bp\_info 109

buff\_free 85

buff\_free\_bottom 85

cat\_cache\_heap\_full 122

cat\_cache\_inserts 120

cat\_cache\_lookups 119

cat\_cache\_overflows 121

catalog\_node 43

catalog\_node\_name 42

data elements (by API element name) *(continued)*

CE\_free 87

CE\_free\_bot 87

client\_db\_alias 54

client\_nname 53

client\_pid 57

client\_platform 58

client\_prdid 53

client\_protocol 58

codepage\_id 48

comm\_private\_mem 78

commit\_sql\_stmts 166

component\_id 38

con\_local\_databases 71

conn\_complete\_time 62

conn\_time 41

connection\_status 89

coord\_agents\_top 76

coord\_node 61

corr\_token 57

country\_code 59

creator 179

cursor\_name 178

db\_conn\_time 41

db\_heap\_top 126

db\_location 43

db\_name 39

db\_path 40

db\_status 42

db2\_status 39

db2start\_time 34

ddl\_sql\_stmts 169

deadlocks 133

degree\_parallelism 191

direct\_read\_reqs 117

direct\_read\_time 118

direct\_reads 115

direct\_write\_reqs 117

direct\_write\_time 118

direct\_writes 116

disconn\_time 41

dl\_conns 140

dynamic\_sql\_stmts 165

execution\_id 56

failed\_sql\_stmts 165

fetch\_count 182

files\_closed 100

host\_prdid 55

idle\_agents 75

input\_db\_alias 193

data elements (by API element name) *(continued)*

int\_auto\_rebinds 170  
int\_commits 171  
int\_deadlock\_rollbacks 173  
int\_rollbacks 172  
int\_rows\_deleted 157  
int\_rows\_inserted 159  
int\_rows\_updated 158  
last\_backup 44  
last\_reset 192  
local\_cons 69  
local\_cons\_in\_exec 70  
lock\_escals 134  
lock\_list\_in\_use 132  
lock\_mode 136  
lock\_object\_name 138  
lock\_object\_type 137  
lock\_status 137  
lock\_timeouts 139  
lock\_wait\_start\_time 145  
lock\_wait\_time 142  
lock\_waits 141  
locks\_held 131  
locks\_held\_top 139  
locks\_waiting 144  
log\_reads 129  
log\_space\_used 130  
log\_writes 130  
MA\_free 86  
MA\_free\_bot 86  
num\_agents 190  
num\_subagents 186  
number\_nodes 89  
open\_loc\_curs 163  
open\_loc\_curs\_blk 163  
open\_rem\_curs 160  
open\_rem\_curs\_blk 161  
operation 176  
outbound\_appl\_id 55  
outbound\_sequence\_no 56  
overflow\_accesses 157  
package\_name 177  
piped\_sorts\_accepted 81  
piped\_sorts\_requested 80  
pkg\_cache\_inserts 124  
pkg\_cache\_lookups 123  
pool\_async\_data\_read\_reqs 107  
pool\_async\_data\_reads 101  
pool\_async\_data\_writes 102  
pool\_async\_index\_reads 104

data elements (by API element name) *(continued)*

pool\_async\_index\_writes 103  
pool\_async\_read\_time 105  
pool\_async\_write\_time 106  
pool\_data\_from\_estore 113  
pool\_data\_l\_reads 93  
pool\_data\_p\_reads 94  
pool\_data\_to\_estore 112  
pool\_data\_writes 95  
pool\_drty\_pg\_steal\_clns 108  
pool\_drty\_pg\_thrsh\_clns 109  
pool\_index\_from\_estore 114  
pool\_index\_l\_reads 96  
pool\_index\_p\_reads 97  
pool\_index\_to\_estore 112  
pool\_index\_writes 98  
pool\_lsn\_gap\_clns 107  
pool\_read\_time 99  
pool\_write\_time 100  
post\_threshold\_sorts 80  
prefetch\_wait\_time 110  
prev\_stop\_time 65  
prev\_uow\_stop\_time 62  
product\_name 38  
query\_card\_estimate 183  
query\_cost\_estimate 183  
RB\_free 88  
RB\_free\_bot 88  
rej\_curs\_blk 162  
rem\_cons\_in 68  
rem\_cons\_in\_exec 69  
rf\_log\_num 149  
rf\_num\_tspaces 149  
rf\_status 149  
rf\_timestamp 148  
rf\_type 148  
rollback\_sql\_stmts 167  
rolled\_back\_appl\_id 147  
rows\_deleted 153  
rows\_inserted 153  
rows\_read 156  
rows\_selected 154  
rows\_updated 154  
rows\_written 155  
sec\_log\_used\_top 127  
sec\_logs\_allocated 129  
section\_number 177  
select\_sql\_stmts 168  
sequence\_no 52  
sequence\_no\_holding\_lk 147

data elements (by API element name) *(continued)*

server\_db2\_type 35  
server\_instance\_name 35  
server\_nname 34  
server\_platform 37  
server\_prdid 36  
server\_version 36  
service\_level 37  
sort\_heap\_allocated 79  
sort\_overflows 83  
sql\_reqs\_since\_commit 173  
sqlca 182  
ss\_exec\_time 185  
ss\_node\_number 184  
ss\_number 184  
ss\_status 185  
ss\_sys\_cpu\_time 191  
ss\_usr\_cpu\_time 191  
static\_sql\_stmts 164  
status 65  
status\_change\_time 48  
stmt\_node\_number 174  
stmt\_operation 176  
stmt\_sorts 181  
stmt\_start 179  
stmt\_stop 180  
stmt\_sys\_cpu\_time 191  
stmt\_text 180  
stmt\_type 175  
stmt\_usr\_cpu\_time 191  
stop\_time 180  
system\_cpu\_time 191  
table\_file\_id 159  
table\_name 151  
table\_schema 152  
table\_type 150  
tablespace\_name 143  
time\_stamp 193  
tot\_log\_used\_top 128  
total\_buffers\_rcvd 90  
total\_buffers\_sent 90  
total\_cons 71  
total\_sec\_cons 78  
total\_sort\_time 82  
total\_sorts 82  
tq\_cur\_send\_spills 188  
tq\_node\_waited\_for 187  
tq\_rows\_read 188  
tq\_rows\_written 189  
tq\_tot\_send\_spills 187

data elements (by API element name) *(continued)*

tq\_wait\_for\_any 186  
ts\_name 148  
uid\_sql\_stmts 168  
uow\_comp\_status 65  
uow\_lock\_wait\_time 144  
uow\_log\_space\_used 130  
uow\_start\_time 63  
uow\_stop\_time 64  
user\_cpu\_time 191  
x\_lock\_escals 135  
data elements (by name) 57  
accepted block cursor requests 162  
accesses to overflowed records 157  
active sorts 84  
agent ID holding lock 145  
agents assigned from pool 75  
agents created due to empty agent pool 76  
agents registered 73  
agents waiting for a token 73  
application agent priority 59  
application creator 179  
application handle (agent ID) 45  
application ID 50  
application ID holding lock 146  
application idle time 66  
application name 49  
application priority type 60  
application status 46  
application status change time 48  
applications connected currently 72  
applications executing in the database currently 72  
authorization ID 52  
binds/precompiles attempted 174  
buffer pool asynchronous data reads 101  
buffer pool asynchronous data writes 102  
buffer pool asynchronous index reads 104  
buffer pool asynchronous index writes 103  
buffer pool asynchronous read requests 107  
buffer pool asynchronous read time 105  
buffer pool asynchronous write time 106  
buffer pool data logical reads 93  
buffer pool data pages from extended storage 113  
buffer pool data pages to extended storage 112  
buffer pool data physical reads 94  
buffer pool data writes 95  
buffer pool index logical reads 96  
buffer pool index pages from extended storage 114  
buffer pool index pages to extended storage 112  
buffer pool index physical reads 97

data elements (by name) (*continued*)

- buffer pool index writes 98
- buffer pool information 109
- buffer pool log space cleaners triggered 107
- buffer pool threshold cleaners triggered 109
- buffer pool victim page cleaners triggered 108
- catalog cache heap full 122
- catalog cache inserts 120
- catalog cache lookups 119
- catalog cache overflows 121
- catalog node network name 42
- catalog node number 43
- client communication protocol 58
- client operating platform 58
- client process ID 57
- client product/version ID 53
- commit statements attempted 166
- committed private memory 78
- configuration NNAME of client 53
- configuration NNAME at monitoring (server)
  - node 34
- connection entries currently free 87
- connection request completion timestamp 62
- connection request start timestamp 62
- connection status 89
- connections involved in deadlock 140
- connects since first database connect 71
- coordinating node 61
- CPU time used 191
- current agents waiting on locks 144
- current number of tablequeue buffers
  - overflowed 188
- cursor name 178
- data definition language (DDL) SQL statements 169
- database activation timestamp 41
- database alias used by application 54
- database country code 59
- database deactivation timestamp 41
- database files closed 100
- database location 43
- database manager type at monitored (server)
  - node 35
- database name 39
- database path 40
- deadlocks detected 133
- degree of parallelism 191
- direct read requests 117
- direct read time 118
- direct reads from database 115
- direct write requests 117

data elements (by name) (*continued*)

- direct write time 118
- direct writes to database 116
- drda correlation token 57
- dynamic SQL statements attempted 165
- exclusive lock escalations 135
- execution elapsed time 185
- failed statement operations 165
- fcv buffers currently free 85
- host product/version ID 55
- ID of code page used by application 48
- input database alias 193
- internal automatic rebinds 170
- internal commits 171
- internal rollbacks 172
- internal rollbacks due to deadlock 173
- internal rows deleted 157
- internal rows inserted 159
- internal rows updated 158
- last backup timestamp 44
- last reset timestamp 192
- local connections 69
- local connections executing in the database
  - manager 70
- local databases with current connects 71
- lock escalations 134
- lock mode 136
- lock object name 138
- lock object type waited on 137
- lock status 137
- lock wait start timestamp 145
- lock waits 141
- locks held 131
- log being rolled forward 149
- log phase 149
- maximum database heap allocated 126
- maximum number of agents registered 74
- maximum number of agents waiting 74
- maximum number of associated agents 77
- maximum number of coordinating agents 76
- maximum number of locks held 139
- maximum secondary log space used 127
- maximum total log space used 128
- message anchors currently free 86
- minimum connection entries 87
- minimum fcv buffers free 85
- minimum message anchors 86
- minimum request blocks 88
- number of agents created 190
- number of agents working on a statement 190

data elements (by name) *(continued)*

- number of agents working on a subsection 186
- number of idle agents 75
- number of lock timeouts 139
- number of log pages read 129
- number of log pages written 130
- number of nodes 89
- number of rollforward table spaces 149
- number of rows read from tablequeues 188
- number of rows written to tablequeues 189
- number of successful fetches 182
- open local cursors 163
- open local cursors with blocking 163
- open remote cursors 160
- open remote cursors with blocking 161
- outbound application ID 55
- outbound sequence number 56
- package cache inserts 124
- package cache lookups 123
- package name 177
- pipel sorts accepted 81
- pipel sorts requested 80
- post threshold sorts 80
- previous transaction stop time 65
- previous unit of work completion timestamp 62
- process or thread id 66
- product identification 38
- product name 38
- query cost estimate 183
- query number of rows estimate 183
- rejected block cursor requests 162
- remote connections executing in the database manager 69
- remote connections to database manager 68
- request blocks currently free 88
- rollback statements attempted 167
- rolled back application 147
- rollforward timestamp 148
- rollforward type 148
- rows deleted 153
- rows inserted 153
- rows read 156
- rows selected 154
- rows updated 154
- rows written 155
- secondary connections 78
- secondary logs allocated currently 129
- section inserts 126
- section lookups 125
- section number 177

data elements (by name) *(continued)*

- select SQL statements executed 168
- sequence number 52
- sequence number holding lock 147
- server instance name 35
- server operating system 37
- server product/version ID 36
- server version 36
- service level 37
- snapshot time 193
- sort overflows 83
- SQL communications area (SQLCA) 182
- SQL dynamic statement text 180
- SQL requests since last commit 173
- start database manager timestamp 34
- statement node 174
- statement operation 176
- statement operation start timestamp 179
- statement operation stop timestamp 180
- statement sorts 181
- statement type 175
- static SQL statements attempted 164
- status of database 42
- status of DB2 instance 39
- stolen agents 77
- subsection node number 184
- subsection number 184
- subsection status 185
- table file ID 159
- table name 151
- table schema name 152
- table space name 143
- table type 150
- tablespace being rolled forward 148
- time waited for prefetch 110
- time waited on locks 142
- total buffer pool physical read time 99
- total buffer pool physical write time 100
- total fcm buffers received 90
- total fcm buffers sent 90
- total lock list memory in use 132
- total number of tablequeue buffers overflowed 187
- total sort heap allocated 79
- total sort time 82
- total sorts 82
- total time unit of work waited on locks 144
- unit of work completion status 65
- Unit of Work log space used 130
- unit of work start timestamp 63
- unit of work stop timestamp 64

data elements (by name) (*continued*)  
   update/insert/delete SQL statements executed 168  
   user authorization level> 60  
   user login ID 56  
   waited for node on a tablequeue 187  
   waiting for any node to send on a tablequeue 186  
 database  
   information 233  
   monitor, resetting 241  
 database activation timestamp, monitor element 41  
 database alias used by application, monitor element 54  
 database connection  
   applications connected currently, monitor element 72  
   applications executing in the database currently, monitor element 72  
   connection request completion timestamp, monitor element 62  
 database country code, monitor element 59  
 database deactivation timestamp, monitor element 41  
 database event monitor 18  
 database files closed, monitor element 100  
 database heap data elements  
   maximum database heap allocated 126  
 database identification data elements  
   catalog node network name 42  
   catalog node number 43  
   database activation timestamp 41  
   database deactivation timestamp 41  
   database location 43  
   database name 39  
   database path 40  
   last backup timestamp 44  
   status of database 42  
 database location, monitor element 43  
 database manager  
   monitor switches, checking 228, 230  
   statistics 232  
 database manager configuration data elements  
   active sorts 84  
   agents assigned from pool 75  
   agents registered 73  
   agents waiting for a token 73  
   applications connected currently 72  
   applications executing in the database currently 72  
   committed private memory 78  
   connection entries currently free 87  
   connection status 89  
   connects since first database connect 71  
   fcm buffers currently free 85  
   database manager configuration data elements (*continued*)  
     local connections 69  
     local connections executing in the database manager 70  
     local databases with current connects 71  
     maximum number of agents registered 74  
     maximum number of agents waiting 74  
     maximum number of associated agents 77  
     maximum number of coordinating agents 76  
     message anchors currently free 86  
     minimum connection entries 87  
     minimum fcm buffers free 85  
     minimum message anchors 86  
     minimum request blocks 88  
     number of idle agents 75  
     number of nodes 89  
     piped sorts accepted 81  
     piped sorts requested 80  
     post threshold sorts 80  
     remote connections executing in the database manager 69  
     remote connections to database manager 68  
     request blocks currently free 88  
     secondary connections 78  
     sort overflows 83  
     stolen agents 77  
     total fcm buffers received 90  
     total fcm buffers sent 90  
     total sort heap allocated 79  
     total sort time 82  
     total sorts 82  
   database manager snapshot 7  
   database manager type at monitored (server) node, monitor element 35  
   database monitor  
     description 266  
   database name, monitor element 39  
   database path, monitor element 40  
   database snapshot 7  
   database status data elements  
     See database identification data elements  
   database system monitor  
     GET DATABASE MANAGER MONITOR SWITCHES 228  
     GET MONITOR SWITCHES 230  
     GET SNAPSHOT 232  
     RESET MONITOR 241  
     UPDATE MONITOR SWITCHES 266

- db\_conn\_time element 41
- db\_heap\_top element 126
- db\_location element 43
- db\_name element 39
- db\_path element 40
- db\_status element 42
- db2 explain 2
- db2\_status 39
- db2batch 27
- db2eva 14, 16, 222
- db2evmon 16, 27, 224
- db2gov 27
- db2start\_time element 34
- ddl\_sql\_stmts element 169
- deadlock event monitor 18
- deadlocks data elements
  - See locks and deadlocks data elements
- deadlocks detected, monitor element 133
- deadlocks element 133
- degree of parallelism, monitor element 191
- degree\_parallelism element 191
- direct read requests, monitor element 117
- direct read time, monitor element 118
- direct reads from database, monitor element 115
- direct write requests, monitor element 117
- direct write time, monitor element 118
- direct writes to database, monitor element 116
- direct\_read\_reqs element 117
- direct\_read\_time element 118
- direct\_reads element 115
- direct\_write\_reqs element 117
- direct\_write\_time element 118
- direct\_writes element 116
- disconn\_time element 41
- dl\_conns element 140
- drda correlation token, monitor element 57
- DROP statement 226
- dynamic SQL statements attempted, monitor element 165
- dynamic\_sql\_stmts element 165

## E

- ESTIMATE SIZE REQUIRED FOR sqlmonss() OUTPUT BUFFER (sqlmonsz) 260
- event analyzer 27
  - See also db2eva
- event analyzer command 222
- event monitor
  - CREATE EVENT MONITOR statement 214

- event monitor (*continued*)
  - DROP statement 226
  - EVENT\_MON\_STATE function 227
  - SET EVENT MONITOR STATE statement 242
- event monitor trace formatter 224
- event monitors
  - activating 15
  - authority required 15
  - autostarting 15
  - blocked 202
  - buffering 201
  - creating 15
  - definition 4, 17
  - disk space 203, 213
  - event types 17
  - example of deadlock monitoring 10
  - file event monitors 201
  - information available 17
  - matching to application 200
  - non-blocked 202
  - output 195
  - partitioned databases 25
  - pipe event monitors 18
  - processing data 203
  - reading the trace 16
  - restarting 204
  - target 202, 204
  - trace 10, 195
  - using 15
  - when written 10, 17
- event types 17
- EVENT\_MON\_STATE function 227
- exclusive lock escalations, monitor element 135
- execution elapsed time, monitor element 185
- execution\_id element 56
- extended storage 110
- extended storage data elements
  - buffer pool data pages from extended storage 113
  - buffer pool data pages to extended storage 112
  - buffer pool index pages from extended storage 114
  - buffer pool index pages to extended storage 112

## F

- failed statement operations, monitor element 165
- failed\_sql\_stmts element 165
- fast communication manager data elements
  - connection entries currently free 87
  - connection status 89
  - fcm buffers currently free 85



fast communication manager data elements (*continued*)

- message anchors currently free 86
- minimum connection entries 87
- minimum fcm buffers free 85
- minimum message anchors 86
- minimum request blocks 88
- number of nodes 89
- request blocks currently free 88
- total fcm buffers received 90
- total fcm buffers sent 90

fcm buffers currently free, monitor element 85

fetch\_count element 182

file event monitors 201

files\_closed element 100

function

- EVENT\_MON\_STATE 227
- EVENT\_MON\_STATE, returning event monitor states 227

## G

gauges 33

GET DATABASE MANAGER MONITOR SWITCHES 228

GET MONITOR SWITCHES 230

GET SNAPSHOT 232

- effect on UPDATE MONITOR SWITCHES 266

GET SNAPSHOT (sqlmonss) 248

GET/UPDATE MONITOR SWITCHES (sqlmon) 244

## H

host product/version ID, monitor element 55

host\_prdid element 55

## I

ID of code page used by application, monitor element 48

idle agent 67

idle\_agents element 75

information available

- from snapshot monitoring 7

information data elements 33

input database alias, monitor element 193

input\_db\_alias element 193

instance connection 9

int\_auto\_rebinds element 170

int\_commits element 171

int\_deadlock\_rollbacks element 173

int\_rollbacks element 172

int\_rows\_deleted element 157

int\_rows\_inserted element 159

int\_rows\_updated element 158

interface, database system monitor

- event monitor commands 213
- event monitor GUI 16
- snapshot monitoring APIs 7, 213
- snapshot monitoring commands 7, 213
- snapshot monitoring GUI 7

internal automatic rebinds, monitor element 170

internal commits, monitor element 171

internal rollbacks due to deadlock, monitor element 173

internal rollbacks, monitor element 172

internal rows deleted, monitor element 157

internal rows inserted, monitor element 159

internal rows updated, monitor element 158

intra-query parallelism data elements

- degree of parallelism 191
- number of agents created 190
- number of agents working on a statement 190

## L

last backup timestamp, monitor element 44

last reset timestamp, monitor element 192

last\_backup element 44

last\_reset element 192

LIST ACTIVE DATABASES 235

LIST APPLICATIONS 237

LIST DCS APPLICATIONS 239

loc\_list\_in\_use 132

local connections executing in the database manager, monitor element 70

local connections, monitor element 69

local databases with current connects, monitor element 71

local\_cons element 69

local\_cons\_in\_exec 70

lock escalations, monitor element 134

lock mode, monitor element 136

lock object name, monitor element 138

lock object type waited on, monitor element 137

lock snapshot 7

lock status, monitor element 137

lock wait data elements

- agent ID holding lock 145
- application ID holding lock 146
- current agents waiting on locks 144

- lock wait data elements (*continued*)
  - lock wait start timestamp 145
  - lock waits 141
  - rolled back application 147
  - sequence number holding lock 147
  - table space name 143
  - time waited on locks 142
  - total time unit of work waited on locks 144
- lock wait start timestamp, monitor element 145
- lock waits, monitor element 141
- lock\_escals element 134
- lock\_mode element 136
- lock\_object\_name element 138
- lock\_object\_type element 137
- lock\_status element 137
- lock\_timeouts element 139
- lock\_wait\_start\_time element 145
- lock\_wait\_time element 142
- lock\_waits 141
- locks
  - current agents waiting on locks, monitor element 144
  - locks held, monitor element 131
  - total lock list memory in use, monitor element 132
  - total time unit of work waited on locks, monitor element 144
- locks and deadlocks data elements
  - connections involved in deadlock 140
  - deadlocks detected 133
  - exclusive lock escalations 135
  - lock escalations 134
  - lock mode 136
  - lock object name 138
  - lock object type waited on 137
  - lock status 137
  - locks held 131
  - maximum number of locks held 139
  - number of lock timeouts 139
  - total lock list memory in use 132
- locks held, monitor element 131
- locks\_held element 131
- locks\_held\_top element 139
- locks\_waiting 144
- log being rolled forward, monitor element 149
- log phase, monitor element 149
- log\_reads element 129
- log\_space\_used element 130
- log\_writes element 130
- logging data elements
  - maximum secondary log space used 127

- logging data elements (*continued*)
  - maximum total log space used 128
  - number of log pages read 129
  - number of log pages written 130
  - secondary logs allocated currently 129
  - Unit of Work log space used 130
- logical view 22

## M

- MA\_free element 86
- MA\_free\_bottom element 86
- maximum database heap allocated, monitor element 126
- maximum number of agents registered, monitor element 74
- maximum number of agents waiting, monitor element 74
- maximum number of associated agents, monitor element 77
- maximum number of coordinating agents, monitor element 76
- maximum number of locks held, monitor element 139
- maximum secondary log space used, monitor element 127
- maximum total log space used, monitor element 128
- memory requirements 23
- message anchors currently free, monitor element 86
- minimum connection entries, monitor element 87
- minimum fcm buffers free, monitor element 85
- minimum message anchors, monitor element 86
- minimum request blocks, monitor element 88
- mon\_heap\_sz 23
- monitor switches
  - query database manager switch settings 4
  - setting explicitly 3
  - setting for a snapshot 5
  - setting implicitly 4
- monitored level 28
- monitored object 28
- monitoring
  - levels 3
- monitoring application 28
- monitoring databases 228, 230
- multiple partition databases
  - event monitors 25
  - snapshot monitoring 23
  - subsections 26
  - tablequeue 26

## N

- nodegroup 26
- non-buffer I/O activity data elements
  - direct read requests 117
  - direct read time 118
  - direct reads from database 115
  - direct write requests 117
  - direct write time 118
  - direct writes to database 116
- num\_agents element 190
- num\_subagents element 186
- number of agents created, monitor element 190
- number of agents working on a statement, monitor element 190
- number of agents working on a subsection, monitor element 186
- number of idle agents, monitor element 75
- number of lock timeouts, monitor element 139
- number of log pages read, monitor element 129
- number of log pages written, monitor element 130
- number of nodes, monitor element 89
- number of rollforward table spaces, monitor element 149
- number of rows read from tablequeues, monitor element 188
- number of rows written to tablequeues, monitor element 189
- number of successful fetches, monitor element 182
- number\_nodes element 89

## O

- open local cursors with blocking, monitor element 163
- open local cursors, monitor element 163
- open remote cursors with blocking, monitor element 161
- open remote cursors, monitor element 160
- open\_loc\_curs element 163
- open\_loc\_curs\_blk element 163
- open\_rem\_curs element 160
- open\_rem\_curs\_blk element 161
- operation element 176
- outbound application ID, monitor element 55
- outbound sequence number, monitor element 56
- outbound\_appl\_id element 55
- outbound\_sequence\_no element 56
- overflow\_accesses 157
- overflows, event monitor 202

## P

- package cache
  - See cache
- package cache data elements
  - package cache inserts 124
  - package cache lookups 123
  - section inserts 126
  - section lookups 125
- package cache inserts, monitor element 124
- package cache lookups, monitor element 123
- package name, monitor element 177
- package\_name element 177
- pipe event monitors
  - defining 18
  - overflows 20
  - using 18
- piped sorts accepted, monitor element 81
- piped sorts requested, monitor element 80
- piped\_sorts\_accepted element 81
- piped\_sorts\_requested element 80
- pkg\_cache\_inserts element 124
- pkg\_cache\_lookups 123
- pool\_async\_data\_read\_reqs element 107
- pool\_async\_data\_reads element 101
- pool\_async\_data\_writes element 102
- pool\_async\_index\_reads element 104
- pool\_async\_index\_writes element 103
- pool\_async\_read\_time element 105
- pool\_async\_write\_time element 106
- pool\_data\_from\_estore element 113
- pool\_data\_l\_reads element 93
- pool\_data\_p\_reads element 94
- pool\_data\_to\_estore element 112
- pool\_data\_writes element 95
- pool\_drty\_pg\_steal\_clns 108
- pool\_drty\_pg\_thrsh\_clns element 109
- pool\_index\_from\_estore element 114
- pool\_index\_l\_reads element 96
- pool\_index\_p\_reads element 97
- pool\_index\_to\_estore element 112
- pool\_index\_writes element 98
- pool\_lsn\_gap\_clns element 107
- pool\_read\_time element 99
- pool\_write\_time element 100
- post threshold sorts, monitor element 80
- post\_threshold\_sorts element 80
- prefetch\_wait\_time element 110
- prefetchers 92
- prev\_stop\_time element 65

prev\_uow\_stop\_time element 62  
previous transaction stop time, monitor element 65  
previous unit of work completion timestamp, monitor element 62  
process or thread id, monitor element 66  
product identification, monitor element 38  
product name, monitor element 38  
product\_name element 38

## Q

query  
    database manager monitor switch settings 4  
    event monitor state 17  
query cost estimate, monitor element 183  
query number of rows estimate, monitor element 183  
query\_card\_estimate element 183  
query\_cost\_estimate element 183

## R

RB\_free element 88  
RB\_free\_bottom element 88  
rej\_curs\_blk element 162  
rejected block cursor requests, monitor element 162  
rem\_cons\_in element 68  
rem\_cons\_in\_exec 69  
remote connections executing in the database manager, monitor element 69  
remote connections to database manager, monitor element 68  
request blocks currently free, monitor element 88  
RESET MONITOR 241  
RESET MONITOR (sqlmrset) 263  
resetting monitor data 21  
rf\_log\_num element 149  
rf\_num\_tspaces element 149  
rf\_status element 149  
rf\_timestamp element 148  
rf\_type element 148  
rollback statements attempted, monitor element 167  
rollback\_sql\_stmts 167  
rolled back application, monitor element 147  
rolled\_back\_appl\_id element 147  
rollforward data elements  
    log being rolled forward 149  
    log phase 149  
    number of rollforward table spaces 149  
    rollforward timestamp 148  
    rollforward type 148

rollforward data elements (*continued*)  
    tablespace being rolled forward 148  
rollforward timestamp, monitor element 148  
rollforward type, monitor element 148  
rows deleted, monitor element 153  
rows inserted, monitor element 153  
rows read, monitor element 156  
rows selected, monitor element 154  
rows updated, monitor element 154  
rows written, monitor element 155  
rows\_deleted element 153  
rows\_inserted element 153  
rows\_read element 156  
rows\_selected element 154  
rows\_updated element 154  
rows\_written element 155

## S

samples  
    event monitor trace 12, 196  
    event monitoring on partitioned databases 25  
    lock snapshot 5  
    monitoring deadlocks with a lock snapshot 10  
    programming to read the data stream 204  
    query event monitor state 17  
    setting switches 21  
    snapshots on partitioned databases 23  
sec\_log\_used\_top element 127  
sec\_logs\_allocated element 129  
secondary connections, monitor element 78  
secondary logs allocated currently, monitor element 129  
section inserts, monitor element 126  
section lookups, monitor element 125  
section number, monitor element 177  
section\_number element 177  
select SQL statements executed, monitor element 168  
select\_sql\_stmts element 168  
sequence number holding lock, monitor element 147  
sequence number, monitor element 52  
sequence\_no element 52  
sequence\_no\_holding\_lk element 147  
server identification data elements  
    configuration NNAME at monitoring (server)  
        node 34  
    database manager type at monitored (server)  
        node 35  
    product identification 38  
    product name 38

- server identification data elements (*continued*)
  - server instance name 35
  - server operating system 37
  - server product/version ID 36
  - server version 36
  - service level 37
  - start database manager timestamp 34
  - status of DB2 instance 39
- server instance name, monitor element 35
- server operating system, monitor element 37
- server product/version ID, monitor element 36
- server status data elements
  - See server identification data elements
- server version, monitor element 36
- server\_db2\_type element 35
- server\_instance\_name element 35
- server\_nname element 34
- server\_platform element 37
- server\_prdid element 36
- server\_version element 36
- service level, monitor element 37
- service\_level element 37
- SET EVENT MONITOR STATE statement 242
- setting switches
  - for a monitoring application 5
- snapshot monitoring
  - APIs 213
  - authority required 6
  - availability of data 10
  - commands 213
  - data element categories 4
  - definition 4
  - information available 7
  - information returned 7
  - interface 7
  - partitioned databases 23
  - request types 7
  - required connection 9
  - sample output 5, 21
  - setting switches 5, 21
  - snapshot types 7
- snapshot monitoring data elements
  - input database alias 193
  - last reset timestamp 192
  - snapshot time 193
- snapshot time, monitor element 193
- snapshot types
  - application 8
  - buffer pool 8
  - database 8
- snapshot types (*continued*)
  - database manager 8
  - lock 8
  - table 8
  - table space 8
- sort overflows, monitor element 83
- sort\_heap\_allocated element 79
- sort\_overflows element 83
- SQL communications area (SQLCA), monitor element 182
- SQL cursors data elements
  - accepted block cursor requests 162
  - open local cursors 163
  - open local cursors with blocking 163
  - open remote cursors 160
  - open remote cursors with blocking 161
  - rejected block cursor requests 162
- SQL dynamic statement text, monitor element 180
- SQL requests since last commit, monitor element 173
- SQL statement
  - CREATE EVENT MONITOR 214, 221
  - DROP 226
  - SET EVENT MONITOR STATE 242, 243
- SQL statement activity data elements
  - binds/precompiles attempted 174
  - commit statements attempted 166
  - data definition language (DDL) SQL statements 169
  - dynamic SQL statements attempted 165
  - failed statement operations 165
  - internal automatic rebinds 170
  - internal commits 171
  - internal rollbacks 172
  - internal rollbacks due to deadlock 173
  - rollback statements attempted 167
  - select SQL statements executed 168
  - SQL requests since last commit 173
  - statement node 174
  - static SQL statements attempted 164
  - update/insert/delete SQL statements executed 168
- SQL statement details data elements
  - application creator 179
  - cursor name 178
  - number of successful fetches 182
  - package name 177
  - query cost estimate 183
  - query number of rows estimate 183
  - section number 177
  - SQL communications area (SQLCA) 182
  - SQL dynamic statement text 180
  - statement operation 176

- SQL statement details data elements (*continued*)
  - statement operation start timestamp 179
  - statement operation stop timestamp 180
  - statement sorts 181
  - statement type 175
- sql\_reqs\_since\_commit element 173
- sqlca element 182
- sqlcode -973 23
- sqlmon.h header file 195
- ss\_exec\_time element 185
- ss\_node\_number element 184
- ss\_number element 184
- ss\_status element 185
- start database manager timestamp, monitor element 34
- statement event monitor 18
- statement node, monitor element 174
- statement operation start timestamp, monitor element 179
- statement operation stop timestamp, monitor element 180
- statement operation, monitor element 176
- statement sorts, monitor element 181
- statement type, monitor element 175
- static SQL statements attempted, monitor element 164
- static\_sql\_stmts 164
- statistics
  - database manager 232
- status element 65
- status of database, monitor element 42
- status of DB2 instance, monitor element 39
- status\_change\_time element 48
- stmt\_node\_number element 174
- stmt\_operations element 176
- stmt\_sorts element 181
- stmt\_start element 179
- stmt\_stop element 180
- stmt\_text element 180
- stmt\_type element 175
- stolen agent 67
- stolen agents, monitor element 77
- stop\_time element 180
- subagent 67
- subagent information data elements
  - active sorts 84
  - pipel sorts accepted 81
  - pipel sorts requested 80
  - post threshold sorts 80
  - process or thread id 66
  - sort overflows 83
  - total sort heap allocated 79

- subagent information data elements (*continued*)
  - total sort time 82
  - total sorts 82
- subsection details data elements
  - current number of tablequeue buffers
    - overflowed 188
  - execution elapsed time 185
  - number of agents working on a subsection 186
  - number of rows read from tablequeues 188
  - number of rows written to tablequeues 189
  - subsection node number 184
  - subsection number 184
  - subsection status 185
  - total number of tablequeue buffers overflowed 187
  - waited for node on a tablequeue 187
  - waiting for any node to send on a tablequeue 186
- subsection node number, monitor element 184
- subsection number, monitor element 184
- subsection status, monitor element 185
- subsections 183
  - definition 26
  - monitoring 26
  - tablequeue 26
- switches
  - See monitor switches

## T

- table activity data elements
  - accesses to overflowed records 157
  - internal rows deleted 157
  - internal rows inserted 159
  - internal rows updated 158
  - rows deleted 153
  - rows inserted 153
  - rows read 156
  - rows selected 154
  - rows updated 154
  - rows written 155
  - table file ID 159
  - table name 151
  - table schema name 152
  - table type 150
- table event monitor 18
- table file ID, monitor element 159
- table name, monitor element 151
- table schema name, monitor element 152
- table snapshot 7
- table space event monitor 18

- table space name, monitor element 143
- table space snapshot 7
- table type, monitor element 150
- table\_file\_id element 159
- table\_name element 151
- table\_schema element 152
- table\_type element 150
- tablequeue 26
- tablespace being rolled forward, monitor element 148
- tablespace\_name element 143
- taking a snapshot
  - issuing get snapshot command 5
  - sample output 5
- time 33
- time waited for prefetch, monitor element 110
- time waited on locks, monitor element 142
- time\_stamp element 193
- timestamp 33
- tools
  - control center 27
  - db2batch 27
  - db2evmon 27
  - db2gov 27
  - event analyzer 27
- tot\_log\_used\_top element 128
- total buffer pool physical read time, monitor element 99
- total buffer pool physical write time, monitor element 100
- total fcm buffers received, monitor element 90
- total fcm buffers sent, monitor element 90
- total lock list memory in use, monitor element 132
- total number of tablequeue buffers overflowed, monitor element 187
- total sort heap allocated , monitor element 79
- total sort time, monitor element 82
- total sorts, monitor element 82
- total time unit of work waited on locks, monitor element 144
- total\_buffers\_rcvd element 90
- total\_buffers\_sent element 90
- total\_cons element 71
- total\_sec\_cons element 78
- total\_sort\_time element 82
- total\_sorts element 82
- tq\_cur\_send\_spills element 188
- tq\_node\_waited\_for element 187
- tq\_rows\_read element 188
- tq\_rows\_written element 189
- tq\_tot\_send\_spills element 187

- tq\_wait\_for\_any element 186
- trace
  - event monitor 10
  - format 195
  - programming to read 204
  - sample 12
  - size 202
  - viewing 14
- transaction event monitor 18
- ts\_name element 148

## U

- uid\_sql\_stmts element 168
- unit of work completion status, monitor element 65
- Unit of Work log space used, monitor element 130
- unit of work start timestamp, monitor element 63
- unit of work stop timestamp, monitor element 64
- uow\_comp\_status element 65
- uow\_lock\_wait\_time element 144
- uow\_log\_space\_used element 130
- uow\_start\_time element 63
- uow\_stop\_time element 64
- UPDATE MONITOR SWITCHES 266
- update/insert/delete SQL statements executed, monitor element 168
- updating switch settings
  - See* setting switches
- user authorization level, monitor element 60
- user login ID, monitor element 56

## V

- view, logical 22

## W

- waited for node on a tablequeue, monitor element 187
- waiting for any node to send on a tablequeue, monitor element 186
- water mark 33

## X

- x\_lock\_escals element 135





---

## Contacting IBM

This section lists ways you can get more information from IBM.

If you have a technical problem, please take the time to review and carry out the actions suggested by the *Troubleshooting Guide* before contacting DB2 Customer Support. Depending on the nature of your problem or concern, this guide will suggest information you can gather to help us to serve you better.

For information or to order any of the DB2 Universal Database products contact an IBM representative at a local branch office or contact any authorized IBM software remarketer.

### Telephone

If you live in the U.S.A., call one of the following numbers:

- 1-800-237-5511 to learn about available service options.
- 1-800-IBM-CALL (1-800-426-2255) or 1-800-3IBM-OS2 (1-800-342-6672) to order products or get general information.
- 1-800-879-2755 to order publications.

For information on how to contact IBM outside of the United States, see Appendix A of the IBM Software Support Handbook. You can access this document by selecting the "Roadmap to IBM Support" item at: <http://www.ibm.com/support/>.

Note that in some countries, IBM-authorized dealers should contact their dealer support structure instead of the IBM Support Center.

### World Wide Web

<http://www.software.ibm.com/data/>  
<http://www.software.ibm.com/data/db2/library/>

The DB2 World Wide Web pages provide current DB2 information about news, product descriptions, education schedules, and more. The DB2 Product and Service Technical Library provides access to frequently asked questions, fixes, books, and up-to-date DB2 technical information. (Note that this information may be in English only.)

### Anonymous FTP Sites

<ftp.software.ibm.com>

Log on as anonymous. In the directory `/ps/products/db2`, you can find demos, fixes, information, and tools concerning DB2 and many related products.

### Internet Newsgroups

`comp.databases.ibm-db2`, `bit.listserv.db2-l`

These newsgroups are available for users to discuss their experiences with DB2 products.

### CompuServe

**GO IBMDB2** to access the IBM DB2 Family forums

All DB2 products are supported through these forums.

To find out about the IBM Professional Certification Program for DB2 Universal Database, go to <a href="http://www.software.ibm.com/data/db2/db2tech/db2cert.html">http://www.software.ibm.com/data/db2/db2tech/db2cert.html</a>
--



Part Number: 10J8164



Printed in the United States of America  
on recycled paper containing 10%  
recovered post-consumer fiber.

S10J-8164-00



10J8164

