

Note

Before using this information and the product it supports, read the general information under Appendix E, "Notices," on page 397.

Edition Notice

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.

You can order IBM publications online or through your local IBM representative.

- To order publications online, go to the IBM Publications Center at www.ibm.com/shop/publications/order
- To find your local IBM representative, go to the IBM Directory of Worldwide Contacts at www.ibm.com/planetwide

To order DB2 publications from DB2 Marketing and Sales in the United States or Canada, call 1-800-IBM-4YOU (426-4968).

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 2007, 2009.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About this book vii

Chapter 1. Introduction to DB2 workload manager concepts 1

Stages of workload management 1
Workload manager administrator authority (WLMADM) 3
Frequently asked questions about DB2 workload manager 4

Chapter 2. Work identification 13

Activities 13
DDL statements for DB2 workload manager . . . 16
Work identification by origin with workloads . . . 17
 Workload assignment 21
 Default workloads 23
 Creating a workload 28
 Altering a workload 29
 Permitting occurrences of a workload to access the database 31
 Preventing occurrences of a workload from accessing the database 31
 Enabling a workload 32
 Disabling a workload 32
 Granting the USAGE privilege on a workload . . 33
 Revoking the USAGE privilege on a workload . 34
 Dropping a workload 34
 Example: Workload assignment 35
 Example: Workload assignment when workload attributes have single values 39
 Example: Workload assignment for a unit of work when multiple workloads exist 41
 Example: Workload assignment when workload attributes have multiple values 44
Work identification by type of work with work classes 45
 Work classes and work class sets 48
 Evaluation order of work classes in a work class set 50
 Assignment of activities to work classes . . . 51
 Work classifications supported by thresholds . . 51
 Creating a work class 52
 Altering a work class 55
 Dropping a work class 55
 Creating a work class set 55
 Altering a work class set 56
 Dropping a work class set 56
 Example: Analyzing workloads by activity type . 57
 Example: Using a work class set to manage specific types of activities 58
 Example: Working with a work class defined with the ALL keyword 59

Chapter 3. Activities management 63

Resource assignment with service classes 63

Default service superclasses and subclasses . . . 66
Activity-to-service class mapping 68
Agent priority of service classes 72
Prefetch priority of service classes 72
Buffer pool priority of service classes 73
States of connections and activities in a service class 74
System-level entities not tracked by service classes 75
Creating a service class 76
Altering a service class 78
Dropping a service class 80
Example: Using service classes 81
Example: Analyzing a service class–related system slowdown 86
Example: Investigating agent usage by service class 88
Control of work with thresholds 88
 Threshold domain and enforcement scope . . . 91
 Threshold evaluation order 92
 Connection thresholds 94
 Activity thresholds 95
 Aggregate thresholds 102
 Creating a threshold 109
 Altering a threshold 110
 Dropping a threshold 110
 Example: Using thresholds 111
Priority aging of ongoing work 113
 Sample priority aging scripts 117
 Remapping activities between service subclasses 123
Apply controls to types of activities with work action sets 125
 How work classes, work class sets, work actions, and work action sets work together and are associated with other DB2 objects 126
 Work actions and work action sets 128
 Work actions and the work action set domain 130
 Thresholds that can be used in work actions . . 133
 Application of work actions to database activities 133
 Workload and work action set comparison . . 135
 Creating a work action set 137
 Altering a work action set 138
 Disabling a work action set 139
 Dropping a work action set 140
 Creating a work action 140
 Altering a work action 143
 Disabling a work action 144
 Dropping a work action 145
 Example: Using a work action set and database threshold 145
 Example: Using work action sets to determine the types of work being run 147

Chapter 4. Monitoring and intervention 149

Real-time monitoring with table functions	149
Example: Using DB2 workload manager table functions	154
Example: Monitoring current system behavior at different levels using DB2 workload manager table functions	155
Example: Obtaining point-in-time statistics from service classes	158
Example: Aggregating data using DB2 workload manager table functions	159
Historical monitoring with WLM event monitors	160
Available monitoring data	166
DB2 workload manager stored procedures	168
Statistics for DB2 workload manager objects	169
Statistics collection and monitoring with priority aging	178
Histograms in workload management	180
Historical analysis tool	188
Collecting workload management statistics using a statistics event monitor	188
Resetting statistics on DB2 workload manager objects	190
Monitoring metrics for DB2 workload manager	191
Workload management table functions and snapshot monitor integration	193
Monitoring threshold violations	194
Collecting data for individual activities	195
Importing activity information into the Design Advisor	197
Cancelling activities	197
Guidelines for capturing information about and investigating a rogue activity	198
Workload management performance modelling	199
Example: Capturing information about an activity for later analysis	199

Chapter 5. Integration with operating system workload managers 203

Integration of AIX Workload Manager with DB2 workload manager	203
Integration of Linux workload management with DB2 workload manager	208

Chapter 6. Tutorial for DB2 workload manager 215

Exercise 1: Getting started with basic monitoring using default DB2 workload manager objects	215
Exercise 2: Isolating activities using service classes and workloads	220
Exercise 3: Using thresholds to control rogue activities and using the threshold violation monitor	227
Exercise 4: Differentiating activities by activity type	230
Exercise 5: Using histograms for service classes	237
Exercise 6: Investigating delays with WLM table functions	246
Exercise 7: Cancelling an ongoing activity	249
Exercise 8: Discovering what types of activities are running on your system	250

Exercise 9: Capturing detailed information about an executing activity	253
Exercise 10: Generating historical data and reports	255
Exercise 11: Using extended aggregates for service classes	258

Chapter 7. Workload management scenarios 265

Workload management sample application	265
Scenario: Investigating a workload-related system slowdown	266
Scenario: Identifying activities that are taking too long to complete	267
Scenario: Tuning a DB2 workload manager configuration when capacity planning data is available	269
Scenario: Tuning a DB2 workload manager configuration when capacity planning information is unavailable	271
Scenario: Identifying activities with low estimated cost and high runtime	276

Chapter 8. Reference 277

Procedures and table functions	277
WLM_CANCEL_ACTIVITY - Cancel an activity	277
WLM_CAPTURE_ACTIVITY_IN_PROGRESS - Collect activity information for activities event monitor	278
WLM_COLLECT_STATS - Collect and reset workload management statistics	280
WLM_GET_ACTIVITY_DETAILS - Return detailed information about a specific activity	281
WLM_GET_QUEUE_STATS table function - Return threshold queue statistics	287
WLM_GET_SERVICE_CLASS_AGENTS_V97 table function - List agents running in a service class	291
WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97 - List workload occurrences	298
WLM_GET_SERVICE_SUBCLASS_STATS_V97 table function - Return statistics of service subclasses	302
WLM_GET_SERVICE_SUPERCLASS_STATS - Return statistics of service superclasses	309
WLM_GET_WORK_ACTION_SET_STATS - Return work action set statistics	311
WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97 - Return a list of activities	313
WLM_GET_WORKLOAD_STATS_V97 table function - Return workload statistics	318
WLM_SET_CLIENT_INFO procedure - Set client information	321
Workload management monitor elements	323
activate_timestamp - Activate timestamp monitor element	323
activity_collected - Activity collected monitor element	324
activity_id - Activity ID monitor element	324

activity_secondary_id - Activity secondary ID monitor element	325	coord_act_est_cost_avg - Coordinator activity estimated cost average monitor element	340
activity_type - Activity type monitor element	325	coord_act_interarrival_time_avg - Coordinator activity arrival time average monitor element	340
act_cpu_time_top - Activity CPU time top monitor element	326	db_work_action_set_id - Database work action set ID monitor element	341
act_exec_time - Activity execution time monitor element	326	db_work_class_id - Database work class ID monitor element	342
act_remapped_in - Activities remapped in monitor element	327	destination_service_class_id - Destination service class ID monitor element	342
act_remapped_out - Activities remapped out monitor element	327	histogram_type - Histogram type monitor element	342
act_rows_read_top - Activity rows read top monitor element	327	last_wlm_reset - Time of last reset monitor element	343
act_total - Activities total monitor element	328	num_threshold_violations - Number of threshold violations monitor element	344
agg_temp_tablespace_top - Aggregate temporary table space top monitor element	328	num_remaps - Number of remaps monitor element	344
arm_correlator - Application response measurement correlator monitor element	328	number_in_bin - Number in bin monitor element	345
bin_id - Histogram bin identifier monitor element	329	parent_activity_id - Parent activity ID monitor element	345
bottom - Histogram bin bottom monitor element	329	parent_uow_id - Parent unit of work ID monitor element	345
concurrent_act_top - Concurrent activity top monitor element	329	prep_time - Preparation time monitor element	346
concurrent_connection_top - Concurrent connection top monitor element	330	queue_assignments_total - Queue assignments total monitor element	346
concurrent_wlo_act_top - Concurrent WLO activity top monitor element	330	queue_size_top - Queue size top monitor element	346
concurrent_wlo_top - Concurrent workload occurrences top monitor element	330	queue_time_total - Queue time total monitor element	347
coord_act_aborted_total - Coordinator activities aborted total monitor element	331	request_exec_time_avg - Request execution time average monitor element	347
coord_act_completed_total - Coordinator activities completed total monitor element	331	rows_fetched - Rows fetched monitor element	348
coord_act_est_cost_avg - Coordinator activity estimated cost average monitor element	331	rows_modified - Rows modified monitor element	348
coord_act_exec_time_avg - Coordinator activities execution time average monitor element	332	rows_returned - Rows returned monitor element	349
coord_act_interarrival_time_avg - Coordinator activity arrival time average monitor element	333	rows_returned_top - Actual rows returned top monitor element	350
coord_act_lifetime_avg - Coordinator activity lifetime average monitor element	334	sc_work_action_set_id - Service class work action set ID monitor element	351
coord_act_lifetime_top - Coordinator activity lifetime top monitor element	334	sc_work_class_id - Service class work class ID monitor element	351
coord_act_queue_time_avg - Coordinator activity queue time average monitor element	335	section_env - Section environment monitor element	352
coord_act_rejected_total - Coordinator activities rejected total monitor element	336	service_class_id - Service class ID monitor element	352
coord_partition_num - Coordinator partition number monitor element	336	service_subclass_name - Service subclass name monitor element	353
cost_estimate_top - Cost estimate top monitor element	336	service_superclass_name - Service superclass name monitor element	354
coord_act_lifetime_avg - Coordinator activity lifetime average monitor element	337	source_service_class_id - Source service class ID monitor element	354
coord_act_queue_time_avg - Coordinator activity queue time average monitor element	338	statistics_timestamp - Statistics timestamp monitor element	355
coord_act_exec_time_avg - Coordinator activities execution time average monitor element	338	temp_tablespace_top - Temporary table space top monitor element	355
request_exec_time_avg - Request execution time average monitor element	339	threshold_action - Threshold action monitor element	356
		threshold_domain - Threshold domain monitor element	356

threshold_maxvalue - Threshold maximum value monitor element	356
threshold_name - Threshold name monitor element	357
threshold_predicate - Threshold predicate monitor element	357
threshold_queue_size - Threshold queue size monitor element	357
thresholdid - Threshold ID monitor element	358
time_completed - Time completed monitor element	358
time_created - Time created monitor element	358
time_of_violation - Time of violation monitor element	359
time_started - Time started monitor element	359
top - Histogram bin top monitor element	359
uow_id - Unit of work ID monitor element	359
uow_total_time_top - UOW total time top monitor element	360
wlo_completed_total - Workload occurrences completed total monitor element	360
work_action_set_id - Work action set ID monitor element	361
work_action_set_name - Work action set name monitor element	361
work_class_id - Work class ID monitor element	361
work_class_name - Work class name monitor element	362
workload_id - Workload ID monitor element	362
workload_name - Workload name monitor element	363
workload_occurrence_id - Workload occurrence identifier monitor element	363
Commands	364
SET WORKLOAD command	364
Configuration parameters	365
wlm_collect_int - Workload management collection interval configuration parameter	365
Catalog views	366
SYSCAT.HISTOGRAMTEMPLATEBINS	366
SYSCAT.HISTOGRAMTEMPLATES	366

SYSCAT.HISTOGRAMTEMPLATEUSE	366
SYSCAT.SERVICECLASSES	367
SYSCAT.THRESHOLDS	369
SYSCAT.WORKACTIONS	371
SYSCAT.WORKACTIONSETS	373
SYSCAT.WORKCLASSES	373
SYSCAT.WORKCLASSESETS	374
SYSCAT.WORKLOADAUTH	375
SYSCAT.WORKLOADCONNATTR	375
SYSCAT.WORKLOADS	376

Appendix A. Naming rules 379

Appendix B. Roles 381

Appendix C. Trusted contexts and trusted connections 383

Appendix D. Overview of the DB2 technical information 387

DB2 technical library in hardcopy or PDF format	387
Ordering printed DB2 books	390
Displaying SQL state help from the command line processor	391
Accessing different versions of the DB2 Information Center	391
Displaying topics in your preferred language in the DB2 Information Center	391
Updating the DB2 Information Center installed on your computer or intranet server	392
Manually updating the DB2 Information Center installed on your computer or intranet server	393
DB2 tutorials	395
DB2 troubleshooting information	395
Terms and Conditions	396

Appendix E. Notices 397

Index 401

About this book

This book provides information on the DB2® workload manager features and functionality that can help you obtain a stable, predictable execution environment that meets your business objectives. Using DB2 workload manager, both requests and resources are managed. This book also provides information on monitoring and performing troubleshooting for the workload on your data server.

Chapter 1. Introduction to DB2 workload manager concepts

A good workload management system helps to efficiently meet goals in the environment where work occurs. You can see examples of the need for a good workload management system all around you.

For example, look at a grocery store. Different activities must be considered: serving customers, stocking shelves, maintaining inventories, and so on. And some simple goals must be set. Store owners want to maximize both the number of customers who move through the store, and the amount that customers purchase, achieving both goals in a way that customers leave both satisfied and wanting to come back. Store owners must also ensure that they have sufficient stock for their customers to buy (but not too much stock, because waste becomes an issue). Store owners also track what their customers purchase, and use this information to create advertisements that are designed to induce their customers to return. Monitoring mechanisms track inventory and send notifications when stocks run low. Security devices are in place to detect shoplifting. Special fast checkout lanes are created so that shoppers who only want to purchase a few items can do so without having to wait behind other customers who are purchasing many items. If all of these goals are met and all of these operational procedures work well, customers are satisfied, and are likely to return rather than to go to another store. These goals and operational procedures are all aspects of workload management.

In a data server environment, you can see even more of a need for effective management of work, especially now that data servers are being stressed like never before. Cash registers generate thousands of data inserts, reports are constantly being generated to determine whether sales targets are being met, batch applications run to load collected data, and administration tasks such as backups and reorganizations run to protect the data and make the server run optimally. All these operations are using the same database system and competing for the same resources.

To ensure the best chance of meeting goals for running a data server, an efficient workload management system is critical.

Stages of workload management

Workload management has three clearly defined stages: identification of the work entering the data server, management of the work when it is running, and monitoring to ensure that the data server is being used efficiently.

A number of aspects must be considered for successful workload management with DB2 workload manager, starting with understanding your goals. In the grocery store example described in Chapter 1, "Introduction to DB2 workload manager concepts," goals might include maximizing customer spending, minimizing shoplifting, and ensuring that customers leave the store satisfied so that they will return again.

In a data server environment, you must also define goals. Sometimes the goals are clear, especially when they originate from service level agreement (SLA) objectives. For example, queries from a particular application can consume no more than 10% of the total processor resource. Goals can also be tied to a particular time of day. For example, an overnight batch utility might have to complete loading data by 8

a.m. so that the daily sales reports are on time. In other situations, the goals can be difficult to quantify. A goal might be to keep the database users satisfied and to prevent aberrant database activity from hampering their day-to-day work. Whether the goals are quantifiable or not, understanding them is critical when considering the following stages of workload management:

Identification

If you want to achieve a goal for some kind of work, you first must be able to identify details about the work. In the grocery store, you can identify shopper information through credit cards and debit cards, or an unpaid-for item through an active security tag on the item. For the data server, you need to decide how you want to identify the work that enters the system. You can use the name of the application that submits the work, the authorization ID that submits the work, or a combination of elements that provide some form of identification.

Management

The management phase includes mechanisms for making steady progress towards your goal, and actions to take if a goal is not being met. An example of a mechanism is managing price checks in fast checkout lanes. Fast checkout lanes should result in faster throughput and satisfied customers, but if a carton of milk has the wrong price and a price check is required, the fast checkout lane could slow down. The problem is managed by performing a fast price check, possibly opening up another checkout lane, and trying to fix the pricing problem so that it does not occur again. On the data server, you might find that overall performance is suffering when a few poorly written SQL statements are running, a surge in volume occurs during peak times, or there is too much competition between different applications for the same resources. The management phase includes mechanisms for assigning resources to achieve your goals, and actions to take if a goal is not being met.

Monitoring

Monitoring is important for a couple of reasons. First, to determine whether you are achieving a goal, you must have a mechanism to track progress toward that goal. Also, monitoring helps to identify the problems that might be preventing you from achieving your goal. In a store, the store manager can watch the flow of customers, automatically be alerted to problems such as shoplifting or dangerously low inventory of a particular sale item, or perform analysis on historical purchase patterns to determine optimal product placement in the store. For a data server, there are often explicit goals for response times of database activities and it is important to have a way to measure this metric, and watch for trends.

The following figure represents the workload management stages:

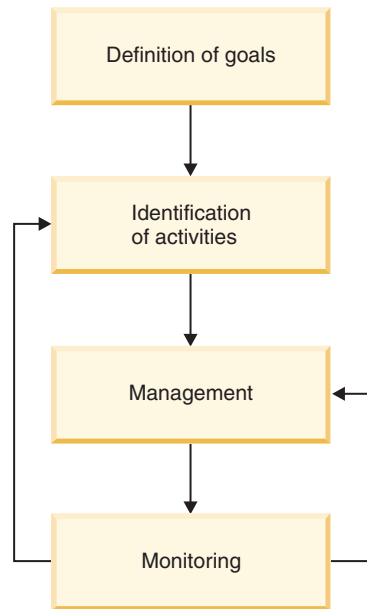


Figure 1. Stages of workload management

Workload manager administrator authority (WLMADM)

You need workload manager administrator authority (WLMADM) authority to manage workload objects for a specific database. This authority enables you to create, alter, drop, comment on, grant access to, and revoke access from DB2 workload manager objects.

The security administrator, who is someone holding SECADM authority, or a user with ACCESSCTRL authority can grant WLMADM authority to a user, group, or role.

WLMADM authority provides the ability to perform the following tasks:

- Issue CREATE, ALTER, COMMENT ON, and DROP statements for the following DB2 workload manager objects:
 - Histogram templates
 - Service classes
 - Thresholds
 - Work action sets
 - Work class sets
 - Workloads
- Issue GRANT and REVOKE statements for workload privileges

This authority is a subset of the database administrator (DBADM) authority.

Ownership of DB2 workload manager objects

Most database objects have owners, and these owners have the authority to alter the objects that they own. Unlike most objects, DB2 workload manager objects do not have owners, to avoid unpredictable effects on your execution environment because of changes to DB2 workload manager objects by those owners. If a

resource allocation setting is changed for a service class, for example, this change affects not only the service class itself but also the resources available to other service classes and hence must remain under control of DB2 workload manager.

The following example illustrates the problem that would be caused by having owners for DB2 workload manager objects. Assume that a service superclass has two user-defined service subclasses, A and B, and that each service subclass has a different owner. Initially, the prefetch priority setting is medium for the default service subclass and for the two service subclasses. If the owner of service subclass A changes its prefetch priority to high and many prefetch requests come from this service subclass, connections to service subclass B and the default subclass have less access to prefetcher services, and the performance of activities running in these service subclasses might suffer.

Frequently asked questions about DB2 workload manager

This section provides you with answers to common questions about DB2 workload manager (WLM) and about how it relates to existing Query Patroller and Governor functionality. Answers to the following questions are available:

- On which DB2 platforms can I use DB2 workload manager?
- Do I need Query Patroller to use DB2 workload manager?
- Why are the new WLM capabilities not integrated into Query Patroller?
- How does this new functionality affect Query Patroller and DB2 Governor?
- Why should I not use just concurrency thresholds to control all my work ? Is that not how Query Patroller works?
- I am not on AIX®, does this mean I do not have any control over processor resource or I/O activity?
- Can I use AIX WLM to manage I/O activity?
- Can I use AIX WLM to manage memory use?
- Why is there not a tool to automatically upgrade from my Query Patroller configuration?
- Is there a way for WebSphere® Application Server to pass the client information fields used by the DB2 workload?
- Why do the service class agent priority settings not seem to be in effect?
- Can I create multiple CONCURRENTDBCOORDACTIVITIES concurrency thresholds for the same set of work?
- Why is my work not assigned to the correct workload?
- Why does the DB2 data server not automatically create AIX service classes?
- Why does DB2 workload manager affect REORGCHK, IMPORT, EXPORT and other CLP commands?
- Is there a way to change the service class to which an activity is assigned while it is executing?
- Much of my batch work is done using CLP scripts under the same ID, how can I go about uniquely identifying these so I can manage them differently from each other?
- When should I use the COLLECT AGGREGATE ACTIVITY DATA clause versus the COLLECT ACTIVITY DATA clause?
- How does DB2 WLM work with the new AIX WPAR feature?
- What is the relationship between the DB2_OPT_MAX_TEMP_SIZE registry variable and a DB2 threshold based on SQLTEMPSPACE?

- How does DB2 workload manager differ from Query Patroller in how it deals with SQL statements invoked from a procedure or any other type of routine?
- Now that Query Patroller and DB2 Governor are deprecated, how do I migrate to DB2 workload manager?
- What are the licensing requirements for DB2 workload manager?

On which DB2 platforms can I use DB2 workload manager?

DB2 workload manager is available on all platforms supported by DB2 9.5 for Linux®, UNIX®, and Windows® or later. The optional tight integration offered between DB2 service classes and operating system service classes is available with AIX and Linux WLM.

Do I need Query Patroller to use DB2 workload manager?

No. Although Query Patroller and DB2 workload manager are both part of the Performance Optimization feature, they are independent of each other. In other words, one does not require Query Patroller to be installed to use DB2 workload manager or vice versa.

Why are the new WLM capabilities not integrated into Query Patroller?

DB2 workload manager represents a shift of emphasis in our workload management strategy to focus primarily on the ability to monitor and control active work once it has entered the DB2 execution environment. In order to provide the degree of control and monitoring desired by our customers for higher volumes of concurrent work requests while it is actually executing, yet with minimal overhead, we made the strategic decision to incorporate any new WLM technology directly into the DB2 engine infrastructure.

The approach of using a cooperative relationship with an auxiliary application to provide workload management for DB2 data server, such as that used with Query Patroller, simply would not be able to provide what was needed to satisfy customer requirements.

How does this new functionality affect Query Patroller and DB2 Governor?

The DB2 workload manager introduces an independent approach to workload management and does not rely on or interact with Query Patroller or DB2 Governor in any way. While Query Patroller and DB2 Governor are still functional, they are deprecated and no longer central to DB2 workload management strategy and no further investment is planned for them in future releases.

To ensure an easy transition, DB2 data server enables Query Patroller and DB2 Governor to coexist with the facilities provided by DB2 workload manager while still providing separate scopes of control. If Query Patroller is present, any work submitted for execution in the default user service class are intercepted and sent to Query Patroller. Work submitted for execution in other service classes defined by the database administrator are not presented to Query Patroller.

When DB2 9.5 or later is first installed, the default user service class is automatically defined and all incoming work is sent to it for execution. This means that any existing Query Patroller installation will continue to function as it did

before. It is only once the administrator introduces new DB2 service classes and begins to guide work away from the default user service class that the work seen by Query Patroller begins to change.

The story is essentially the same for DB2 Governor which, although it can watch agents in any service class, it is permitted to adjust the agent priority only for agents in the default user service class.

Note that DB2 workload manager is aware of and can control all work within DB2 including that within the default user service class. When Query Patroller is used, it is recommended to limit the use of DB2 workload manager to control work in the default user service class in order to avoid potential conflicts between Query Patroller and DB2 workload manager. It is always safe to use the monitoring features of DB2 workload manager.

Why should I not use just concurrency thresholds to control all my work ? Is that not how Query Patroller works?

While you can emulate the approach taken by Query Patroller by categorizing work by its estimated cost, mapping it to different service sub classes, and applying different concurrency thresholds on each service sub class, this is neither the recommended approach nor the best starting point. This approach does not deal with all the different types of work that execute within DB2 data server, only with DML SQL statements. Achieving a stable execution environment requires that all work executing within DB2 data server is controlled to one degree or another.

With DB2 workload manager, it is possible to separate and isolate competing workloads from each other and then affect their individual response times by changing the resources available to them. This is done by using DB2 service classes and manipulating the processor and prefetcher priorities each service class receives. It is recommended that you start here as this provides the groundwork for controlling all work executing within DB2 data server.

If you cannot separate work by its source (via a DB2 workload), then you can map all incoming work to a common service super class and use a DB2 work action set to separate work by different characteristics and assign it to different service sub classes. At this point, you can manipulate the resources available to each service class to achieve your objectives. Note that not all types of activities can be recognized within a work action set and any unrecognized ones will not be mapped to a different service class; they will remain in the one originally assigned to them.

If resource manipulation does not achieve the desired results, you can selectively apply other features of DB2 workload manager as needed until you achieve your objectives. This includes the application of DB2 thresholds, including concurrency thresholds. As most concurrency thresholds (such as the `CONCURRENTDBCOORDACTIVITIES` threshold, for example) coordinate activities across all database partitions, they impose a higher overhead on the activities that they manage. Introducing a concurrency threshold adds complexity to the execution environment; if care is not taken in the definition, unexpected or unintended results may be the consequence.

I am not on AIX, does this mean I do not have any control over processor resource or I/O activity?

Users on all platforms have the ability to control processor resource and prefetcher I/O activity between service classes using SQL (for CREATE and ALTER SERVICE CLASS statements, for example). To control CPU usage, users can use the agent priority attribute of the DB2 service class to set a relative processor priority for all threads that run in that service class. On AIX, users can also use this approach or they can choose to take advantage of AIX WLM for more advanced processor usage management. For prefetcher I/O activity, users on all platforms can set the prefetcher priority attribute of a DB2 service class to a value of high, medium or low. All service classes run with a medium prefetch priority by default.

Can I use AIX WLM to manage I/O activity?

Currently, AIX WLM does not support I/O activity controls at the thread level. Because DB2 Version 9.5 and later use a threaded model, it is not possible to use AIX WLM to control disk I/O activity. You can control DB2 prefetcher I/O activity by using the PREFETCH PRIORITY attribute of any DB2 service class.

Can I use AIX WLM to manage memory use?

DB2 data server uses primarily shared memory which is accessed by more than one agent from different service classes. For this reason, it is not possible to divide memory allocation between different service classes using AIX WLM.

Why is there not a tool to automatically upgrade from my Query Patroller configuration?

The main reason that a tool has not been provided to automatically upgrade your Query Patroller configuration to DB2 workload manager is that the type of controls and mechanisms available are different between the two. In Query Patroller, the primary control mechanism is to establish one or more query classes based on a range of estimated query costs (provided in timerons, in query compiler unit of measurement) and to set the desired concurrency rate for each class. The idea behind this approach is to control system resources by classifying and controlling when work of different size (different cost) is permitted to enter the system. With the DB2 WLM capabilities a number of additional control mechanisms are available, in addition to the control of concurrency, which enable you to approach the same workload management problems in different but more effective ways. For example, if the issue you face is one of conflict over processor usage, it is now possible to explicitly allocate CPU resources between competing work such that each grouping of work gets only the processor resource allocation that it requires and no more. This can be done without having to determine the query compiler estimates for all queries, determining the best set of estimate ranges to use, or the right concurrency rates.

In summary, while it is indeed possible to emulate the Query Patroller approach to solving workload management problems using DB2 WLM, this is not necessarily desirable. It is quite possible to solve many of the more common scenarios in a simpler way that is more effective for your overall system workload. Rather than perpetuate the Query Patroller approach, DB2 WLM provides a new opportunity to review your overall approach to workload management in light of the additional capabilities provided to determine whether it is the best one for your environment.

Another important reason not to move directly from the Query Patroller approach to that offered by DB2 WLM is to reduce the impact and risk to your environment when upgrading to DB2 9.5 or later. By enabling Query Patroller and DB2 workload manager to co-exist in the same environment, you can upgrade from one to the other in a controlled, granular manner.

Is there a way for WebSphere Application Server to pass the client information fields used by the DB2 workload?

WebSphere Application Server Version 6.0 and Version 6.1 can set or pass in the CLIENT INFO fields to DB2 data server, either explicitly by your applications (see: Passing client information to a database) or implicitly by having WebSphere Application Server do it for you (see: Implicitly set client information).

Why do the service class agent priority settings not seem to be in effect?

The service class agent priority setting does not take effect until an agent begins work on activities in that service class. An idle agent keeps the priority of the service class it last worked for until it joins a different service class. Another reason may be that the AGENTPRI dbm config parameter is set. Even though this parameter is deprecated as of Version 9.5, it does take precedence over the WLM service class setting. To use the WLM setting, reset the AGENTPRI config parameter to its default value, which is -1. On AIX, the instance owner must have CAP_NUMA_ATTACH and CAP_PROPAGATE capabilities to set a higher relative priority for agents in a service class.

Can I create multiple CONCURRENTDBCOORDACTIVITIES concurrency thresholds for the same set of work?

The simple answer to this question is yes. You can create one or more CONCURRENTDBCOORDACTIVITIES concurrency thresholds that apply to the same set of activities by defining them at the level of the database, the service class in which the work executes, or within a work action set applied at the database level. Be aware that each new concurrency threshold that applies to an activity implies additional overhead to enforce that concurrency threshold.

The more complex answer includes the following caution: verify that you actually need to use concurrency thresholds at all, let alone multiple ones. There may be simpler ways to address the scenario you are facing by using one or more of the other mechanisms and controls provided by DB2 workload manager. If you find yourself introducing one or more concurrency thresholds, you may have bypassed a simpler approach to address the problem. In general, concurrency thresholds for activities should be used at the database level, via a work action set, for disruptive activities that affect the entire system or go across service class boundaries, while concurrency thresholds at the service class level can be used to ensure proper sharing of resources between one service class and another (although a more effective technique may be to use the CONCURRENTWORKLOADACTIVITIES threshold on the workloads that contribute to the service class). There should rarely, if ever, be a case where you need to define a concurrency threshold for CONCURRENTDBCOORDACTIVITIES at the database level by itself.

Why is my work not assigned to the correct workload?

There are a number of reasons why a connection may not be mapped to the desired workload. The most common ones are the failure to grant USAGE

privilege on the workload, incorrect spelling of the case sensitive connection attributes, or the existence of a matching workload definition that is positioned earlier in the evaluation order.

Before a connection can be assigned to a workload, the connection attributes must match those of the workload definition, and the session authorization ID must have USAGE privilege on the workload. A common omission is to create the workload but not to grant USAGE privilege on the workload to users (See GRANT (Workload Privileges) statement). Only users with ACCESSCTRL, SECADM, or WLMADM authority can grant workload usage privilege to other users. Users with ACCESSCTRL, DATAACCESS, DBADM, SECADM, or WLMADM authority have implicit usage privilege on all workloads.

Connection attributes for workloads are case sensitive. For example: If the system user ID is uppercased, then the SYSTEM_USER connection attribute you specify must be in uppercase as well.

To establish why a connection is not being mapped to the expected workload, you should gather some information. Which workload is the work being mapped to? Is that workload before or after the one that you thought would be used when you look at the workload definitions in the order of evaluation? (Hint: try selecting the workload definitions ordered in ascending order by the value of the EVALUATIONORDER column in SYSCAT.WORKLOADS).

If you do not know what the connection attributes are for the target connection, you can find out the values for the connection in a number of different ways:

- Issue a query against the system using the WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97 table function while the connection is active
- Open a cursor on a connection and use the WLM_CAPTURE_ACTIVITY_IN_PROGRESS stored procedure against that cursor to have the activity information captured to the activities event monitor (Hint: do not forget to create and activate the activities information event monitor)
- Turn on the collection of detailed activity information for the workload being used by the connection, issue one statement in order to capture the activity information, and then turn off the collection.

Why does the DB2 data server not automatically create AIX service classes?

While having the DB2 data server automatically create corresponding AIX WLM service classes when DB2 service classes are created might reduce administrative overhead for system administrators, this is not available for a number of reasons:

- AIX WLM provides a wide variety of configuration options from which you can craft an environment suited for your unique needs. If DB2 data server were to automatically create AIX WLM service classes, then either the full variety of options provided by AIX WLM service classes would have to be surfaced within the DB2 service class DDL statements, greatly increasing DDL complexity, or the AIX WLM service classes would have to be standardized using a limited set of features which would not permit full utilization of AIX WLM features.
- DB2 service classes are created once in a database by SQL statements and are then automatically available on all database partitions when they are started. AIX WLM service class definitions need to be defined on each AIX machine that participates in the database, including any new ones that are added later.

- The use of tight integration with AIX WLM is an optional feature of DB2 workload manager which can be enabled or disabled at any time.

In the end, we decided that it is better for DB2 data servers not to create AIX WLM service classes when a DB2 service class is created. We believe that this gives the DB2 data server and our customers maximum flexibility.

Why does DB2 workload manager affect REORGCHK, IMPORT, EXPORT and other CLP commands?

These CLP commands are affected by DB2 workload manager thresholds, because the database engine cannot distinguish system requests originating with these utilities from other requests directly initiated by users within the CLP interactive front-end.

Is there a way to change the service class to which an activity is assigned while it is executing?

Yes, you can change the service subclass an activity is executing in to another service subclass within the same parent service superclass by defining a CPUTIMEINSC or SQLROWSINSC threshold with the REMAP ACTIVITY action on the original service subclass. Initially, DB2 workload manager maps an activity to a service class based on the relevant workload definition for the connection, modifies it as required if a work action set exists on that service class, and then sets up the DB2 agent to execute in the assigned service class. When an activity violates a threshold that has a REMAP ACTIVITY action defined, the agent remaps itself to the specified target service subclass (under the same superclass) once the threshold violation has been detected and the activity continues executing in the new service subclass.

Much of my batch work is done using CLP scripts under the same ID, how can I go about uniquely identifying these so I can manage them differently?

You have a couple of options:

An enhancement has been added to CLP so that the client application name is automatically set to the CLP script filename, with a **CLP** prefix preceding it (the value of this field at the server can be seen in the CURRENT CLIENT_APPLNAME special register). For example, if the CLP script filename is **batch.db2**, the CURRENT CLIENT_APPLNAME special register value is set to **CLP batch.db2** by CLP when that script is run. With this feature, it is possible for different CLP scripts to be associated with different workloads based on the client application name.

For example, to create a workload for CLP file **batch1.db2**, you can issue the following DDL statement:

```
CREATE WORKLOAD batch1 CURRENT CLIENT_APPLNAME ('CLP batch1.db2')
SERVICE CLASS class1
```

To create a workload for CLP file **batch2.db2**, you can issue the following DDL statement:

```
CREATE WORKLOAD batch2 CURRENT CLIENT_APPLNAME ('CLP batch2.db2')
SERVICE CLASS class2
```

Since these two batch files are associated with different workloads, they can be assigned to different service classes and managed differently.

Another option is the new stored procedure `WLM_SET_CLIENT_INFO`, which permits you to set the values of any of the client information fields at the server using a simple `CALL SQL` statement. By inserting a `CALL` statement into any of your existing `CLP` scripts, you can uniquely identify them using these fields and map them to different workload definitions.

For more information, see .

When should I use the `COLLECT AGGREGATE ACTIVITY DATA` clause versus the `COLLECT ACTIVITY DATA` clause?

The answer depends on why the monitoring is desired and what is to be done with the information.

Aggregate activity information is about the entire set of work that has executed within a service class and captures characteristics of this set; it does not capture details about individual activities. For normal operational monitoring, using the `COLLECT AGGREGATE ACTIVITY DATA` clause is preferred because it is very light-weight, can be gathered automatically by an event monitor for a historical record, and provides important information on overall response time patterns. If further insight is required on the type of work within a service class, it is also possible to use the `COUNT ACTIVITY` or `COLLECT AGGREGATE ACTIVITY DATA` actions within a DB2 work action set to gather more granular information about different types of work executing in a service class with minimal overhead.

In contrast, activity information contains detailed information about each and every activity that executes within the scope covered by the `COLLECT ACTIVITY DATA` clause. This clause can be specified on DB2 workloads, DB2 service classes, DB2 work action sets and DB2 thresholds. It permits further in-depth analysis of the individual activities that are captured, in order to understand the flow and type of SQL statements submitted by a new application, for example, or to look into performance tuning opportunities with tools such as the Explain facility or the Design Advisor. Because it captures much more information for each activity affected by it, the impact of using this clause is higher on affected activities than other monitoring methods and it should be carefully controlled.

How does DB2 WLM work with the new AIX WPAR feature?

All aspects of the DB2 workload manager will work within an AIX WPAR but because AIX WPARs do not support the use of AIX WLM features, the option to tightly integrate DB2 service classes with AIX WLM service classes is of no benefit in this environment.

What is the relationship between the `DB2_OPT_MAX_TEMP_SIZE` registry variable and DB2 thresholds based on `SQLTEMPSPACE`?

There is no direct relationship between these two things. The `DB2_OPT_MAX_TEMP_SIZE` registry variable is a directive to the query compiler to limit the amount of temporary table space that a query can use. This can cause the optimizer to choose a plan that is more expensive (potentially less efficient) but which uses less space in the system temporary table spaces. A DB2 threshold based on `SQLTEMPSPACE` does not affect the type of plan chosen by the

optimizer. It simply causes DB2 data server to monitor the usage of system temporary table space by that query at each database partition and generates a threshold violation if the stated limit is exceeded during normal processing.

How does DB2 workload manager differ from Query Patroller in how it deals with SQL statements invoked from a procedure or any other type of routine?

Query Patroller cannot schedule, hold, or queue SQL statements issued from within a routine. Since the QP query class relies on queuing to effect its control, this is a significant limitation.

DB2 workload manager enables control of all queries regardless of origin and while it does offer queuing as a secondary control mechanism, the primary control mechanism is resource control and prioritization via a DB2 service class. This means that you can allocate processor resource between DB2 service classes and possibly avoid using a queue at all.

Now that Query Patroller and DB2 Governor are deprecated, how do I migrate to DB2 workload manager?

Following the introduction of DB2 workload manager as the strategic workload management solution in DB2 Version 9.5, Query Patroller and the DB2 Governor have been deprecated and might be removed in a future release.

Query Patroller and DB2 Governor are still supported in this release, but you should begin adopting the new features and capabilities of DB2 workload manager, including those introduced in this release. Note that with DB2 workload manager, you have many more options, and you should explore them, which might require you to rethink your approach to controlling work on your DB2 data server in current workload management terms. The article Best Practices: DB2 Workload Management contains a section on the approach you should take when adapting DB2 workload manager. Task topics for migrating from Query Patroller and Governor are also available, as is additional information in this FAQ topic.

What are the licensing requirements for DB2 workload manager?

DB2 workload manager is licensed as part of the Performance Management Feature, and use of the full set of workload management functionality requires that you obtain this additional license.

Regardless of licensing, some workload management functionality is always used by your DB2 data server. The following is always available:

- You can use and alter the default service classes and workloads.
- You can create, alter and drop histogram templates.
- You can use the DB2 workload manager table functions and stored procedures.
- You can create, activate, stop and drop WLM event monitors.
- You can grant, alter and drop workload privileges.

The following functionality requires a license for the Performance Management Feature:

- The creation of any service class, workload, threshold, or work action set.

Chapter 2. Work identification

The first stage of implementing a DB2 workload manager solution identifies the work that runs on your data server. You identify work either by its origin or by its type.

Use workloads to identify work by where it originates or who submits it. You identify work by the application name or system authorization ID that submitted it, for example.

Work can be identified by type by pinpointing certain characteristics of the work with work classes. You identify statements such as those that only modify data on the data server, such as INSERT, UPDATE or DELETE statements, for example.

The following figure shows a number of different sources of work, coming from different users, groups, and applications.




Figure 2. Various sources and types of database activities on a data server

Activities

One way that you can monitor and control workloads is on the basis of individual activities. Each time your DB2 data server executes the access plan for an SQL or XQuery statement or executes the load utility a corresponding activity is created.

For workload monitoring, commonly used monitor elements provide information in terms of activity units. For example, you can obtain information about the workload volume and response time from monitor elements such as the activity execution time monitor element (`coord_act_exec_time`) and the high watermark for the concurrent activities (`concurrent_act_top`).

For workload control, most workload controls and thresholds apply to each activity. For example, the `ACTIVITYTOTALTIME` threshold controls the maximum time that your data server can spend processing an activity.

Statements or commands that trigger activities on your data server

The following statements or commands trigger activities on your data server:

- All DML statements

- All DDL statements
- The CALL statement
- The load utility

The life cycle of activities


The life cycle of an activity for a DML statement does not include processing that occurs before or outside of access plan execution. This implies that activity-based monitoring does not cover operations such as connecting to the database or compiling SQL into an access plan.

During its life cycle, an activity can spend time in various states, which are reported by the activity_state event monitor element. Some of the states an activity can be in are:

- EXECUTING - This state indicates that the coordinator agent is working on the activity. An activity that encounters a lock wait situation is reported as executing.
- IDLE - This state indicates that the coordinator agent is waiting for the next request from a client.
- QUEUED - Some thresholds include a built-in queue. This state indicates that the activity is waiting in the queue for its turn to begin executing.

Monitoring data for the activity is aggregated at the end of the lifetime of an activity.

The following figure shows how the lifetime of a long running query breaks down into queue time and execution time:



Types of SQL statements and application development

This section describes what activities are created for various SQL statements and identifies the start and end points in the lifetime of these activities. You can use this information to understand how SQL statements are monitored and controlled through activities.

SELECT statements: A SELECT statement is represented by one activity. This includes any cursor requests such as FETCH operations and subselects or subqueries. The activity starts when your DB2 data server starts processing an OPEN cursor statement or request, and ends when your data server has completed processing for a CLOSE cursor statement or request.

SELECT statements using WITH HOLD cursors: When a WITH HOLD cursor is used, an application can open a cursor within one unit of work and close the cursor in a subsequent unit of work. The cursor remains open for multiple units of work. The corresponding activity exists for as long as the cursor is open, because the life cycle of the activity ends only after the cursor is closed.

CALL statement and stored procedures: A CALL statement itself is represented by one activity, but the payload of the stored procedure can spawn nested activities as follows:

Table 1. Contents of stored procedures and activities they create

Contents of stored procedure	Additional activities created
A single SQL statement	One
No SQL statements in the stored procedure	0
SQL procedures, multiple SQL statements, and looping logic	Multiple activities, one corresponding to each invocation of each statement
A call to another stored procedure	Activities for that stored procedure

The activity associated with the CALL statement starts when your DB2 data server starts processing the statement or request and ends after the stored procedure processing is complete.

Triggers and UDFs: When a SQL statement calls a trigger or UDF, no additional activity is created. The work done by that trigger or UDF is accrued to the activity for the SQL statement that called it. Cases where the trigger or UDF executes additional SQL statements are handled like any statement execution, that is, an activity is created for each statement.

PREPARE statement: No activity is created, because activities are not created until an access plan is executed.

Nested activities

Nested activities do not significantly affect activity-based monitoring and control of workloads, but some additional information applies.

Activities that can have nested activities within them are:

- A stored procedure
- An anonymous block
- An autonomous routine
- A DML activity that executes a UDF
- A load from cursor (a load activity that has the cursor activity nested within it)
- A DML activity that is subject to a trigger that contains any of the activities above as part of the trigger definition

Nested activities are reported in monitoring information as follows:

- A nested activity is indicated by a non-zero parent UOW ID and a non-zero parent activity ID.
- A nested activity is not counted towards histograms or any statistics derived from histograms.
- Data for a nested activity is not also reported as part of the metrics for the parent activity. For example, if a procedure executed by a CALL statement performs an insert which consumes 10 seconds of processor time, that processor time is counted only towards the processor time metric for the insert activity and does not count towards the processor time metric for the parent CALL activity.

Workload control considers nested activities as follows:

- An activity nested inside a UDF or trigger does not contribute to the CONCURRENTDBCOORDACTIVITIES threshold.
- A cursor activity nested within a load activity does not contribute to the CONCURRENTDBCOORDACTIVITIES threshold.

Activities and the load utility

Running the load utility will generate several activities, one of which is a load activity and several others that are of type READ, WRITE, or OTHER. In the case of a load from cursor, an additional activity for the cursor the load activity is loading from is created. This cursor activity is a nested activity of the load activity.

DDL statements for DB2 workload manager

DB2 workload manager DDL statements consist of the CREATE, ALTER, and DROP statements you use to work with service classes, workloads, work class sets, work action sets, thresholds, and histograms.

The DB2 workload manager DDL statements are as follows:

- CREATE SERVICE CLASS, ALTER SERVICE CLASS, and DROP SERVICE CLASS
- CREATE WORKLOAD, ALTER WORKLOAD, and DROP WORKLOAD
- GRANT USAGE ON WORKLOAD and REVOKE USAGE ON WORKLOAD
- CREATE THRESHOLD, ALTER THRESHOLD, and DROP THRESHOLD
- CREATE WORK CLASS SET, ALTER WORK CLASS SET, and DROP WORK CLASS SET
- CREATE WORK ACTION SET, ALTER WORK ACTION SET, and DROP WORK ACTION SET
- CREATE HISTOGRAM TEMPLATE, ALTER HISTOGRAM TEMPLATE, and DROP HISTOGRAM TEMPLATE

Workload management DDL statements differ from other DB2 DDL statements:

- Only one uncommitted DB2 workload manager DDL statement is permitted at a time across all database partitions. If an uncommitted DB2 workload manager DDL statement exists, subsequent DB2 workload manager DDL statements wait until the uncommitted DB2 workload manager DDL statement is either committed or rolled back. DB2 workload manager DDL statements are processed in the order in which they are issued.
- Every DB2 workload manager DDL statement must be followed by a COMMIT or ROLLBACK statement.
- A DB2 workload manager DDL statement cannot be issued in an XA transaction. After a connection issues a DB2 workload manager DDL statement, the same connection must issue a COMMIT or ROLLBACK statement immediately after the DB2 workload manager DDL statement. With XA transactions, it is possible for multiple connections to join a transaction, and any of the connections can commit or roll back the transaction. In this situation, it is impossible to ensure that the workload management environment would be correctly implemented.
- DB2 for z/OS[®] does not recognize DB2 Database for Linux, UNIX, and Windows DB2 workload manager DDL statements.

Work identification by origin with workloads

Workloads identify incoming work based on its source so that it can later be managed. The source is determined using the attributes of the database connection under which the work is submitted.

The connection attributes are evaluated when the connection is established, and the connection is assigned to a given workload, creating a new occurrence of that workload. If any of the connection attributes change during the life of that connection, the workload assignment is reevaluated at the start of the next unit of work after the change and, if a new workload definition is to be assigned, the old workload occurrence for the previous assigned workload is ended and a new occurrence is started for the newly assigned workload definition. While each connection is assigned to one and only one workload at any one time, it is possible for there to be multiple connections assigned to the same workload at the same time resulting in the concurrent existence of multiple workload occurrences related to that definition.

Workload reevaluation occurs at the beginning of each unit of work in the event that the value of a connection attribute or the workload definition itself changes during the unit of work. This reevaluation might result in the connection being associated with a new workload, creating a different workload occurrence. For more information, see “Workload assignment” on page 21

For example, to assign all connections created by the application Accounts to a workload REPORTING, which maps the activities under those connections to run in the Marketing service class, issue a CREATE WORKLOAD statement such as the following:

```
CREATE WORKLOAD REPORTING APPLNAME('Accounts') SERVICE CLASS Marketing
```

This creates the following workload:



Figure 3. The REPORTING workload

To assign all activities created by the application Accounts under the connections that belong to the session user group Deptmgr to the SUMMARY workload, which maps the activities to the HumanResources service class, issue a statement such as the following:

```
CREATE WORKLOAD SUMMARY SESSION_USER_GROUP('Deptmgr') APPLNAME('Accounts')  
SERVICE CLASS HumanResources
```

A This creates the following workload:




Figure 4. The SUMMARY workload

You can view your workloads by querying the SYSCAT.WORKLOADS view and you can view the connection attributes that you specified for each workload by querying the SYSCAT.WORKLOADCONNATTR view. You can view who is authorized to use a workload by querying the SYSCAT.WORKLOADAUTH view.

SYSDEFAULTUSERWORKLOAD is the default workload. Any connection that is not assigned to a custom-defined workload during workload evaluation is assigned to this default workload, which ensures that all database connections are associated with a workload. Work assigned to the default workload SYSDEFAULTUSERWORKLOAD is executed in the SYSDEFAULTUSERCLASS service class by default.

Supported database connection attributes

You must specify at least one database connection attribute in the workload, and each connection attribute can have one or more values. If you do not specify a value for a specific connection attribute for the workload, the data server does not examine that attribute during workload evaluation.

Table 2. Connection attributes in a workload definition

Connection attribute	Description
Address	The actual communication address used by the client to communicate with the database server. The only protocol supported is TCP/IP. The address must be an IPv4 address, an IPv6 address, or a secure domain name.
Application name	The name of the application running at the client, as known to the data server. The application name is equivalent to the value shown in the Application name field in the system monitor output. See the appl_name monitor element for more information.
System authorization ID	The authorization ID of the user who connected to the database, as set in the SYSTEM_USER special register. You can change the value of SYSTEM_USER by connecting as a user with a different authorization ID.
Session authorization ID	The authorization ID that is used for the current session of the application, as set in the SESSION_USER special register. You can change the value of SESSION_USER using the SET SESSION AUTHORIZATION statement.
Group of session authorization ID	The groups to which the current session user belongs.
Role of session authorization ID	The roles granted to the current session user. For more information, see: <ul style="list-style-type: none"> • Roles • GRANT ROLE statement • REVOKE ROLE statement

Table 2. Connection attributes in a workload definition (continued)

Connection attribute	Description
Client user ID	The client user ID from the client information as set in the CURRENT CLIENT_USERID (or CLIENT USERID) special register. You can change the value of the client user ID by using the sqleseti (set client information) API or the WLM_SET_CLIENT_INFO procedure.
Client application name	The application name from the client information as set in the CURRENT CLIENT_APPLNAME (or CLIENT APPLNAME) special register. The client application name is equivalent to the value shown in the TP Monitor client application name field in the system monitor output. You can change the value of the client application name by using the sqleseti API or the WLM_SET_CLIENT_INFO procedure.
Client workstation name	The workstation name from the client information as set in the CURRENT CLIENT_WRKSTNNAME (or CLIENT WRKSTNNAME) special register. You can change the value of the client workstation name by using the sqleseti API or the WLM_SET_CLIENT_INFO procedure.
Client accounting string	The accounting string from the client information as set in the CURRENT CLIENT_ACCTNG (or CLIENT ACCTNG) special register. You can change the value of the client accounting string by using the sqleseti API or the WLM_SET_CLIENT_INFO procedure.

Use of wild cards in connection attributes

Some connection attributes support the specification of an asterisk (*) as a wild card in the CREATE WORKLOAD and ALTER WORKLOAD statements. You can use wild cards in situations where a connection attribute can take on several similar values, which can be matched by a regular expression using wild cards, without defining connection attributes for each of the possible values.

The wild card asterisk (*) matches zero or more characters. If you need to match an asterisk, use a double asterisk (**) to specify the asterisk as a literal character.

For example: If you have several accounts receivable applications (*accrec01*, *accrec02* ... *accrec15*) that you all want to belong to the same workload for equal treatment by DB2 workload manager, define the *CURRENT CLIENT_APPLNAME('accrec*')* connection attribute to match all of these applications when you create or alter your workload. Similarly, an *acc*rec* accounts receivable application (a name that includes an asterisk character) is matched by the *CURRENT CLIENT_APPLNAME('acc**rec')* connection attribute.

The following workload connection attributes support the use of wild cards:

- APPLNAME
- CURRENT CLIENT_ACCTNG
- CURRENT CLIENT_APPLNAME
- CURRENT CLIENT_USERID
- CURRENT CLIENT_WRKSTNNAME

Set client information to identify requests

In a three-tier client/server environment, the database connection is established by the application server that is working on behalf of the clients. The application server can use the `sqleseti` API or the `WLM_SET_CLIENT_INFO` procedure to pass client information to the DB2 data server; otherwise, only the information about the application server is passed, and that information is likely to be the same for all client requests that are routed through this application server. By specifying client information attributes such as the client user ID, client application name, client workstation name, and client accounting string in the workload definition, you can assign users running from different clients to different workloads (and to different service classes).

Connection attribute evaluation order

As you analyze the usage characteristics of your environment, you can use the `CREATE WORKLOAD` statement to create your own workloads and map them to specific service classes. When you create the workload, you define both the values that are used to evaluate the connection attributes during workload assignment and the order in which the workload is evaluated relative to other workloads. Because more than one workload can match incoming connection attributes, being able to change the evaluation order enables you to determine which matching workload is chosen. Whether or not the session user has the `USAGE` privilege on the workload also determines which matching workload is chosen. For more information, see “Workload assignment” on page 21.

The following figure shows multiple requests being evaluated against workloads in the order A, B, C, and D, then assigned to specific workloads and executed in the applicable service class. Requests that cannot be matched to an existing workload are matched to the `SYSDEFAULTUSERWORKLOAD` workload and executed by default in the `SYSDEFAULTUSERCLASS` service superclass. For information about the types of activities that run in the default maintenance class and default system class, see “Default service superclasses and subclasses” on page 66.




Figure 5. Service classes and workloads

Workload assignment

At the beginning of the first unit of work after a database connection is established, the data server assigns the connection to a workload by evaluating the connection attributes of each workload that is enabled.

The order in which the workloads are evaluated is determined by the EVALUATIONORDER column value of each workload in the SYSCAT.WORKLOADS table. If a workload with matching connection attributes is found, the data server checks whether the current session user has the USAGE privilege on the workload. If the user has the USAGE privilege on the matching

workload, the workload assignment is complete, and the connection is assigned to that workload. If the user does not have the USAGE privilege on the matching workload, the data server continues to evaluate workloads until it finds a matching workload on which the session user has the USAGE privilege. If no matching workload is found, the data server attempts to use the SYSDEFAULTUSERWORKLOAD workload. If the current session user does not have the USAGE privilege on that workload, SQL4707N is returned, and the unit of work is rejected. Otherwise, the connection is assigned to the SYSDEFAULTUSERWORKLOAD workload.

You can set the evaluation order by using the POSITION keyword of the CREATE WORKLOAD or ALTER WORKLOAD statement, as follows:

- By specifying the absolute position of the workload in the evaluation order, as shown in the following example:

```
CREATE WORKLOAD...POSITION AT 2
```

POSITION AT 2 means that the workload is to be positioned second in the evaluation order. A matching workload that is positioned higher in the evaluation order is evaluated first. That is, if the workloads at both position 2 and position 3 match, the workload at position 2 is evaluated before the workload at position 3.

If the position that you specify on the CREATE WORKLOAD or ALTER WORKLOAD statement is greater than the total number of existing workloads, the workload is positioned next to last in the evaluation order, before the SYSDEFAULTUSERWORKLOAD workload. The effect is the same as specifying POSITION LAST on the CREATE WORKLOAD or ALTER WORKLOAD statement.

- By using the POSITION BEFORE *workload-name* or POSITION AFTER *workload-name* keyword, where *workload-name* is an existing workload. This keyword specifies the position of a new or altered workload relative to another workload in the evaluation order, as shown in the following example:

```
ALTER WORKLOAD...POSITION BEFORE workload2
```

If you do not specify the POSITION keyword, by default, the new workload is positioned after the other defined workloads in the evaluation order but before the SYSDEFAULTUSERWORKLOAD workload, which is always considered last.

Workload reassignment

A connection can switch service superclasses at unit of work boundaries by changing workloads. In DB2 workload manager, a unit of work boundary is the point when a connection disassociates with its current transaction. The following events cause a unit of work boundary: Commit, rollback, XA end (success), XA commit, and XA rollback.

The workload assignment is reevaluated at the beginning of a new unit of work if the data server detects that one of the following events occurred:

- A relevant connection attribute changed. See the table in “Work identification by origin with workloads” on page 17 for a list of connection attributes that you can specify in a workload definition. Workload reevaluation also occurs if the current session authorization ID changes because the database connection switches because of a trusted context. For more information, see Trusted contexts and trusted connections.
- You created or altered a workload.

- You granted the USAGE privilege on a workload to a user, a group, or a role or revoked the USAGE privilege on a workload from a user, group, or role.

A connection cannot be reassigned to a different workload while an activity that spans a unit of work boundary is still active. An activity can be an operation that maintains resources across multiple UOWs, such as a load operation, a stored procedure or table function, or a WITH HOLD cursor. The current workload occurrence runs until all activities complete. The workload reassignment then occurs at the beginning of the next unit of work.

An attempted workload assignment or reassignment results in an SQL4707N error if either of the following cases exists:

- The data server attempts to assign the connection to a workload that is dispermitted access to the database. For more information, see “Preventing occurrences of a workload from accessing the database” on page 31.
- The data server attempts to assign the connection to the SYSDEFAULTUSERWORKLOAD workload, but the current session user does not have the USAGE privilege on this workload.

If you have ACCESSCTRL, DATAACCESS, DBADM, SECADM, or WLMADM authority, you can assign your database connection to the SYSDEFAULTADMWORKLOAD workload, the default administrator workload. See “Taking corrective action using the default administration workload” on page 26 for more information.

XA transactions and workload reassignment

XA calls such as XA_END (success), XA commit, and XA rollback issue a DB2 COMMIT or ROLLBACK, which indicates the end of a unit of work. Because workload reevaluation can occur at the beginning of a unit of work, these XA calls can initiate workload reevaluation, although the reason for workload reevaluation is not directly related to the XA transaction itself.

Default workloads

The default user workload SYSDEFAULTUSERWORKLOAD provides a workload for your data server to which all connections are assigned initially. The default administration workload SYSDEFAULTADMWORKLOAD permits you to take corrective administrative action that cannot otherwise be performed. Both workloads are created at database creation time and you cannot drop them.

The default user workload (SYSDEFAULTUSERWORKLOAD)

Connections that are assigned to the default user workload are mapped to the default user service superclass SYSDEFAULTUSERCLASS, which provides the default execution environment. You can map connections to user-defined service classes by creating user defined workloads. In addition, you can alter SYSDEFAULTUSERWORKLOAD so that it maps connections to a different service class than SYSDEFAULTUSERCLASS.

You can view the SYSDEFAULTUSERWORKLOAD workload by querying the SYSCAT.WORKLOADS table.

The following table shows shows the columns returned for the SYSDEFAULTUSERWORKLOAD workload in the SYSCAT.WORKLOADS view, along with values and whether you can modify these values. See “Workload

assignment” on page 21 for information on how to assign a connection to the SYSDEFAULTUSERWORKLOAD workload.

Table 3. SYSDEFAULTUSERWORKLOAD entry in SYSCAT.WORKLOADS

Column	Value	Modifiable using the ALTER WORKLOAD statement if you have DBADM or WLMADM authority (and SQLADM for COLLECT clauses)
WORKLOADID	1	No
WORKLOADNAME	SYSDEFAULTUSERWORKLOAD	No
EVALUATIONORDER	Second last one	No
CREATE_TIME	Timestamp of database creation	No
ALTER_TIME	Timestamp of the last update of the workload definition	No (but the data server modifies this column when you update the workload definition)
ENABLED	Y	No
ALLOWACCESS	Y	Yes
SERVICECLASSNAME	SYSDEFAULTSUBCLASS	Yes
PARENTSERVICECLASSNAME	SYSDEFAULTUSERCLASS	Yes
COLLECTAGGACTDATA	N	Yes
COLLECTACTDATA	N	Yes
COLLECTACTPARTITION	C	Yes
COLLECTDEADLOCK	W	Yes
COLLECTLOCKTIMEOUT	W	Yes
COLLECTLOCKWAIT	N	Yes
LOCKWAITVALUE	0	Yes
COLLECTACTMETRICS	N	Yes
COLLECTUOWDATA	N	Yes
EXTERNALNAME	NULL	No
REMARKS	BLANK	Yes

For more information, see SYSCAT.WORKLOADS.

The default administration workload (SYSDEFAULTADMWORKLOAD)

This workload permits ACCESSCTRL, DATAACCESS, DBADM, SECADM, or WLMADM users to query the database and perform administrative or monitoring tasks at any time, but is typically used in cases when:

- The workload to which the administrator is assigned is not permitted to access the database (that is, the DISALLOW DB ACCESS keyword of the CREATE WORKLOAD or ALTER WORKLOAD statement was specified for the workload).
- A threshold was violated, preventing the administrator from performing work on the database.

The SYSDEFAULTADMWORKLOAD workload differs from other workloads in the following ways:

- You cannot drop or disable it.
- You cannot specify DISALLOW DB ACCESS for it.
- None of the thresholds apply to occurrences of this workload and the activities in it.
- You can run this workload only in the SYSDEFAULTUSERCLASS service superclass. See “Default service superclasses and subclasses” on page 66 for more information.
- You can assign a connection to this workload by using the SET WORKLOAD command from the CLP interface, or by invoking the WLM_SET_CLIENT_INFO stored procedure (and specifying SYSDEFAULTADMWORKLOAD for the **client_workload** parameter). For more information, see “Taking corrective action using the default administration workload” on page 26.

You can view the SYSDEFAULTADMWORKLOAD workload by querying the SYSCAT.WORKLOADS table. The following table shows the columns returned for the SYSDEFAULTADMWORKLOAD workload in the SYSCAT.WORKLOADS catalog view, along with values and whether you can modify these values:

Table 4. SYSDEFAULTADMWORKLOAD entry in SYSCAT.WORKLOADS

Column	Value	Modifiable using the ALTER WORKLOAD statement if you have DBADM or WLMADM authority (and SQLADM for COLLECT clauses)
WORKLOADID	2	No
WORKLOADNAME	SYSDEFAULTADMWORKLOAD	No
EVALUATIONORDER	Last one	No
CREATE_TIME	Timestamp of database creation	No
ALTER_TIME	Timestamp of the last update of the workload definition	No (but the data server modifies this column when you update the workload definition)
ENABLED	Y	No
ALLOWACCESS	Y	No
SERVICECLASSNAME	SYSDEFAULTSUBCLASS	No
PARENTSERVICECLASSNAME	SYSDEFAULTUSERCLASS	No
COLLECTAGGACTDATA	N	Yes
COLLECTACTDATA	N	Yes
COLLECTACTPARTITION	C	Yes
COLLECTDEADLOCK	W	Yes
COLLECTLOCKTIMEOUT	W	Yes
COLLECTLOCKWAIT	N	Yes
LOCKWAITVALUE	0	Yes
COLLECTACTMETRICS	N	Yes
COLLECTUOWDATA	N	Yes
EXTERNALNAME	NULL	No
REMARKS	BLANK	Yes

For more information, see SYSCAT.WORKLOADS.

Taking corrective action using the default administration workload

The default administration workload SYSDEFAULTADMWORKLOAD is a special DB2-supplied workload definition that is not subject to any DB2 thresholds. Use this workload to take corrective action that cannot otherwise be performed, such as altering prohibitive threshold definitions that prevent all activities from running in a workload.

Use the SET WORKLOAD command (or the WLM_SET_CLIENT_INFO procedure) to assign a connection to the default administration workload SYSDEFAULTADMWORKLOAD.

Although you require no special authority to use the SET WORKLOAD command, you require ACCESSCTRL, DATAACCESS, DBADM, SECADM, or WLMADM authority to assign a connection to the default administration workload. Otherwise, SQL0552N is returned during workload assignment.

Because this workload is not affected by thresholds, it has limited workload management control and is not recommended for use in submitting regular day-to-day work.

To assign a connection to the default administration workload, issue the SET WORKLOAD command as follows:

```
SET WORKLOAD TO SYSDEFAULTADMWORKLOAD
```

When the command takes effect depends on when you issue it:

- If you issue the SET WORKLOAD TO SYSDEFAULTADMWORKLOAD command before the connection to the database, after the connection is established, it is assigned to SYSDEFAULTADMWORKLOAD at the beginning of the first unit of work.
- If you issue the SET WORKLOAD TO SYSDEFAULTADMWORKLOAD command at the beginning of a unit of work, after a connection to the database is established, the connection is assigned to SYSDEFAULTADMWORKLOAD when the first request that is not an sqleseti (Set Client Information) request is submitted.
- If you issue the SET WORKLOAD TO SYSDEFAULTADMWORKLOAD command at the middle of a unit of work, after a connection is established, the connection is assigned to SYSDEFAULTADMWORKLOAD at the beginning of the next unit of work.

When a connection is assigned to SYSDEFAULTADMWORKLOAD, workload reassignment is performed at the beginning of the next unit of work if either of the following situations occurs:

- You revoke SYSADM or DBADM authority from the session user. In this situation, SQL0552N is returned.
- You issue a SET WORKLOAD TO AUTOMATIC command. This command indicates that the next unit of work should not be assigned to the SYSDEFAULTADMWORKLOAD workload and that a normal workload evaluation is to be performed at the beginning of the next unit of work. For more information, see “Workload assignment” on page 21.

Example

The following example shows how you can use the SYSDEFAULTADMWORKLOAD workload to take corrective action when no other corrective action is possible.

If you create a severely prohibitive concurrency threshold so that no activities can execute, because the threshold is always being exceeded, the same threshold can prevent you from correcting the problem. To be able to alter the prohibitive threshold, you must first set the workload so that the work runs in the default administration workload. Because activities running in this workload are not subject to thresholds, you can correct the problem and set the workload (for your ID) back to the default behavior.

The threshold that is the cause of the problem is created accidentally with the following statement. Concurrency should have been set to 100 but was set to 0. This threshold effectively prevents any activity from executing:

```
CREATE THRESHOLD PROHIBITIVE FOR DATABASE ACTIVITIES
  ENFORCEMENT DATABASE WHEN CONCURRENTDBCOORDACTIVITIES > 0
  STOP EXECUTION
```

Note: This statement is intended only to show you how a severely prohibitive threshold might be created. You should not issue this statement.

If you attempt to execute even just a simple SELECT statement, an error is returned, because concurrency is set to 0:

```
SELECT * FROM SYSIBM.SYSTABLES
```

```
SQL4712N The threshold "PROHIBITIVE" has been exceeded. Reason code = "6".
SQLSTATE=5U026
```

Before you can take corrective action, you must set the workload to the default administration workload:

```
SET WORKLOAD TO SYSDEFAULTADMWORKLOAD
```

This statement can be issued only by someone with ACCESSCTRL, DATAACCESS, DBADM, SECADM, or WLMADM authority and causes any connection to be assigned to the SYSDEFAULTADMWORKLOAD workload, where activities are not subject to the prohibitive threshold.

The problem can now be corrected by altering the threshold so that activities can run:

```
ALTER THRESHOLD PROHIBITIVE WHEN CONCURRENTDBCOORDACTIVITIES > 100 STOP EXECUTION
```

Once corrected, change the workload back so that the connection will no longer be assigned to SYSDEFAULTADMWORKLOAD but to whatever workload it was assigned to before:

```
SET WORKLOAD TO AUTOMATIC
```

The same SELECT statement used before should now complete successfully:

```
SELECT * FROM SYSIBM.SYSTABLES
```

```
...
```

```
DB20000I The SQL command completed successfully.
```

Creating a workload

Use a CREATE WORKLOAD statement to add a workload to the catalogs.

To create a workload, you require WLMADM or DBADM authority.

See the following topics for more information about prerequisites:

- “DDL statements for DB2 workload manager” on page 16
- Naming rules

To create a workload:

1. Specify one or more of the following properties for the workload using the CREATE WORKLOAD statement:
 - The name of the workload.
 - The connection attributes. The incoming connection must supply matching connection attributes to those that you specified for the workload for a match to occur. For more information, see “Work identification by origin with workloads” on page 17. When specifying the connection attributes, note that values are ORed and attributes are ANDeD: for example, UserID (bob OR sue OR frank) AND Application (SAS).
 - A value that indicates whether occurrences of this workload are permitted to access the database. By default, occurrences of this workload are permitted to access the database.
 - A value that indicates whether the workload is enabled or disabled. By default, the workload is enabled.
 - The service class under which occurrences of this workload are to be executed. The SYSDEFAULTUSERCLASS service superclass is the default.

If you specify a user-defined service superclass and do not map the workload to run in a user-defined service subclass under the service superclass, the workload occurrences will run in the SYSDEFAULTSUBCLASS service subclass of the service superclass.

Note: You cannot specify the SYSDEFAULTSUBCLASS service subclass under any service superclass, including the SYSDEFAULTUSERCLASS service superclass.

If you do not want occurrences of the workload to run in the SYSDEFAULTSUBCLASS service subclass, you can map the workload for execution in a user-defined service subclass through the workload. You can also use a work action to map the workload to a different service subclass (for more information, see “Work actions and work action sets” on page 128).

- The position of the workload relative to other workloads when cached in the memory. The position of the new workload determines the order in which it is evaluated during workload assignment. By default, the new workload is positioned last, which means that it is evaluated last, immediately before the default user workload is considered. For more information, see “Workload assignment” on page 21.
- The monitoring activity metrics collection level for activities submitted by connections associated with this workload. The default activity metrics collection setting for a workload is NONE. Note that the effective activity collection setting for activities is the combination of both the workload activity metrics collection level and the **mon_act_metrics** database configuration parameter.

- The type of activity information to collect. By default, no information for activities associated with the workload is sent to an activities event monitor.
 - The aggregate activity information to collect. The aggregate activity information used for the workload only changes after the CREATE WORKLOAD operation is committed.
 - The lock timeout events information to collect. By default, data about a lock event is sent to the locking event monitor, if one is active, when the lock event occurs, but previous lock timeout events are not sent (WITHOUT HISTORY).
 - The deadlock information to collect. By default, data about a deadlock event is sent to the locking event monitor, if one is active, when the deadlock event occurs, but previous deadlock events are not sent (WITHOUT HISTORY).
 - The lock wait information to collect. By default, no lock wait information is collected if a lock is not acquired within the set wait time.
 - The unit of work information for each transaction associated with this workload to send to the unit of work event monitor, if one is active, when a unit of work ends. By default, no unit of work information is sent.
 - The histogram templates that the workload should use as templates for its histograms. The histogram templates specified are reflected in the SYSCAT.HISTOGRAMTEMPLATEUSE view. For more information on histograms and histogram templates, see “Histograms in workload management” on page 180.
2. Commit your changes. When you commit your changes the workload is added to the SYSCAT.WORKLOADS view. Committing the change causes a workload reevaluation to take place at the beginning of the next unit of work of each application. Depending on which workload is chosen, the application might be reassigned to a different workload.

After you create a workload, you might need to grant the USAGE privilege on it to one or more session users. (Session users with WLMADM or DBADM authority have an implicit privilege to use any workload.) Even if a connection provides an exact match to the connection attributes of the workload, if the session user does not have the USAGE privilege on the workload, the data server does not consider the workload when performing workload evaluation. For more information, see “Granting the USAGE privilege on a workload” on page 33.

Altering a workload

An ALTER WORKLOAD statement changes a workload in the catalogs.

To alter a workload, you require SQLADM, WLMADM or DBADM authority. To specify any clause other than a COLLECT clause, the authorization id must include WLMADM or DBADM authority.

See “DDL statements for DB2 workload manager” on page 16 for more information about prerequisites.

To alter a workload:

1. Specify one or more of the following properties for the workload using the ALTER WORKLOAD statement:
 - The connection attributes. You can add connection attributes to and drop connection attributes from the workload definition unless it is the SYSDEFAULTUSERWORKLOAD or SYSDEFAULTADMWORKLOAD workload. The incoming connection must supply matching connection

attributes to those that you specified for the workload for a match to occur. For more information, see “Work identification by origin with workloads” on page 17. To see the connection attributes for a workload, query the SYSCAT.WORKLOADCONNATTR view.

- A value that indicates whether an occurrence of this workload is permitted to access the database. By default, an occurrence of this workload is permitted to access the database. You cannot remove database access from the SYSDEFAULTADMWORKLOAD workload.
 - A value that indicates whether the workload is enabled or disabled. By default, the workload is enabled. You cannot disable the SYSDEFAULTUSERWORKLOAD or the SYSDEFAULTADMWORKLOAD workload.
 - The service class under which occurrences of this workload are to be executed. The SYSDEFAULTUSERCLASS service superclass is the default. If you specify a user-defined service superclass, you can specify a service subclass under the service superclass. You cannot specify the SYSDEFAULTSUBCLASS subclass under any service superclass, including the SYSDEFAULTUSERCLASS service superclass. In addition, you cannot specify the SYSDEFAULTSYSTEMCLASS or SYSDEFAULTMAINTENANCECLASS service superclass.
 - The position of the workload relative to other workloads, which determines the order in which the workload is evaluated during workload assignment. You cannot specify the position of the SYSDEFAULTUSERWORKLOAD or the SYSDEFAULTADMWORKLOAD workload. For more information, see “Workload assignment” on page 21.
 - The type of activity information to collect. By default, no information for activities associated with the workload is sent to an activities event monitor.
 - The monitoring activity metrics collection level for activities submitted by connections associated with this workload. Note that the effective activity collection setting for activities is the combination of both the workload activity metrics collection level and the **mon_act_metrics** database configuration parameter.
 - The aggregate activity information to collect. The aggregate activity information used for the workload only changes after the ALTER WORKLOAD operation is committed.
 - The lock timeout event information to send to the locking event monitor, if one is active, when a lock event occurs.
 - The deadlock information to send to the locking event monitor, if one is active, when a deadlock event occurs.
 - The lock wait information to collect.
 - The unit of work information for each transaction associated with this workload to send to the unit of work event monitor, if one is active, when a unit of work ends.
 - The histogram templates that the workload should use as templates for its histograms. The histogram templates specified are reflected in the SYSCAT.HISTOGRAMTEMPLATEUSE view. For more information on histograms and histogram templates, see “Histograms in workload management” on page 180.
2. Commit your changes. When you commit your changes the workload is updated in the SYSCAT.WORKLOADS view. The committed change causes a workload reevaluation to take place at the beginning of the next unit of work of each application. Depending on which workload is chosen, the application might be reassigned to a different workload.

You might need to grant the USAGE privilege on it to one or more session users. (Session users with DBADM authority have an implicit privilege to use any workload.) Even if a connection provides an exact match to the connection attributes of the workload, if the session user does not have the USAGE privilege on the workload, the data server does not associate the connection with the workload to create an occurrence of the workload. For more information, see “Granting the USAGE privilege on a workload” on page 33.

Permitting occurrences of a workload to access the database

If you have a workload that is not permitted to access the database but now want to permit occurrences of that workload to run, alter the workload so that it is permitted to access the database. By default, when a workload is created, it is permitted to access the database.

To alter a workload so that it can access a database, you require WLMADM or DBADM authority.

See “DDL statements for DB2 workload manager” on page 16 for more information about prerequisites.

When you prevent a workload from accessing the database, the data server still examines that workload when performing workload assignment. However, all occurrences of that workload are rejected with an error. To permit a workload to access the database:

1. Use the ALLOW DB ACCESS option of the ALTER WORKLOAD statement to permit the workload to access the database. For example, to permit a workload called WL1 to access the database, specify the following statement:

```
ALTER WORKLOAD WL1 ALLOW DB ACCESS
```
2. Commit your changes. When you commit your changes workload is updated in the SYSCAT.WORKLOADS view.

Altering a workload to permit its occurrences to access the database takes effect when the data server analyzes the next unit of work for that workload. For example, if you specified DISALLOW DB ACCESS for workload A and alter the workload by specifying ALLOW DB ACCESS, new occurrences of workload A are permitted to execute. Previously, any occurrence of workload A would have been rejected with an error.

Preventing occurrences of a workload from accessing the database

Use this task to control which workloads can access the database. Before a workload occurrence begins to run, the data server checks whether the workload is permitted to access the database. If you dispermit the workload occurrence from accessing the database, an error is returned indicating that the workload occurrence is rejected.

To prevent a workload from accessing the database, you require WLMADM or DBADM authority.

See “DDL statements for DB2 workload manager” on page 16 for more information about prerequisites.

Preventing a workload occurrence differs from disabling a workload. When you disable a workload, the workload definition is not cached in memory and is therefore not considered for workload assignment. To prevent a workload from accessing a database:

1. Use the `DISALLOW DB ACCESS` option of the `ALTER WORKLOAD` statement, as shown in the following example:

```
ALTER WORKLOAD workload-name DISALLOW DB ACCESS ...
```
2. Commit your changes. When you commit your changes, the workload is updated in the `SYSCAT.WORKLOADS` view.

Altering a workload to prevent its occurrences from accessing a database takes effect at the beginning of the next unit of work for workload occurrences that are already running. For example, if you specify `ALLOW DB ACCESS` for workload A and alter the workload by specifying `DISALLOW DB ACCESS`, occurrences of workload A that are already running receive an SQL error at the beginning of the next unit of work. New occurrences of workload A are rejected.

Enabling a workload

The DB2 data server checks the connection attributes specified for a workload against the connection attributes of the current session. The data server does not consider a disabled workload when it looks for a matching workload.

To alter a workload, you require `WLMADM` or `DBADM` authority.

See “DDL statements for DB2 workload manager” on page 16 for more information about prerequisites.

By default, a workload is enabled when you create it. If you create a workload as disabled, you must enable it for the data server to consider the workload when it performs workload evaluation.

To enable a workload:

1. Identify the workload that you want to enable. You can display the set of disabled workloads by querying the `SYSCAT.WORKLOADS` view, as shown in the following example:

```
SELECT * FROM SYSCAT.WORKLOADS WHERE ENABLED='N'
```
2. Use the `ALTER WORKLOAD` statement to enable the disabled workload:

```
ALTER WORKLOAD...ENABLE
```

If the `ALTER WORKLOAD` statement is successful, the definition for the workload is written to the database catalog.

3. Commit your changes. When you commit your changes the workload is updated in the `SYSCAT.WORKLOADS` view.

Enabling a workload takes effect at the beginning of the next unit of work. At that point, a workload reevaluation occurs, and the data server considers the newly enabled workload when it performs workload reevaluation.

Disabling a workload

Use this task to prevent specific workloads from being considered during workload assignment. If you disable a workload, the data server does not consider it when it looks for a matching workload. Instead, the data server assigns the unit

of work to the next matching workload. If no custom-defined workload matches, the work is assigned to the default workload.

To create or alter a workload, you require WLMADM or DBADM authority.

See “DDL statements for DB2 workload manager” on page 16 for more information about prerequisites.

To disable a workload:

1. Use the DISABLE option of the ALTER WORKLOAD statement to disable the workload:

```
ALTER WORKLOAD...DISABLE
```

2. Commit your changes. When you commit your changes, the workload is updated in the SYSCAT.WORKLOADS view.

Disabling a workload takes effect at the beginning of the next unit of work. At that point, a workload reevaluation occurs, and the connection is assigned to the next enabled workload that matches the connection attributes and for which there is authorization.

Granting the USAGE privilege on a workload

For a workload to be associated with a connection, the session user must have the USAGE privilege on that workload. Users with the ACCESSCTRL, DATAACCESS, DBADM, SECADM, or WLMADM authority implicitly have the USAGE privilege on all workloads.

To use the GRANT USAGE ON WORKLOAD statement, you require ACCESSCTRL, SECADM, or WLMADM authority

See “DDL statements for DB2 workload manager” on page 16 for more information about prerequisites.

When the data server finds a workload that matches the attributes of an incoming connection, the data server checks whether the session user has the USAGE privilege on that workload. If the session user does not have the USAGE privilege on that workload, the data server looks for the next matching workload. (In other words, the workloads for which the session user does not have the USAGE privilege are treated as if they do not exist.) Therefore, the workload USAGE privilege gives you the ability to further control which workload among the matching workloads a user, group, or role should be assigned to. For example, you can define more than one workload with the same connection attributes and grant the USAGE privilege on each of these workloads to only certain users, groups, or roles. For more information, see “Workload assignment” on page 21.

The client can set the client user ID, client application name, client workstation name, and client accounting string (which are some of the connection attributes that are used to assign a connection to a workload) without authorization. Therefore, the workload USAGE privilege also permits you to control which session user has the authority to use a workload.

You can view the USAGE privilege information by querying the SYSCAT.WORKLOADAUTH view.

If you create a database without the RESTRICT option, the USAGE privilege on the SYSDEFAULTUSERWORKLOAD workload is granted to PUBLIC at database

creation time. Otherwise, you must explicitly grant the USAGE privilege on this workload to non-WLMADM and non-DBADM users. If the session user does not have the USAGE privilege on any of the workloads, including SYSDEFAULTUSERWORKLOAD, SQL4707N is returned when the data server attempts to associate a workload with the database connection.

To grant the USAGE privilege on a workload:

1. Use the GRANT USAGE ON WORKLOAD statement. You can grant the USAGE privilege to specific users, groups, roles, or PUBLIC. For example, to grant the USAGE privilege on the ACCOUNTS workload to the CPA group, you would issue the following statement:

```
GRANT USAGE ON WORKLOAD ACCOUNTS TO GROUP CPA
```

You cannot grant the USAGE privilege on the SYSDEFAULTADMWORKLOAD workload. The SYSDEFAULTADMWORKLOAD workload can only be used by ACCESSCTRL, DATAACCESS, DBADM, SECADM, or WLMADM users who issue the SET WORKLOAD TO SYSDEFAULTADMWORKLOAD command.

2. Commit your changes. When you commit your changes, the SYSCAT.WORKLOADAUTH view is updated. Until the GRANT statement is committed, the data server cannot consider the workload when performing workload assignment for the newly authorized users, groups, or roles.

Revoking the USAGE privilege on a workload

Use the REVOKE USAGE ON WORKLOAD statement to revoke the USAGE privilege on a workload.

To use the REVOKE USAGE ON WORKLOAD statement, you require ACCESSCTRL, SECADM, or WLMADM authority.

See “DDL statements for DB2 workload manager” on page 16 for more information about prerequisites.

You cannot explicitly revoke the USAGE privilege on the SYSDEFAULTADMWORKLOAD workload. Only ACCESSCTRL, DATAACCESS, DBADM, SECADM, or WLMADM users who issue the SET WORKLOAD TO SYSDEFAULTADMWORKLOAD command can use this workload. Therefore, REVOKE USAGE ON WORKLOAD statements do not work for SYSDEFAULTADMWORKLOAD.

To revoke the USAGE privilege on a workload:

1. Use the REVOKE USAGE ON WORKLOAD statement. You can revoke the USAGE privilege from specific users, groups, roles, or PUBLIC. For example, to revoke the USAGE privilege on the ACCOUNTS workload from PUBLIC, you would specify the following statement:

```
REVOKE USAGE ON WORKLOAD ACCOUNTS FROM PUBLIC
```

2. Commit your changes. When you commit your changes, the SYSCAT.WORKLOADAUTH view is updated. Until the REVOKE statement is committed, the data server considers the workload when performing workload assignment.

Dropping a workload

Dropping a workload removes it from the database catalog.

To drop a workload, you require WLMADM or DBADM authority.

See “DDL statements for DB2 workload manager” on page 16 for more information about prerequisites.

To drop a workload:

1. Disable the workload by specifying the ALTER WORKLOAD statement. See “Disabling a workload” on page 32 for more information. Disabling the workload prevents new occurrences of the workload from being able to run against the database.
2. Ensure that no occurrences of this workload are running by using the WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97 table function. For more information, see WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97 table function.
The WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97 table function returns the application handles corresponding to the active workload occurrences. You can use the FORCE APPLICATION command to terminate the applications using the application handles.
3. Drop the workload by specifying the DROP WORKLOAD statement. For example, to drop the ACCTNG workload, specify the following statement:

```
DROP WORKLOAD ACCTNG
```
4. Commit your changes. When you commit your changes, the workload is removed from the SYSCAT.WORKLOADS view. In addition, authorization information for the workload is removed from the SYSCAT.WORKLOADAUTH view.

Example: Workload assignment

At the beginning of the first unit of work after a database connection is established, the data server assigns the connection to a workload by evaluating the connection attributes of each workload that is enabled. Workload reevaluation occurs at the beginning of each unit of work if the value of a connection attribute or the workload definition itself changes during the unit of work, or if the usage granted on the workload changes.

The following figure shows a workload assignment. Users in the Marketing group who submit queries through AppA are assigned to the APPAQUERIES workload. They are not assigned to the PAYROLL workload, even though PAYROLL is positioned before APPAQUERIES, because the definition of workload PAYROLL specifies the SESSION_USER GROUP keyword as Finance. Users in the Finance group who submit queries through AppA are assigned to the FINANCE workload. They are not assigned to the PAYROLL workload, even though it is more specific and specifies both AppA and Finance in its definition, because the FINANCE workload is positioned before the PAYROLL workload. Users in the Marketing group who submit queries through AppB are assigned to the SYSDEFAULTUSERWORKLOAD workload, because none of the connection attributes specified in the FINANCE, PAYROLL, or APPAQUERIES workload definitions match the AppB application or Marketing group.




Figure 6. Example of workload assignment

1 In the preceding figure, the CREATE WORKLOAD statements are as follows:

```
CREATE WORKLOAD PAYROLL APPLNAME ('AppA') SESSION_USER GROUP ('FINANCE')
SERVICE CLASS SC1
```

```
CREATE WORKLOAD APPAQUERIES APPLNAME('AppA') POSITION LAST
SERVICE CLASS SC2
```

```
CREATE WORKLOAD FINANCE SESSION_USER GROUP ('FINANCE') SERVICE CLASS SC1
POSITION BEFORE PAYROLL
```

In a three-tier client/server environment, the database connection is established by the application server that is working on behalf of the clients. The application server can use the sqleseti (set client information) API to pass client information to the DB2 data server; otherwise, only the information from the application server is passed, and that information is likely to be the same for all client requests that are routed through this application server. When the data server assigns units of work from different clients to different workloads (and to different service classes), the data server uses the client information attributes (that is, the client user ID, client application name, client workstation name, and client accounting string) as criteria for associating a unit of work with a workload.

The following figure shows an example of a three-tier environment where queries are submitted by different user applications, (marketing.exe, auditing.exe, and reporting.exe), through an application server that establishes a connection to the database using the session user APPUSER. Three workloads are defined: one for queries submitted by marketing.exe, one for queries submitted by reporting.exe,

and one for the rest of the queries. As shown in the figure, to assign queries submitted by `marketing.exe` to the `MARKETING` workload, the application server calls the `sqleseti` API to set the value of the `CURRENT CLIENT_APPLNAME` special register to `marketing.exe`. Similarly, to assign queries submitted by `reporting.exe` to the `REPORTING` workload, the server calls `sqleseti` to set the value of the `CURRENT CLIENT_APPLNAME` special register to `reporting.exe`. Note that in the figure, when the server calls `sqleseti` to set the `CURRENT CLIENT_USERID` special register to `Lidia` (with nothing else changing; that is, the client application name is still set to `reporting.exe`), no workload reassignment occurs because there is no workload defined specifically with the `CURRENT CLIENT_USERID` set to `Lidia`.

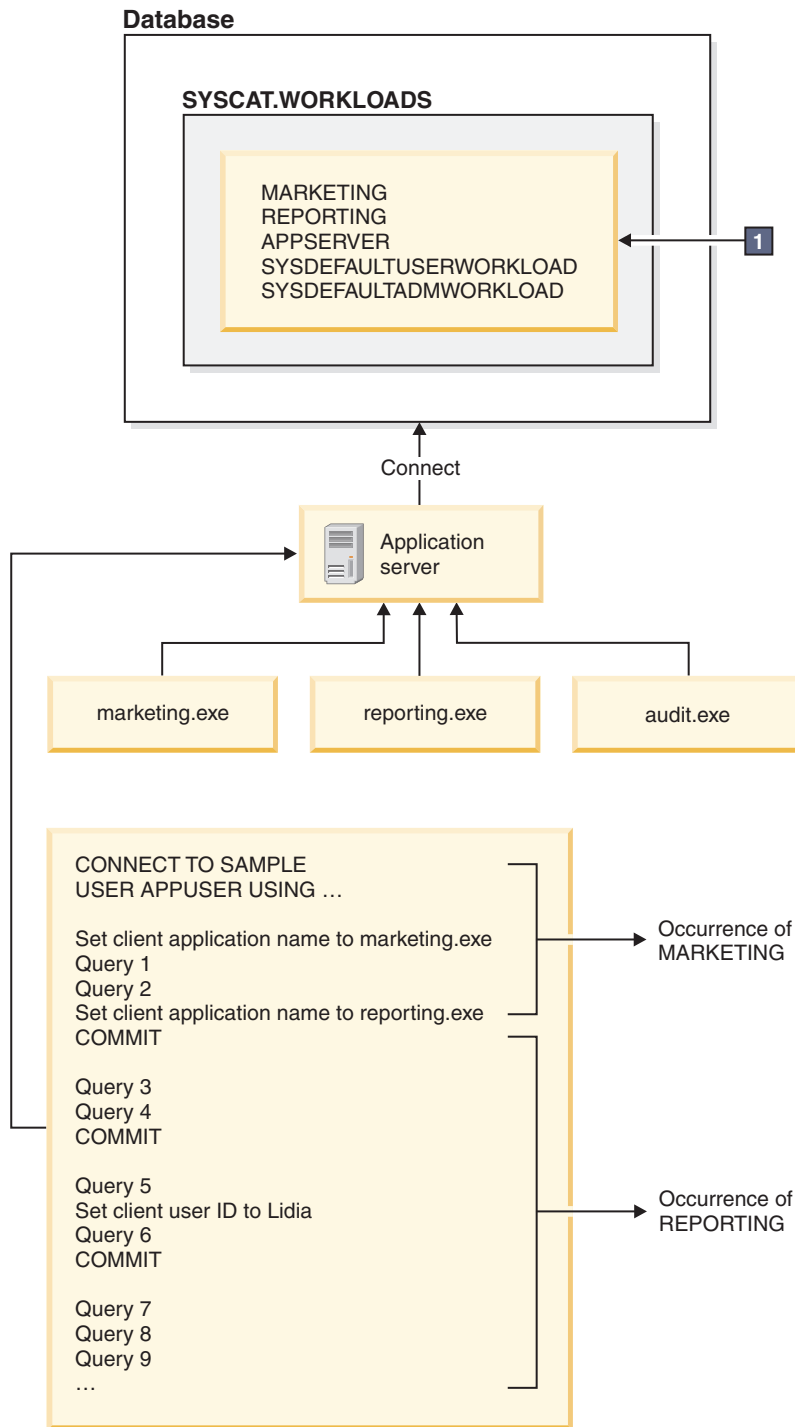


Figure 7. Example of workload assignment in a three-tier environment

The following statements are used to define the workloads specified in box **1** in the previous figure:

```
CREATE WORKLOAD MARKETING SESSION_USER ('APPUSER')
CURRENT_CLIENT_APPLNAME ('marketing.exe') SERVICE CLASS SC2
POSITION AT 1
```

```
CREATE WORKLOAD REPORTING SESSION_USER ('APPUSER')
CURRENT_CLIENT_APPLNAME ('reporting.exe') SERVICE CLASS SC4
```

POSITION AFTER MARKETING

```
CREATE WORKLOAD APPSERV SESSION_USER ('APPUSER')
SERVICE CLASS SC1
```

Example: Workload assignment when workload attributes have single values

The example in this topic shows how the data server performs workload assignment. In this example, only one value is specified for each workload connection attribute.

Assume that the following workloads exist in the catalog:

Table 5. Workloads in the catalog

Evaluation order	Workload name	ADDRESS	APPLNAME	SYSTEM_USER	SESSION_USER	SESSION_USER GROUP	SESSION_USER ROLE	CURRENT_CLIENT_USERID	CURRENT_CLIENT_APPLNAME	CURRENT_CLIENT_WRKSTNNAME	CURRENT_CLIENT_ACCTNG
1	REPORTS		AppA								
2	INVENTORY REPORT		AppB	LYNN		ACCOUNTING	TELEMKTR				
3	SALES REPORT		AppC	KATE	KATE		SALESREP				
4	AUDIT REPORT		AppB			ACCOUNTING	FINANALYST				
5	EXPENSE REPORT		AppA	TIM			EXPENSE APPROVER				
6	AUDIT RESULT				LYNN			LYNN			Audit Group

Assume that a database connection with the following attributes is established:

Table 6. Database connection attributes

ADDRESS	APPLNAME	SYSTEM_USER	SESSION_USER	SESSION_USER GROUP	SESSION_USER ROLE	CURRENT_CLIENT_USERID	CURRENT_CLIENT_APPLNAME	CURRENT_CLIENT_WRKSTNNAME	CURRENT_CLIENT_ACCTNG
9.26.53.111	AppA	TIM	TIM	FINANCE	FINANALYST, EXPENSE APPROVER	NULL	NULL	NULL	Business account

When the first unit of work is submitted, the data server checks each workload in the catalog, starting with the first workload in the list, and processes the workloads in ascending order until it finds a workload with matching attributes. When a matching workload is found, the unit of work runs under an occurrence of that workload. When determining which workload to assign the connection to, the data server compares the connection attributes in deterministic order.

The data server first checks the REPORTS workload for a match. The REPORTS workload is first in the list.

Table 7. REPORTS workload in the catalog

Evaluation order	Workload name	ADDRESS	APPLNAME	SYSTEM_USER	SESSION_USER	SESSION_USER GROUP	SESSION_USER ROLE	CURRENT_CLIENT_USERID	CURRENT_CLIENT_APPLNAME	CURRENT_CLIENT_WRKSTNNAME	CURRENT_CLIENT_ACCTNG
1	REPORTS		AppA								

The data server checks the connection attributes in the following deterministic order:

1. APPLNAME. The value of APPLNAME, AppA, for the database connection matches the value of APPLNAME for the REPORTS workload.
2. SYSTEM_USER, which is not set in the workload definition. Any value (including a null value) is considered a match.
3. SESSION_USER, which is not set in the workload definition. Any value is considered a match.

4. SESSION_USER GROUP, which is not set in the workload definition. Any value is considered a match.
5. SESSION_USER ROLE, which is not set in the workload definition. Any value is considered a match.
6. CURRENT_CLIENT_USERID, which is not set in the workload definition. Any value is considered a match.
7. CURRENT_CLIENT_APPLNAME, which is not set in the workload definition. Any value is considered a match.
8. CURRENT_CLIENT_WRKSTNNAME, which is not set in the workload definition. Any value is considered a match.
9. CURRENT_CLIENT_ACCTNG, which is not set in the workload definition. Any value is considered a match.

In this situation, because of the explicit and implicit matches between the connection attributes of the REPORTS workload and the information passed on the connection, the data server selects the REPORTS workload as a potential match. After selecting a workload, the data server then checks whether the session user has the USAGE privilege on the workload. Assuming that the session user TIM has the USAGE privilege on the REPORTS workload, that workload is used for the connection. If, however, TIM does not possess the USAGE privilege on the REPORTS workload, the data server continues by checking the INVENTORYREPORT workload for a match.

Assume that you want TIM to be assigned to the EXPENSEREPORT workload because that workload has additional connection attributes specified. In this situation, you would alter the evaluation order of the workloads to position EXPENSEREPORT before REPORTS in the workload list:

```
ALTER WORKLOAD EXPENSEREPORT POSITION AT 1
```

You could also use the following SQL statement to achieve the same result:

```
ALTER WORKLOAD EXPENSEREPORT BEFORE REPORTS
```

To ensure that the ALTER WORKLOAD statement takes effect, you must immediately issue a COMMIT statement after the ALTER WORKLOAD statement. The effect of the ALTER WORKLOAD statement on the catalog is as follows:

Table 8. Workloads in the catalog after repositioning the EXPENSEREPORT workload

Evaluation order	Workload name	APPLNAME	SYSTEM_USER	SESSION_USER	SESSION_USER GROUP	SESSION_USER ROLE	CURRENT_CLIENT_USERID	CURRENT_CLIENT_APPLNAME	CURRENT_CLIENT_WRKSTNNAME	CURRENT_CLIENT_ACCTNG
1	EXPENSE REPORT	AppA	TIM			EXPENSE APPROVER				
2	REPORTS	AppA								
3	INVENTORY REPORT	AppB	LYNN		ACCOUNTING	TELEMKTR				
4	SALES REPORT	AppC	KATE	KATE		SALESREP				
5	AUDIT REPORT	AppB			ACCOUNTING	FINANALYST				
6	AUDIT RESULT			LYNN			LYNN			Audit Group

If TIM does not already have the USAGE privilege on the EXPENSEREPORT workload, you must issue the following statements (the COMMIT statement ensures that the GRANT statement takes effect):

```
GRANT USAGE ON WORKLOAD EXPENSEREPORT TO USER TIM
COMMIT
```

At the beginning of the next unit of work, workload reassignment occurs, and the data server assigns the connection from TIM to the EXPENSEREPORT workload.

In addition, new units of work submitted by other connections that have the same attributes are also associated with the EXPENSEREPORT workload.

Example: Workload assignment for a unit of work when multiple workloads exist

The example in this topic shows how the data server performs workload evaluation to assign the connection to an existing workload.

Assume that the following workloads are defined in the catalog:

Table 9. Workloads in the catalog

Evaluation order	Workload name	APPLNAME	SYSTEM _USER	SESSION _USER	SESSION _USER GROUP	SESSION _USER ROLE	CURRENT CLIENT _USERID	CURRENT CLIENT _APPLNAME	CURRENT CLIENT _WRKSTNNAME	CURRENT CLIENT _ACCTNG
1	EXPENSE REPORT	AppB	TIM			EXPENSE APPROVER				
2	REPORTS	AppB								
3	INVENTORYREPORT	AppA	LYNN		ACCOUNTING	TELEMKTR				
4	SALES REPORT	AppC	KATE	KATE		SALESREP				
5	AUDIT REPORT	AppA			ACCOUNTING	FINANALYST				
6	AUDIT RESULT			LYNN			LYNN			Audit Group

Suppose that a database connection with the following attributes is established:

Table 10. Database connection attributes

APPLNAME	SYSTEM _USER	SESSION _USER	SESSION _USER GROUP	SESSION _USER ROLE	CURRENT CLIENT _USERID	CURRENT CLIENT _APPLNAME	CURRENT CLIENT _WRKSTNNAME	CURRENT CLIENT _ACCTNG
AppA	LYNN	LYNN	ACCOUNTING	FINANALYST, SALESREP	LYNN	NULL	wrkstn2	Audit group

When the first unit of work is submitted, the data server checks each workload in the catalog in ascending evaluation order and stops when it finds a workload whose connection attributes match those supplied by the connection. When it checks the workloads, the data server compares the connection attributes in deterministic order.

First, the data server checks the EXPENSEREPORT workload:

Table 11. EXPENSEREPORT workload in the catalog

Evaluation order	Workload name	APPLNAME	SYSTEM _USER	SESSION _USER	SESSION _USER GROUP	SESSION _USER ROLE	CURRENT CLIENT _USERID	CURRENT CLIENT _APPLNAME	CURRENT CLIENT _WRKSTNNAME	CURRENT CLIENT _ACCTNG
1	EXPENSEREPORT	AppB	TIM			EXPENSE APPROVER				

Because the APPLNAME attribute in the workload definition is AppB but the APPLNAME attribute passed by the connection is AppA, no match is possible. The data server proceeds to the REPORTS workload, which is second in the list:

Table 12. REPORTS workload in the catalog

Evaluation order	Workload name	APPLNAME	SYSTEM _USER	SESSION _USER	SESSION _USER GROUP	SESSION _USER ROLE	CURRENT CLIENT _USERID	CURRENT CLIENT _APPLNAME	CURRENT CLIENT _WRKSTNNAME	CURRENT CLIENT _ACCTNG
2	REPORTS	AppB								

Again, the APPLNAME attribute in the workload definition is AppB, which does not match AppA. The data server proceeds to the third workload in the list, INVENTORYREPORT:

Table 13. INVENTORYREPORT workload in the catalog

Evaluation order	Workload name	APPLNAME	SYSTEM_USER	SESSION_USER	SESSION_USER GROUP	SESSION_USER ROLE	CURRENT_CLIENT_USERID	CURRENT_CLIENT_APPLNAME	CURRENT_CLIENT_WRKSTNNAME	CURRENT_CLIENT_ACCTNG
3	INVENTORYREPORT	AppA	LYNN		ACCOUNTING	TELEMKTR				

The data server checks for a match between the submitted connection attributes and the INVENTORYREPORT workload. The attributes are checked in the following order:

1. APPLNAME. Both the workload definition and the connection have a value of AppA, so a match occurs.
2. SYSTEM_USER. Both the workload definition and the connection have a value of LYNN, so a match occurs.
3. SESSION_USER. The connection passed a value of LYNN. Because the SESSION_USER attribute is not set for the workload, any value, including a null value, that is passed by the connection matches.
4. SESSION_USER GROUP. Both the workload definition and the connection have a value of ACCOUNTING, so a match occurs.
5. SESSION_USER ROLE. The workload definition specifies the value TELEMKTR, but the connection supplied the values of FINANALYST and SALESREP. No match occurs for this attribute.

The data server stops trying to match the INVENTORYREPORT workload and the connection attributes and proceeds to the fourth workload in the list, SALESREPORT:

Table 14. SALESREPORT workload in the catalog

Evaluation order	Workload name	APPLNAME	SYSTEM_USER	SESSION_USER	SESSION_USER GROUP	SESSION_USER ROLE	CURRENT_CLIENT_USERID	CURRENT_CLIENT_APPLNAME	CURRENT_CLIENT_WRKSTNNAME	CURRENT_CLIENT_ACCTNG
4	SALESREPORT	AppC	KATE	KATE		SALESREP				

Because the APPLNAME of the SALESREPORT workload definition is AppC, no match occurs with the connection (which passed a value of AppA for APPLNAME). The data server then proceeds to the fifth workload in the list, AUDITREPORT:

Table 15. AUDITREPORT workload in the catalog

Evaluation order	Workload name	APPLNAME	SYSTEM_USER	SESSION_USER	SESSION_USER GROUP	SESSION_USER ROLE	CURRENT_CLIENT_USERID	CURRENT_CLIENT_APPLNAME	CURRENT_CLIENT_WRKSTNNAME	CURRENT_CLIENT_ACCTNG
5	AUDITREPORT	AppA			ACCOUNTING	FINANALYST				

The data server compares the attributes of the AUDITREPORT workload and the connection in the deterministic order:

1. APPLNAME. Both the workload definition and the connection have a value of AppA, so a match occurs.
2. SYSTEM_USER. The connection passed a value of LYNN. Because the SYSTEM_USER attribute is not set for the workload, any value passed by the connection matches.
3. SESSION_USER. The connection passed a value of LYNN. Because the SESSION_USER attribute is not set for the workload, any value passed by the connection matches.
4. SESSION_USER GROUP. Both the workload and the connection have a value of ACCOUNTING for this attribute, so a match occurs.
5. SESSION_USER ROLE. Both the workload and the connection have a value of FINANALYST for this attribute, so a match occurs.

6. CURRENT_CLIENT_USERID. Because the CURRENT_CLIENT_USERID attribute is not set for the workload, any value passed by the connection matches.
7. CURRENT_CLIENT_APPLNAME. Because the CURRENT_CLIENT_APPLNAME attribute is not set for the workload, any value passed by the connection matches.
8. CURRENT_CLIENT_WRKSTNNAME. Because the CURRENT_CLIENT_WRKSTNNAME attribute is not set for the workload, any value passed by the connection matches.
9. CURRENT_CLIENT_ACCTNG. Because the CURRENT_CLIENT_ACCTNG attribute is not set for the workload, any value passed by the connection matches.

After processing all the connection attributes and finding a matching workload, the data server checks whether the session user has the USAGE privilege on the workload. Assume that LYNN does not have the USAGE privilege on the AUDITREPORT workload. In this situation, although all of the connection attributes match, this workload is not associated with the connection. The data server proceeds to the sixth workload in the evaluation list, AUDITRESULT:

Table 16. AUDITRESULT workload in the catalog

Evaluation order	Workload name	APPLNAME	SYSTEM_USER	SESSION_USER	SESSION_USER GROUP	SESSION_USER ROLE	CURRENT_CLIENT_USERID	CURRENT_CLIENT_APPLNAME	CURRENT_CLIENT_WRKSTNNAME	CURRENT_CLIENT_ACCTNG
6	AUDITRESULT			LYNN			LYNN			Audit Group

The data server compares the attributes of the AUDITRESULT workload and the connection in the deterministic order:

1. APPLNAME. Because the APPLNAME attribute is not set for the workload, any value passed by the connection matches.
2. SYSTEM_USER. Because the SYSTEM_USER attribute is not set for the workload, any value passed by the connection matches.
3. SESSION_USER. Both the workload and the connection have a value of LYNN for this attribute, so a match occurs.
4. SESSION_USER GROUP. Because the SESSION_USER GROUP attribute is not set for the workload, any value passed by the connection matches.
5. SESSION_USER ROLE. Because the SESSION_USER ROLE attribute is not set for the workload, any value passed by the connection matches.
6. CURRENT_CLIENT_USERID. Both the workload and the connection have a value of LYNN for this attribute, so a match occurs.
7. CURRENT_CLIENT_APPLNAME. Because the CURRENT_CLIENT_APPLNAME attribute is not set for the workload, any value passed by the connection matches.
8. CURRENT_CLIENT_WRKSTNNAME. Because the CURRENT_CLIENT_WRKSTNNAME attribute is not set for the workload, any value passed by the connection matches.
9. CURRENT_CLIENT_ACCTNG. Both the workload and the connection have a value of Audit Group for this attribute, so a match occurs.

After processing all of the connection attributes and finding a matching workload, the data server checks whether the session user has the USAGE privilege on the workload. In this situation, assume that the session user LYNN has the USAGE

privilege on the AUDITRESULT workload. Because all of the connection attributes match and the session user has the USAGE privilege, the connection is assigned to the AUDITRESULT workload.

Example: Workload assignment when workload attributes have multiple values

The example in this topic shows how the data server performs workload assignment. In this example, some of the workload definitions permit more than one value for a connection attribute.

Assume that the following workloads are defined in the catalog:

Table 17. Workloads in the catalog

Evaluation order	Workload name	APPLNAME	SYSTEM_USER	SESSION_USER	SESSION_USER GROUP	SESSION_USER ROLE	CURRENT CLIENT_USERID	CURRENT CLIENT_APPLNAME	CURRENT CLIENT_WRKSTNNAME	CURRENT CLIENT_ACCTNG
1	ITEMINQ		KYLE, GEORGE		RETAIL, SALES					
2	DAILY TRANS REPORT	AppC		KYLE, CAROL	SALES, ACCOUNTING					
3	SALES SUMMARY	AppA, AppB				ACCOUNTANT, FINANALYST				

Assume that a database connection with the following attributes is established:

Table 18. Database connection attributes

APPLNAME	SYSTEM_USER	SESSION_USER	SESSION_USER GROUP	SESSION_USER ROLE	CURRENT CLIENT_USERID	CURRENT CLIENT_APPLNAME	CURRENT CLIENT_WRKSTNNAME	CURRENT CLIENT_ACCTNG
AppC	LINDA	KYLE	SALES	ACCOUNTANT	LINDA	NULL	NULL	Business Account

When the first unit of work is submitted, the data server checks each workload in the catalog in ascending evaluation order and stops when it finds a workload whose connection attributes match those supplied by the connection. When it checks the workloads, the data server compares the connection attributes in deterministic order.

First, the data server checks the ITEMINQ workload:

Table 19. ITEMINQ workload in the catalog

Evaluation order	Workload name	APPLNAME	SYSTEM_USER	SESSION_USER	SESSION_USER GROUP	SESSION_USER ROLE	CURRENT CLIENT_USERID	CURRENT CLIENT_APPLNAME	CURRENT CLIENT_WRKSTNNAME	CURRENT CLIENT_ACCTNG
1	ITEMINQ		KYLE, GEORGE		RETAIL, SALES					

The data server checks for a match between the submitted connection attributes and the ITEMINQ workload. The attributes are checked in the following order:

1. APPLNAME. Because the APPLNAME attribute is not set for the workload, any value, including a null value, that is passed by the connection matches.
2. SYSTEM_USER. The connection passed a value of LINDA. However, the ITEMNO workload values are KYLE and GEORGE. No match occurs for this attribute.

The data server stops trying to match the ITEMNO workload and the connection and proceeds to the second workload in the list, DAILYTRANSREPORT:

Table 20. DAILYTRANSREPORT workload in the catalog

Evaluation order	Workload name	APPLNAME	SYSTEM_USER	SESSION_USER	SESSION_USER GROUP	SESSION_USER ROLE	CURRENT CLIENT_USERID	CURRENT CLIENT_APPLNAME	CURRENT CLIENT_WRKSTNNAME	CURRENT CLIENT_ACCTNG
2	DAILYTRANSREPORT	AppC		KYLE, CAROL	SALES, ACCOUNTING					

The data server compares the attributes of the DAILYTRANSREPORT workload and the connection in deterministic order:

1. APPLNAME. Both the workload definition and the connection have a value of AppC, so a match occurs.
2. SYSTEM_USER. Because the SYSTEM_USER attribute is not set for the workload, any value, including a null value, that is passed by the connection matches.
3. SESSION_USER. The SESSION_USER value passed on the connection is KYLE, which is a match with one of the workload SESSION_USER values. If the connection had passed CAROL, this would also be a match because both KYLE and CAROL are specified as part of the DAILYTRANSREPORT workload definition.
4. SESSION_USER GROUP. The SESSION_USER GROUP value passed on the connection is SALES, which matches the SALES value specified for the workload SESSION_USER GROUP attribute. If the connection had passed ACCOUNTING, this would also be a match because both SALES and ACCOUNTING are specified in the workload definition.
5. SESSION_USER ROLE. Because the SESSION_USER ROLE attribute is not set for the workload, any value passed by the connection matches.
6. CURRENT_CLIENT_USERID. Because the CURRENT_CLIENT_USERID attribute is not set for the workload, any value passed by the connection matches.
7. CURRENT_CLIENT_APPLNAME. Because the CURRENT_CLIENT_APPLNAME attribute is not set for the workload, any value passed by the connection matches.
8. CURRENT_CLIENT_WRKSTNNAME. Because the CURRENT_CLIENT_WRKSTNNAME attribute is not set for the workload, any value passed by the connection matches.
9. CURRENT_CLIENT_ACCTNG. Because the CURRENT_CLIENT_WRKSTNNAME attribute is not set for the workload, any value passed by the connection matches.

After processing all of the connection attributes and finding a matching workload for the connection, the data server checks whether the session user has the USAGE privilege on the workload. In this situation, assume that the session user KYLE has the USAGE privilege on the DAILYTRANSREPORT workload. Because all connection attributes match and the session user has the USAGE privilege, the connection is assigned to the DAILYTRANSREPORT workload.

Work identification by type of work with work classes

In addition to using connection attributes that focus on the origin of activities with workloads, you can identify activities based on the type of work through the creation of a work class set containing a work class.

A work class is a method of categorizing individual database activities based on attributes of the activities. If a work class has a work action defined for it, the work action will be applied to the work class and determines how the activities in the work class are managed. For more information, see “Apply controls to types of activities with work action sets” on page 125.

The following table shows the type keywords available for work classes and the SQL statements that correspond to the different keywords. Except for the load utility, all the statements in the table below are intercepted immediately before

execution in the processing of an EXECUTE, EXECUTE IMMEDIATE, or OPEN request. The load utility, when issued from a client, might issue requests before starting the actual load operation on the data server.

Table 21. Work types

Work type keyword	Applicable SQL statements
READ, including SET statements with embedded READ SQL	<ul style="list-style-type: none"> All SELECT statements (select into, values into, full select) <p>Exception: SELECT statements containing a DELETE, INSERT, or UPDATE are not included.</p> <ul style="list-style-type: none"> All XQuery statements
WRITE, including SET statements with embedded WRITE SQL	<ul style="list-style-type: none"> All UPDATE statements (searched, positioned) All DELETE statements (searched, positioned) All INSERT statements (values, subselect) All MERGE statements All SELECT statements containing a DELETE, INSERT, or UPDATE statement
CALL	<p>CALL statement</p> <p>The CALL statement is only classified under the CALL and ALL work class types.</p> <p>Note: Both anonymous blocks and autonomous routines are classified as CALL statements.</p>
DML, including SET statements with embedded READ or WRITE SQL	All statements that are classified under the READ and WRITE work class types.
DDL	<ul style="list-style-type: none"> All ALTER statements All CREATE statements COMMENT statement DECLARE GLOBAL TEMPORARY TABLE statement DROP statement FLUSH PACKAGE CACHE statement All GRANT statements REFRESH TABLE All RENAME statements All REVOKE statements SET INTEGRITY statement
LOAD	<p>Load utility</p> <p>The load utility is only classified under the LOAD and ALL work class types.</p>

Table 21. Work types (continued)

Work type keyword	Applicable SQL statements
ALL	<p>All database activity.</p> <p>Note: If the action is a threshold, the database activity that the threshold is applied to depends on the type of threshold. For example, if the threshold type is ESTIMATEDSQLCOST, only DML activity with an estimated cost (in timerons) is affected by the threshold.</p> <p>For more information, see “Example: Working with a work class defined with the ALL keyword” on page 59.</p>

The following figure shows a hierarchical view of the work type keywords:

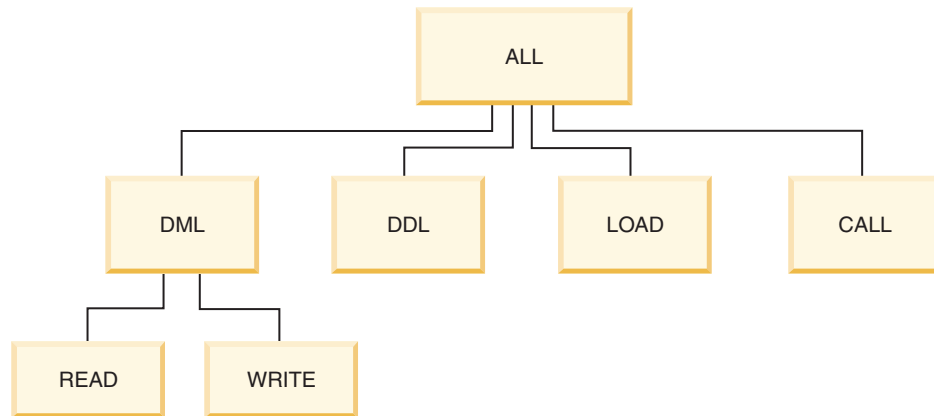


Figure 8. Work type keywords

SQL statements that do not fall under any of the available keywords are not classified, and behave as though no work class and work class set exists. For example, if the statement is SET SCHEMA and the only work class in the work class set has a work type of DML, that statement is not classified and no work action can be applied to it. So, if the action is MAP, the SET SCHEMA activity runs in the default service subclass (SYSDEFAULTSUBCLASS). If the action is a threshold, no threshold is applied to the activity.

Additional identification

Work classes also permit you to use predictive elements in the identification for DML work (or READ and WRITE statements). Predictive elements are useful because they provide information about database activities that can be used to take action before these activities start consuming resources on the data server. The following table provides information about predictive elements supported by work classes:

Table 22. Characteristics for predictive identification

Predictive element	Description
Estimated cost	Uses the estimated cost available from the DB2 compiler to include DML within a given timeron range (for example, create a work class for all large queries with an estimated cost over 1 000 000 timerons)
Estimated cardinality	Uses the estimated rows returned (cardinality) from the DB2 compiler to include DML within a given range of rows returned (for example, create a work class for large queries that are estimated to return more than 500 000 rows)

You can also identify activities by using the schema name of the procedure that a CALL statement calls.

Based on workload attributes and work class types, you can identify work and prepare it for the next stage, the management of the work.

For more information on working with work classes and work class sets, refer the following topics:

Work classes and work class sets

A work class is a method of categorizing individual database activities based on attributes of the activities. Work classes are grouped into work class sets, which can be shared by different work action sets.

Examples of database activity attributes which can determine which work class an activity is associated with include: activity type (DDL, DML, LOAD), the estimated cost (where available), the estimated cardinality (where available), and the schema (where available).

Work classes

A work class has the following attributes:

- The work class name, which must be unique in the work class set.
- The database activity attributes, which consist of the following information:
 - The type of database activity that falls into this work class. Using predefined keywords (for example, CALL, READ, WRITE, DML, DDL, LOAD, or ALL), you can classify database requests into different categories. Different types of database activities can be associated with a work class depending on its work type. For example, the WRITE keyword includes updates, deletes, inserts, merges, and selects that contain a delete, insert, or update. For more information, see “Work identification by type of work with work classes” on page 45.
 - The range information that further categorizes DML or XQuery types of database activity:
 - The type of range to specify (either timeron cost or cardinality). Specifying a range of values is optional. For example, when you specify a range for a work class, you can specify that all queries with an estimated cost of less than 100 timerons be processed differently than other queries.
 - The bottom of the range.
 - The top of the range.

- The schema of the routine to be called. Specifying the schema is optional. When defining a work class, you can use the schema attribute to further classify CALL statements according to the schema of the procedure being called. For example, if you specify SCHEMA1 for the schema of a work class and the work type is CALL, all CALL statements calling a SCHEMA1 procedure are classified in that work class. If you specify the schema for a work class type other than CALL or ALL, the error SQL0628N is returned.
- The evaluation order of the work class (or position of the work class in the work class set). For more information, see “Evaluation order of work classes in a work class set” on page 50.
- An automatically generated class identifier that uniquely identifies the work class.

You can create work classes in two ways:

- Create a new work class set to contain the new work class using the WORK CLASS keyword of the CREATE WORK CLASS SET statement.
- Add the new work class to an existing work class set using the ADD keyword of the ALTER WORK CLASS SET statement

You can alter work classes by using the ALTER WORK CLASS keyword of the ALTER WORK CLASS SET statement.

You can drop work classes from a work class set using the DROP WORK CLASS keyword of the ALTER WORK CLASS SET statement, or by using the DROP WORK CLASS SET statement to drop the work class set.

You can view your work classes by querying the SYSCAT.WORKCLASSES view.

Work class sets

You use work class sets to group one or more work classes. A work class set consists of the following attributes:

- A unique descriptive name for the work class set
- Any comments that you want to supply for the work class set
- Zero or more work classes (although a work class can only exist in a work class set, a work class set does not have to contain any work classes)
- An automatically generated ID that uniquely identifies the work class set

You create a new work class set using the CREATE WORK CLASS SET statement. You can create an empty work class set and add work classes later, or you can create a work class set that contains one or more work classes.

You change an existing work class set in the following ways using the ALTER WORK CLASS SET statement:

- Add work classes to the work class set.
- Change work class attributes for work classes in the work class set.
- Drop work classes from the work class set.

You cannot change any work class set attributes.

Drop a work class set using the DROP WORK CLASS SET statement.

You can view your work class sets by querying the SYSCAT.WORKCLASSSETS catalog view.

The following figure shows an example of work classes in a work class set.




Figure 9. Example of work classes and a work class set

For a work class set to be effective on the system, you must define a work action set and associate it with the work class set. By using a work action set, you can associate a work class set to either a service superclass or the database to indicate what action should be applied to the database activities that fall within the classification. If you do not create a work action set for the work class set, the data server ignores the work class set.

Evaluation order of work classes in a work class set

A work class set can have multiple work classes that match with a database activity. To select which work class from a work class set an activity should fall under, the data server goes through the work classes according to the evaluation order, stopping at the first work class that matches the activity.

If no matching work class exists, the database activity does not belong to any work class, and no work action is applied to that activity.

You can affect the evaluation order of work classes in a work class set when you create or alter a work class set. When you create or alter a work class set, you determine the position at which a work class is placed in the work class set using one of the following three methods:

- Specify the absolute position of the work class in the list.
For example, POSITION AT 2. In this situation, the work class is placed in the second position in the work class set, and the work class that was at the second position is now the third, the third work class is now the fourth, and so on. If the position specified for the work class by the CREATE WORK CLASS SET or ALTER WORK CLASS SET statement is greater than the total number of work classes in the work class set, the work class is positioned last in the list.
- Use the POSITION BEFORE or POSITION AFTER keyword to specify the position of the work class relative to work classes already in the work class set.
- Omit the position when creating a work class.
In this situation, the new work class is positioned at the end of the list. The position you specify for the work class in the work class set list is not

necessarily the actual value of the EVALUATIONORDER column in the SYSCAT.WORKCLASSES view. The data server automatically assigns the order value to prevent gaps.

Work classes are processed in the order they are received, which can affect the evaluation order. For example, assume that you issue the following statement:

```
ALTER WORK CLASS SET WCS ALTER WORK CLASS C1 POSITION AT 1
ALTER WORK CLASS C2 POSITION AT 1
```

As a result, the C1 work class has a evaluation order of 2 and the C2 work class has an evaluation order of 1 because C2 was the last work class processed.

Assignment of activities to work classes

If a work class set, through a work action set, is associated with either a database or a service superclass, just prior to execution in processing of an execute, execute immediate, or open request, or just before the execution of the load utility, the database activity is checked to determine if it matches any of the criteria specified in the work classes within the work class set.

The work classes are sorted within the work class set, by their evaluation order. Based on this evaluation order, the database activity is checked against each work class based on the attributes of the database activity (such as the activity type and cardinality) until there is a match or the list of work classes in the work class set has been exhausted.

Assume that the following work classes are in a work class set:

- Evaluation order: 1; work class name: MyLoad; work class type: LOAD
- Evaluation order: 2; work class name: SmallRead; work class type: READ; other attributes: estimated cost < 300 timerons
- Evaluation order: 3; work class name: AllDML; work class type: DML
- Evaluation order: 4; work class name: LargeRead; work class type: READ; other attributes: estimated cost > 301 timerons
- Evaluation order: 5; work class name: MyDDL; work class type: DDL

If a SELECT statement with an estimated cost of 200 timerons is received, it is assigned to the SmallRead work class. If a DDL activity (such as CREATE TABLE) arrives, it is assigned the MyDDL work class. If a SELECT statement with an estimated cost of 500 timerons arrives, it is assigned to the AllDML work class because AllDML is positioned before the LargeRead work class. For more information, see “Example: Working with a work class defined with the ALL keyword” on page 59.

Work classifications supported by thresholds

Although any of the threshold types that can be used in work actions can be associated with any work class, not all types of database activities are supported for all of those threshold types.

For example, if you create a work class for DDL, then associate that work class with an ESTIMATEDSQLCOST threshold work action, that threshold will not apply to any of the requests that are classified under DDL because DDL statements do not have an estimated cost. If you create a work class for ALL, then associate that work class with an ESTIMATEDSQLCOST threshold work action, although all database activities belong to the ALL work class, the threshold will only apply to the database activities that have an estimated cost.

The following tables show which work class categories are supported by which threshold types:

Table 23. Work classification supported by thresholds

	"ACTIVITYTOTALTIME threshold" on page 96	"ESTIMATEDSQLCOST threshold" on page 98	"CONCURRENTDBCOORDACTIVITIES threshold" on page 103	"CPUTIME threshold" on page 96
READ, including SET statements with embedded READ SQL	Yes	Yes	Yes	Yes
WRITE, including SET statements with embedded WRITE SQL	Yes	Yes	Yes	Yes
CALL	Yes	No	No	Yes
DML, including SET statements with embedded READ or WRITE SQL	Yes	Yes	Yes	Yes
DDL	Yes	No	Yes	No
LOAD	Yes	No	Yes	No
ALL	Yes	Some	Yes	Some

Table 24. Work classification supported by thresholds (continued)

	"SQLROWSREAD threshold" on page 99	"SQLROWSRETURNED threshold" on page 101	"SQLTEMPSPACE threshold" on page 102
READ, including SET statements with embedded READ SQL	Yes	Yes	Yes
WRITE, including SET statements with embedded WRITE SQL	Yes	Yes	Yes
CALL	No	No (see note)	No
DML, including SET statements with embedded READ or WRITE SQL	Yes	Yes	Yes
DDL	No	No	No
LOAD	No	No	No
ALL	Some	Some	Some

Note: Although the statements in the procedure called may return rows, because the rows are not returned as a result of the CALL statement they are not controlled by the SQLROWSRETURNED threshold.

Creating a work class

To create a work class, use the CREATE WORK CLASS SET statement or the ALTER WORK CLASS SET statement.

To create a work class, you require WLMADM or DBADM authority.

For additional prerequisites, see the following topics:

- "DDL statements for DB2 workload manager" on page 16
- Naming rules

To create a work class:

1. Create a work class at the same time you create a new work class set or add the new work class to an existing work class set:
 - To create a new work class that is added to a new work class set, use the WORK CLASS keyword of the CREATE WORK CLASS SET statement.
 - To create a new work class that is added to an existing work class set, use the ADD WORK CLASS keyword of the ALTER WORK CLASS SET statement.

Specify one or more of the following properties for the new work class:

- A name for the work class. This name must be unique in the work class set.
- Attributes for the work class. These attributes are used to associate an activity with the work class:
 - The type of work that the work class is to be used for. Use the WORK TYPE parameter to specify this characteristic.

- READ, which represents non-updating SELECT activities, and all XQuery activities.

When you specify the READ keyword, you can also specify an optional for-from-to-clause argument. Use this argument to specify a range for either the cost of the statement in timerons, or its cardinality (that is, the number of rows returned). You must specify a numeric value for the first value. For the second value, you can specify either a numeric value, or the value UNBOUNDED to indicate that you do not want to impose an upper limit on either the cost or cardinality of the activity. You can also specify this argument for the WRITE keyword, the DML keyword, and the ALL keyword.

For example, to associate SELECT activities that have a cost of 5000 timerons or more with this work class, you would specify:

```
WORK TYPE READ FOR TIMERONCOST FROM 5000 TO UNBOUNDED
```

- WRITE, which represents SQL activities that update data in the database.

For example, to associate data writing activities that update between 50 and 100 rows with this work class, you would specify:

```
WORK TYPE WRITE FOR CARDINALITY FROM 50 TO 100
```

- CALL, which represents CALL activities.

When you specify the CALL keyword, you can also specify the ROUTINES IN SCHEMA keyword to indicate that only CALL activities to routines in a specific schema should be associated with this work class. For example, if you only want to associate calls to routines in the ACCOUNTS schema to this work class, you would specify:

```
WORK TYPE CALL ROUTINES IN SCHEMA ACCOUNTS
```

- DML, which represents SQL activities covered by both the READ and WRITE keywords.

For example, to associate all DML activities that have a cost in timerons from 500 to 1000 with this work class, you would specify:

```
WORK TYPE DML FOR TIMERONCOST FROM 500 TO 1000
```

- DDL, which represents the following activities:
 - ALTER
 - CREATE
 - COMMENT
 - DECLARE GLOBAL TEMPORARY TABLE
 - DROP

- FLUSH PACKAGE CACHE
- GRANT
- REFRESH TABLE
- RENAME
- REVOKE
- SET INTEGRITY

For example, to associate all DDL activities with this work class, you would specify:

```
WORK TYPE DDL
```

- LOAD, which represents a LOAD activity.

For example, to associate LOAD activities to this work class, you would specify:

```
WORK TYPE LOAD
```

- ALL, which represents all the work types indicated by all the preceding keywords.

When you specify ALL for a work class type, you can also specify the ROUTINES IN SCHEMA keyword to indicate that only CALL activities to routines in a specific schema should be associated with this work class. You can also specify the for-from-to-clause argument to indicate that all DML activities that have an estimated timeron cost or cardinality specified fall into this class. For example, to associate both DML activities that have a cardinality of 300 to 1500 rows and routines that are called from the NEWHIRES schema to this work class, you would specify the following statement. Because this work class has a type of ALL, it would also apply to other activities that do not have a schema or cardinality, such as LOAD activities and DDL activities.

```
WORK TYPE ALL FOR CARDINALITY FROM 300 TO 1500 ROUTINES
IN SCHEMA NEWHIRES
```

- Optional. The position of the work class in the work class set. The position of the work class in the work class set determines the order in which the work class is evaluated when classifying an activity to a work class. When work class assignment occurs, the data server first determines the work class set associated with the object (either a service superclass or the database), then selects the first matching work class in the work class set that has a work action associated with it. Use the POSITION keyword to specify one of the following:

- LAST. The work class is placed at the end of the list of work classes in the work class set. For example:

```
WORK TYPE ... POSITION LAST
```

- BEFORE *work-class-name*. The work class is to be created in the work class set and positioned before the specified work class. For example:

```
WORK TYPE ... POSITION BEFORE LARGEDDL
```

- AFTER *work-class-name*. The work class is to be created in the work class set and positioned after the specified work class. For example:

```
WORK TYPE ... POSITION AFTER LARGEDDL
```

- AT *integer*. The work class is to be created in the work class set in the position specified by the integer value. For example:

```
WORK TYPE ... POSITION AT 3
```

2. Commit your changes. When you commit your changes, the work class is added to the SYSCAT.WORKCLASSES view.

Altering a work class

If you need to alter a work class, use the ALTER WORK CLASS SET statement.

To alter a work class, you require WLMADM or DBADM authority.

See “DDL statements for DB2 workload manager” on page 16 for additional prerequisites.

To alter a work class:

1. Use the ALTER keyword of the ALTER WORK CLASS SET statement to change one or more of the following properties. See “Creating a work class” on page 52 for an explanation of the supported values for these properties.
 - The FOR keyword. For example, you can change the value specified for the FOR keyword from CARDINALITY to TIMERONCOST.
 - The FROM *from-value* TO *to-value* argument. For example, you can change the argument from FROM 50 TO 100 to FROM 500 TO 1500.
 - The ROUTINES IN SCHEMA or the ROUTINES IN ALL keywords, for CALL activities. For example, if the work class currently does not specify a schema, you can add one. You can also specify the keyword ALL, so that the work class applies to all CALL statements, regardless of the schema of the routine. ALL is the default.
 - The POSITION keyword, followed by the keywords LAST, BEFORE, AFTER, or AT. If you specify POSITION BEFORE or POSITION AFTER, you also need to specify the work class that you want to use to position your altered work class. If you specify POSITION AT, you need to include the position number. For example, you can move a work class from the last position to any position by using the AT keyword, or from any position to the last position by using the LAST keyword.
2. Commit your changes. When you commit your changes, the work class is updated in the SYSCAT.WORKCLASSES view.

Dropping a work class

If you no longer require a work class, you can drop it from the work class set.

To drop a work class, you require WLMADM or DBADM authority.

See “DDL statements for DB2 workload manager” on page 16 for additional prerequisites.

To drop a work class:

1. Use the DROP keyword of the ALTER WORK CLASS SET statement. You cannot drop a work class if any work action in any work action set associated with this work class set has a dependency on the work class you want to drop. In this situation, you must first drop all dependent work actions before dropping the work class.
2. Commit your changes. When you commit your changes, the work class is removed from the SYSCAT.WORKCLASSES view.

Creating a work class set

To create a work class set, use the CREATE WORK CLASS SET statement.

To create a work class set, you require WLMADM or DBADM authority.

For additional prerequisites, see the following topics:

- “DDL statements for DB2 workload manager” on page 16
- Naming rules

To create a work class set:

1. Specify the following properties for the work class set using the CREATE WORK CLASS SET statement:
 - A name for the work class set. The name you specify must be unique in the database.
 - Optional: One or more work classes for the work class set. For more information, see “Creating a work class” on page 52.
2. Commit your changes. When you commit your changes, the work class set is added to the SYSCAT.WORKCLASSETS view.

Altering a work class set

You cannot change the work class set attributes after you create a work class set. However, you can add, alter, and drop work classes in the work class set using the ALTER WORK CLASS SET statement.

To alter a work class set, you require WLMADM or DBADM authority.

For additional prerequisites, see the following topics:

- “DDL statements for DB2 workload manager” on page 16
 - Naming rules
1. If you want to add work class to the work class set, use the ADD keyword. For information about the keywords that you can specify when adding a work class, see “Creating a work class” on page 52.
 2. If you want to alter a work class, use the ALTER keyword. For information about altering a work class, see “Altering a work class” on page 55.
 3. If you want to drop a work class, use the DROP keyword. For information about dropping a work class from a work class set, see “Dropping a work class” on page 55. If you want to drop all the work classes from the work class set, you can drop the work class set itself. For more information, see “Dropping a work class set.”
 4. Commit your changes. When you commit your changes, the SYSCAT.WORKCLASSES view is updated to show any added, altered, or dropped work class.

Dropping a work class set

Use the DROP WORK CLASS SET statement to drop a work class set.

To drop a work class set, you require WLMADM or DBADM authority.

You can only drop a work class set if no work action sets are associated with it. If you want to drop the work class set, you must first drop its dependent work action sets.

To drop a work class set:

1. Use the DROP WORK CLASS SET statement.

- Commit your changes. When you commit your changes the work class set is removed from the SYSCAT.WORKCLASSETS view. In addition, all work classes that were part of the work class set are removed from the SYSCAT.WORKCLASSES view.

Example: Analyzing workloads by activity type

You can use DB2 workload manager table functions to examine the workloads in your environment according to the types of activities being run.

In some situations, you might be interested in the behavior of a certain type of activities, such as LOAD activities. For example, you can observe how many LOAD activities are currently in the system as follows:

```
SELECT COUNT(*)
FROM TABLE(WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97(CAST(NULL AS BIGINT), -2))
AS ACTS
WHERE ACTIVITY_TYPE = 'LOAD'
```

You can obtain a count of how many activities of a specific type have been submitted since the last reset of the DB2 workload manager statistics by using the WLM_GET_WORK_ACTION_SET_STATS table function, as shown in the following example. Assume that the READCLASS and LOADCLASS work classes exist for activities of type READ and activities of type LOAD. The * represents all activities that do not fall into the READCLASS or LOADCLASS work class.

```
SELECT SUBSTR(WORK_ACTION_SET_NAME,1,18) AS WORK_ACTION_SET_NAME,
       SUBSTR(CHAR(DBPARTITIONNUM),1,4) AS PART,
       SUBSTR(WORK_CLASS_NAME,1,15) AS WORK_CLASS_NAME,
       LAST_RESET,
       SUBSTR(CHAR(ACT_TOTAL),1,14) AS TOTAL_ACTS
FROM TABLE(WLM_GET_WORK_ACTION_SET_STATS(' ', -2)) AS WASSTATS
ORDER BY WORK_ACTION_SET_NAME, WORK_CLASS_NAME, PART
```

WORK_ACTION_SET_NAME	PART	WORK_CLASS_NAME	LAST_RESET	TOTAL_ACTS
AdminActionSet	0	ReadClass	2005-11-25-18.52.49.343000	8
AdminActionSet	1	ReadClass	2005-11-25-18.52.50.478000	0
AdminActionSet	0	LoadClass	2005-11-25-18.52.49.343000	2
AdminActionSet	1	LoadClass	2005-11-25-18.52.50.478000	0
AdminActionSet	0	*	2005-11-25-18.52.50.478000	0
AdminActionSet	1	*	2005-11-25-18.52.50.478000	0

You can view the average lifetime of LOAD activities by creating a work action set to map LOAD activities to a specific service subclass. For example, suppose you map LOAD activities to the service subclass LOADSERVICECLASS under the service superclass MYSUPERCLASS. Then, you can query the WLM_GET_SERVICE_SUBCLASS_STATS_V97 table function:

```
SELECT SUBSTR(SERVICE_SUPERCLASS_NAME,1,19) AS SUPERCLASS_NAME,
       SUBSTR(SERVICE_SUBCLASS_NAME,1,18) AS SUBCLASS_NAME,
       SUBSTR(CHAR(DBPARTITIONNUM),1,4) AS PART,
       CAST(COORD_ACT_LIFETIME_AVG / 1000 AS DECIMAL(9,3)) AS AVGLIFETIME
FROM TABLE
(WLM_GET_SERVICE_SUBCLASS_STATS_V97('MYSUPERCLASS', 'LOADSERVICECLASS', -2))
AS SCSTATS
ORDER BY SUPERCLASS_NAME, SUBCLASS_NAME, PART
```

SUPERCLASS_NAME	SUBCLASS_NAME	PART	AVGLIFETIME
SYSDEFAULTUSERCLASS	LOADSERVICECLASS	0	4691.242
SYSDEFAULTUSERCLASS	LOADSERVICECLASS	1	4644.740
SYSDEFAULTUSERCLASS	LOADSERVICECLASS	2	4612.431
SYSDEFAULTUSERCLASS	LOADSERVICECLASS	3	4593.451

Example: Using a work class set to manage specific types of activities

The following example shows how to use a work class set to manage DML activities.

Assume that you have a large number of applications running on your NONAME database each day and lately a few performance issues have been occurring. To deal with some of these issues, you decide that you need to be able to control the number of large queries (that is, any query that has an estimated cost of greater than 9999 timerons or an estimated cardinality of greater than 9999 rows) that can run simultaneously on the database.

To control the number of large queries that can run on the database, you would do the following:

1. Create a MYWORKCLASSET work class set that contains two work classes: one for queries with a large estimated cost and one for queries with a large estimated cardinality. For example:

```
CREATE WORK CLASS SET MYWORKCLASSET
(WORK CLASS LARGEESTIMATEDCOST WORK TYPE DML
FOR TIMERONCOST FROM 10000 TO UNBOUNDED,
WORK CLASS LARGECARDINALITY WORK TYPE DML
FOR CARDINALITY FROM 10000 TO UNBOUNDED)
```

2. Create a DATABASEACTIONS work action set that contains two work actions that are to be applied to the work classes in the MYWORKCLASSET work class set at the database level

```
CREATE WORK ACTION SET DATABASEACTIONS FOR DATABASE
USING WORK CLASS SET LARGEQUERIES
(WORK ACTION ONECONCURRENTQUERY ON WORK CLASS LARGEESTIMATEDCOST
WHEN CONCURRENTDBCOORDACTIVITIES > 1 AND QUEUEDACTIVITIES > 1 STOP EXECUTION,
WORK ACTION TWOCONCURRENTQUERIES ON WORK CLASS LARGECARDINALITY
WHEN CONCURRENTDBCOORDACTIVITIES > 2 AND QUEUEDACTIVITIES > 3 STOP EXECUTION)
```

In addition, several large administrative applications run daily against the database, and you want these applications to run in one resource pool. To accomplish this goal, you would create a service superclass called ADMINAPPS for these applications. For each application, you would create a workload to map it to the ADMINAPPS service superclass.

Because it is important that the queries (SELECT statements) run quickly, you decide to create a service subclass called SELECTS in the ADMINAPPS service superclass for these queries.

To map the SELECT statements to the SELECTS service subclass:

1. Create a SELECTDML work class set that contains a work class for all SELECT statements that do not update the database:

```
CREATE WORK CLASS SET SELECTDML (WORK CLASS SELECTCLASS WORK TYPE READ)
```

2. Create an ADMINAPPSACTIONS work action set. This work action set contains a work action that is to be applied to the work class in work class set SELECTDML at the service superclass level

```
CREATE WORK ACTION SET ADMINAPPSACTIONS FOR SERVICE CLASS ADMINAPPS
USING WORK CLASS SET SELECTDML
(WORK ACTION MAPSELECTS ON WORK CLASS SELECTCLASS MAP ACTIVITY TO SELECTS)
```

Example: Working with a work class defined with the ALL keyword

This example shows how to work with a work class defined as ALL, which potentially covers all recognized activities in the database.

When a work class with the type of ALL is used with a mapping work action, all recognized database activity is mapped to the service subclass specified in the work action. If a work class with the work type of ALL is used with a threshold work action, the threshold type determines which database activities the threshold applies to. Consider the following example.

Assume that you create a work class set called Example with the following work classes. The evaluation order of the work class is as follows:

1. SMALLDML, which is for all DML-type SQL that has an estimated cost of less than 1000 timerons.
2. LOADUTIL, which is for the load utility.
3. ALLACTIVITY, which is for all database activity

ALLACTIVITY is the last work class evaluated, and covers database activities that do not correspond to the first three work classes.

The DDL for creating this work class set is:

```
CREATE WORK CLASS SET EXAMPLE
(WORK CLASS SMALLDML WORK TYPE DML FOR TIMERONCOST FROM 0 TO 999,
WORK CLASS LOADUTIL WORK TYPE LOAD,
WORK CLASS ALLACTIVITY WORK TYPE ALL)
```

Assume that you have a service superclass called EXAMPLESERVICECLASS, and it has two service subclasses called SMALLACTIVITY and OTHERACTIVITY. You want to set up the system so that all small database activities run in the SMALLACTIVITY service subclass, and all other recognized database activities, except for the load utility, run in the OTHERACTIVITY service subclass. You do not want to remap the load utility to any other service subclass, but instead want it to run in the default service subclass.

To accomplish these goals, you would set up a work action set, SERVICECLASSACTIONS for the EXAMPLESERVICECLASS service superclass. The SERVICECLASSACTIONS work action set would contain the following work actions.

Table 25. SERVICECLASSACTIONS work action set

Work action	Work class applied to	Action
MAPDML	SMALLDML	Maps to the SMALLACTIVITY service subclass
COUNTLOAD	LOADUTIL	Counts the number of LOAD activities
MAPOTHER	ALLACTIVITY	Maps to the OTHERACTIVITY service subclass

The DDL to create this work action set is:

```
CREATE WORK ACTION SET SERVICECLASSACTIONS FOR SERVICE CLASS EXAMPLESERVICECLASS
USING WORK CLASS SET EXAMPLE
(WORK ACTION MAPDML ON WORK CLASS SMALLDML MAP ACTIVITY TO SMALLACTIVITY,
WORK ACTION COUNTLOAD ON WORK CLASS LOADUTIL COUNT ACTIVITY,
WORK ACTION MAPOTHER ON WORK CLASS ALLACTIVITY MAP ACTIVITY TO OTHERACTIVITY)
```

Using this configuration, all small DML runs under the SMALLACTIVITY service subclass. The COUNTLOAD work action is applied to the LOADUTIL work class, which runs under the default service subclass. All other recognized database activities run under the OTHERACTIVITY service subclass.

Note: If the ALLACTIVITY work class were at the top of the evaluation order, all recognized activities would be mapped to the OTHERACTIVITY service subclass.

Now assume that you want to define a work action set for the database and apply thresholds that control what is permitted to run concurrently on the system. You could create a work action set called DATABASEACTIONS that contains the following work actions. The DML for creating this work action set is:

```
CREATE WORK ACTION SET DATABASEACTIONS FOR DATABASE USING WORK CLASS SET EXAMPLE
(WORK ACTION CONCURRENTSMALLDML ON WORK CLASS SMALLDML
WHEN CONCURRENTDBCOORDACTIVITIES > 1000 AND QUEUEDACTIVITIES > 10000
COLLECT ACTIVITY DATA STOP EXECUTION,
WORK ACTION CONCURRENTLOAD ON WORK CLASS LOADUTIL
WHEN CONCURRENTDBCOORDACTIVITIES > 2 AND QUEUEDACTIVITIES > 10
COLLECT ACTIVITY DATA STOP EXECUTION,
WORK ACTION CONCURRENTOTHER ON WORK CLASS ALLACTIVITY
WHEN CONCURRENTDBCOORDACTIVITIES > 100 AND QUEUEDACTIVITIES > 100
COLLECT ACTIVITY DATA STOP EXECUTION,
WORK ACTION MAXCOSTALLOWED ON WORK CLASS ALLACTIVITY
WHEN ESTIMATEDSQLCOST > 1000000 COLLECT ACTIVITY DATA STOP EXECUTION)
```

Table 26. DATABASEACTIONS work action set

Work action	Work class applied to	Threshold type and value	Action
CONCURRENTSMALLDML	SMALLDML	Concurrency up to 1000 statements; queue up to 10 000 statements	<ul style="list-style-type: none"> • Stop execution • Collect activity data
CONCURRENTLOAD	LOADUTIL	Concurrency up to 2 occurrences; queue up to 10 occurrences	<ul style="list-style-type: none"> • Stop execution • Collect activity data
CONCURRENTOTHER	ALLACTIVITY	Concurrency up to 100 activities; queue up to 100 activities	<ul style="list-style-type: none"> • Stop execution • Collect activity data
MAXCOSTALLOWED	ALLACTIVITY	Estimated SQL cost up to 1 000 000 timerons	<ul style="list-style-type: none"> • Stop execution • Collect activity data

When these work actions are applied, up to 1000 small DML-type SQL statements (because of the SMALLDML work class) can run at a time, and up to 10 000 of these statements can be queued. Only two occurrences of the load utility can run at a time, and up to 10 occurrences can be queued. Only 100 activities that are not LOAD and are not small DML are permitted to run at a time, and only 100 of these activities can be queued at a time. In all situations, if the queued threshold is violated, the database activity is not permitted to run and an error message is returned.

In addition, the MAXCOSTALLOWED work action is applied to the ALLACTIVITY work class. This means that a database activity with an estimated

cost (that is, DML and XQueries statements) of more than 1 000 000 timerons is not permitted to run. Although the MAXCOSTALLOWED work action is applied to the ALLACTIVITY work class, this work action only affects database activities that have an estimated cost greater than 1 000 000 timerons. This work action does not affect activities that do not have an estimated cost, such as DDL.

Chapter 3. Activities management

Once you have identified the work running on your data server, you are ready to actively manage this work by assigning resources and imposing controls.

Resource assignment with service classes

A service class defines an execution environment in which work can run. This execution environment allocates available resources and can include thresholds that determine how work is permitted to run.

All work runs in a service class and you use workloads to assign work to service superclasses, or you assign work to service subclasses in a service superclass by using workloads, the REMAP ACTIVITY threshold action, or the MAP ACTIVITY work action. When you define a workload, you indicate the service class where work associated with that workload runs. By default, a default user workload also exists (SYSDEFAULTUSERWORKLOAD) that maps work to the default user service class (SYSDEFAULTUSERCLASS), so that any work not explicitly mapped to a user defined service class using a user defined workload will run in the default user service class.

Without service classes, requests cannot be organized into recognizable, logical groupings, as is shown in the following figure.




Figure 10. Unorganized work

You can create different service superclasses to provide the execution environment for different types of work, then assign the applicable requests to the service superclasses. Assume that you have applications from two separate lines of business, finance and inventory. Each line of business would have its own applications to fulfill its responsibilities to the organization. You can organize the requests into categories that make sense for your workload management objectives. In the following figure, different service superclasses are assigned to different lines of business.




Figure 11. Work organized by service classes

In the previous figure, the activities in both service superclasses are further subdivided. The service class provides a two-tier hierarchy: a service superclass and service subclasses underneath. This hierarchy permits for a more complex division of execution environment and better emulates a real-world model. Unless specified otherwise, service subclasses inherit characteristics from the service superclass. Use the service subclasses to further subdivide work in the service superclass.

Prioritization and resource control

When you create or alter a service class object, you can define a number of resource controls:

Table 27. Resource control afforded by service classes

Control	Description
Agent priority	This control sets a processor priority level for the agent threads running in a service class. This priority flows through to the operating system as a relative (delta) priority to other threads and processes running in the data server. Note: This control cannot be set when outbound correlator is in use.
Prefetch priority	This control assigns a priority to the prefetch requests, which affects the order in which they are addressed by the data server.
Buffer pool priority	This control assigns a buffer pool priority to service classes which affects how likely pages fetched by activities in a service class are to be swapped out.

Table 27. Resource control afforded by service classes (continued)

Control	Description
Outbound correlator	<p>This control permits a workload to have some of its resources controlled by a operating system workload manager like AIX Workload Manager or Linux workload management. The tag flows through the agent to the external workload manager and maps to a resource group defined with the manager.</p> <p>When DB2 workload manager is used in conjunction with an operating system workload manager, additional controls are available. With AIX Workload Manager, you can control the amount of processor resource allocated to each service class by setting a minimum, maximum, or relative share of processor resource for each service class. With Linux workload management, you can control the amount of CPU resource by setting shares for each service class relative to the Linux default class.</p> <p>Note: This control cannot be set when agent priority is in use.</p>

Service subclasses

Although the service superclass is the highest tier for work, activities run only in service subclasses. Each service superclass has a default service subclass defined to run activities that you do not assign to an explicitly defined subclass. This default subclass is created when the service superclass is created. You can create additional subclasses in a service class as you require them to further isolate work. Except for histograms and the COLLECT ACTIVITY DATA, COLLECT AGGREGATE ACTIVITY DATA and COLLECT AGGREGATE REQUEST DATA options, a service subclass inherits the attributes of its service superclass, unless otherwise specified. The resources of the superclass are shared by all subclasses in it.

You can define only a single level of subclasses (that is, you cannot define a subclass under another subclass, only under a service superclass).

The following figure is an example of a custom DB2 workload manager configuration using workloads and service classes:

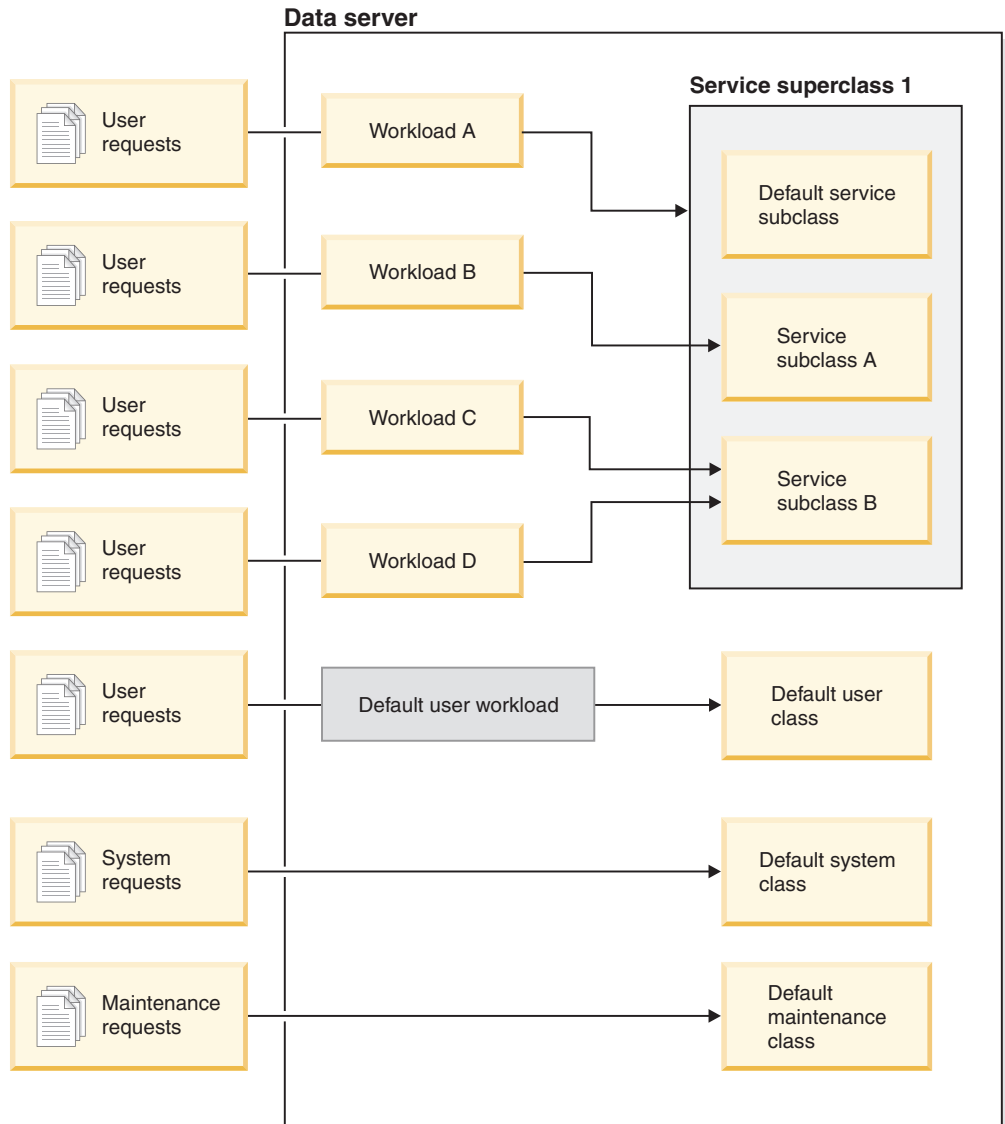


Figure 12. A custom DB2 workload manager configuration using workloads and service classes

As user requests enter the data server, they are identified as belonging to a given workload and assigned to a service superclass or subclass. There are also system requests (for example, prefetches) that run under a special default system service class (SYSDEFAULTSYSTEMCLASS) and DB2-driven maintenance requests (such as an automatic RUNSTATS from the health monitor) that run under a default maintenance service class (SYSDEFAULTMAINTENANCECLASS).

You can view your service classes by querying the SYSCAT.SERVICECLASSES catalog view.

Default service superclasses and subclasses

Each new database or upgraded database has three predefined default service superclasses: the default user class, the default maintenance class, and the default system class.

You cannot disable or drop any of the default service superclasses.

All of the default service superclasses are created with one default service subclass. You cannot create additional service subclasses for the default service superclasses. The default service subclass is always created with the name `SYSDEFAULTSUBCLASS`, as follows:




Figure 13. Two-tier service class hierarchy

All work issued by connections to a default service superclass are processed in the default service subclass of that service superclass.

Default service superclasses and their default service subclasses are dropped only when the database is dropped. They cannot be dropped using the `DROP SERVICE CLASS` statement.

Default user service superclass (SYSDEFAULTUSERCLASS)

By default, all user activities run in the `SYSDEFAULTUSERCLASS`.

Default maintenance service superclass (SYSDEFAULTMAINTENANCECLASS)

The default maintenance service superclass tracks the internal DB2 connections that perform database maintenance and administration tasks. Connections from the DB2 asynchronous background processing (ABP) agents are mapped to this service superclass. ABP agents are internal agents that perform database maintenance tasks. Asynchronous index cleanup (AIC) is an example of an ABP-driven task. ABP agents automatically reduce their resource consumption and number of subagents when the number of user connections increases on the data server. Utilities that are issued by user connections are mapped using regular service classes. You cannot implement service class thresholds on `SYSDEFAULTMAINTENANCECLASS`.

The internal connections tracked by the default maintenance service superclass include:

- ABP connections (including AIC)
- Health monitor initiated backup
- Health monitor initiated RUNSTATS
- Health monitor initiated REORG

Default system service superclass (SYSDEFAULTSYSTEMCLASS)

The default system service superclass tracks internal DB2 connections and threads that perform system-level tasks. You cannot define service subclasses for this service superclass, nor can you associate any workloads or work actions with it. In addition, you cannot implement service class

thresholds on SYSDEFAULTSYSTEMCLASS. The DB2 threads and connections tracked by the default system service superclass include:

- ABP daemon
- Query Patroller (QP) connections
- Self Tuning Memory Manager (STMM)
- Prefetcher engine dispatchable units (EDUs) (db2pfchr)
- Page cleaner EDUs (db2pclnr)
- Log reader EDUs (db2loggr)
- Log writer EDUs (db2loggw)
- Log file reader EDUs (db2lfr)
- Deadlock detector EDUs (db2dlock)
- Event monitors (db2evm)
- Event monitor file writers (db2fw)
- Connections performing system level tasks

A Query Patroller connection is an internal connection to the DB2 data server issued by the QP controller (the server component of QP) when QP is started. This connection is established as QP is starting up, and after QP has successfully started, the connection is mapped to the default system service superclass. Whilst QP is starting up, the connection may temporarily be mapped to another service class as part of the normal workload mapping process. During this period, the connection is subject to all controls and thresholds of the service class it is temporarily mapped to.

Activity-to-service class mapping

All database connections are assigned to a workload at the beginning of the first unit of work. When a workload occurrence is started, all activities running under that workload occurrence are mapped to service classes based on the service class name specified in the workload definition.

The data server assigns a connection to a workload definition if the connection meets the criteria defined for that workload definition. For example, you can set up a DB2 workload manager implementation so that all connections from application A belong to the workload definition Alpha, while all connections from application B belong to the workload definition Beta.

If the workload occurrence is assigned to a service superclass, activities submitted for that workload occurrence can be reassigned to a user-defined service subclass in that service superclass using a work action set.

You can use the workload to map activities from a connection to a service superclass by specifying the SERVICE CLASS keyword of the CREATE WORKLOAD statement. Assuming that no work class or work action applies to the activity, the activity is run in the default service subclass of the service superclass. You can also use a workload to map activities from a connection to a service subclass in the service superclass by specifying the UNDER keyword for the SERVICE CLASS keyword of the CREATE WORKLOAD statement. In this situation, the connection still belongs to the service superclass, but all activities issued from that connection are automatically mapped to the service subclass specified in the workload definition.

After an activity has been mapped to a service subclass and has begun executing, it stays in that service subclass, unless you remap it to another service subclass (with a threshold). Remapping is the process by which you can change the resource assignments for an activity through a different activity-to-service subclass mapping. Both the source and the target service subclasses must exist under the same superclass. After the activity is remapped, it continues to execute in the new service subclass.

Only the coordinator agent does service superclass mapping for the connection. If the coordinator agent spawns subagents, the subagents inherit the superclass mapping of the coordinator agent.

The following figure shows the relationship between connections, workloads, and service superclasses. Connections that meet the definition of workload A are mapped to service superclass 1; connections that meet the definition of workloads B or C are mapped to service superclass 2; connections that meet the definition of workload D are mapped to the SYSDEFAULTUSERCLASS service superclass.

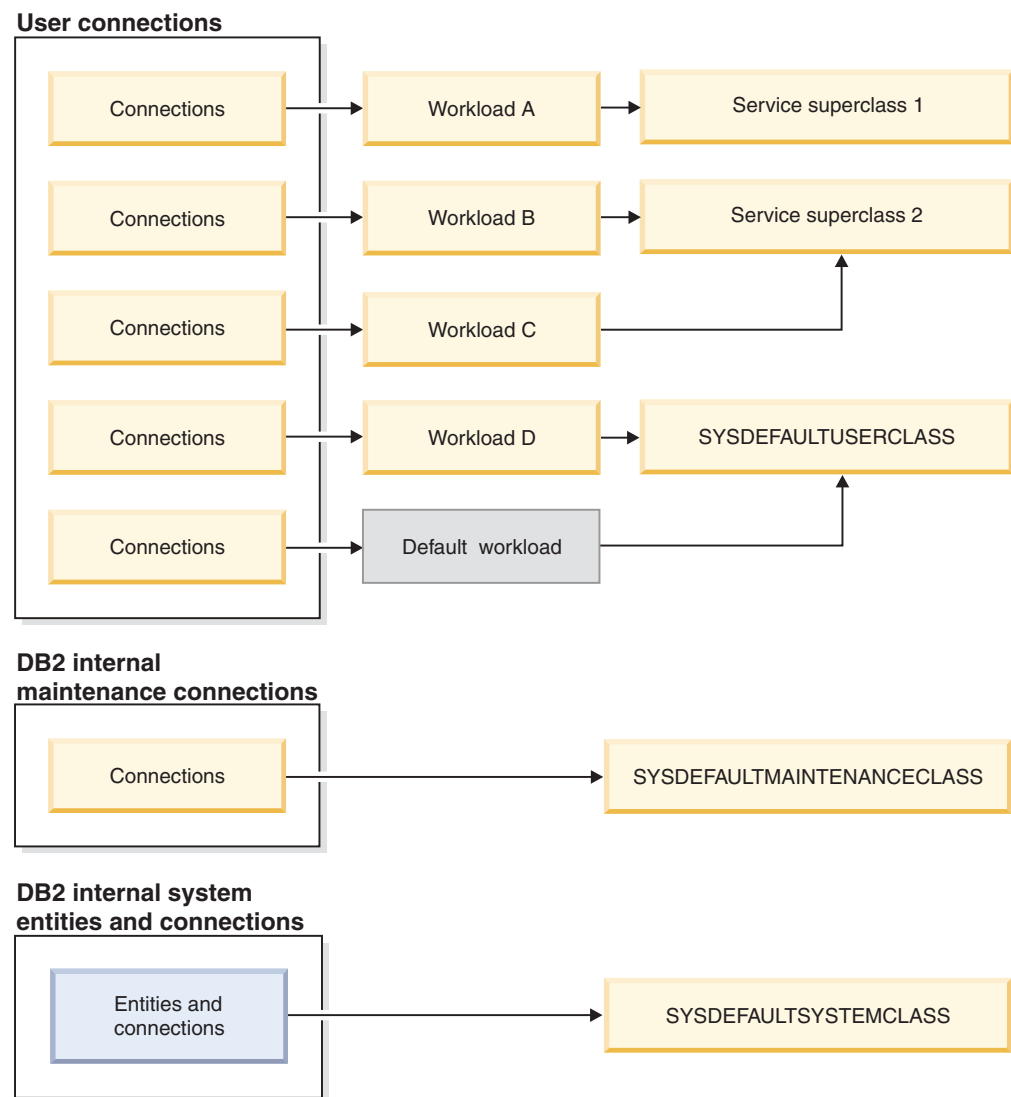


Figure 14. Mapping of database connections to a service superclass

Further differentiating between activities

If you have a more complex DB2 workload manager configuration, you might want to handle activities differently based on either the activity type or some other activity attribute. For example, you might want to do one of the following actions:

- Put DML in a different service subclass than DDL.
- Put all read-type queries with an estimated cost of less than 100 timerons in a different service subclass than all the other read-type queries.

In a more complex configuration you can set up the workload to map activities from the connection to the service superclass. Then, using work actions (contained in a work action set that is applied to the service superclass), you can remap activities, based on their type or attribute, to specific service subclasses in a service superclass.

Specifically, you could apply a work action set that contains a MAP ACTIVITY work action to the service superclass. All activities that are both mapped to the service superclass and match a work class to which a MAP ACTIVITY work action is associated are mapped to the service subclass specified by the work action.

If a workload maps an activity to a service subclass, that activity is not affected by any work action in a work action set that is applied to the service superclass.

- An activity can be mapped to one service subclass in a service superclass by a workload.
- A work action that maps the activity to a different service subclass in the same service superclass also applies to the activity.

If an activity is not mapped to a service subclass through a workload or a work action, the activity is mapped to the default subclass (SYSDEFAULTSUBCLASS) of the service superclass for that activity.

When database activities have been mapped to their respective service superclasses and service subclasses, you can implement controls on all the activities in a particular service class. Statistics are available at the service-class level that you can use to monitor database activities in that service class.

The following figure shows requests to the database being mapped to a service superclass or service subclass through workloads. For information on how work actions are used to map activities to a service subclass, see “Work actions and work action sets” on page 128

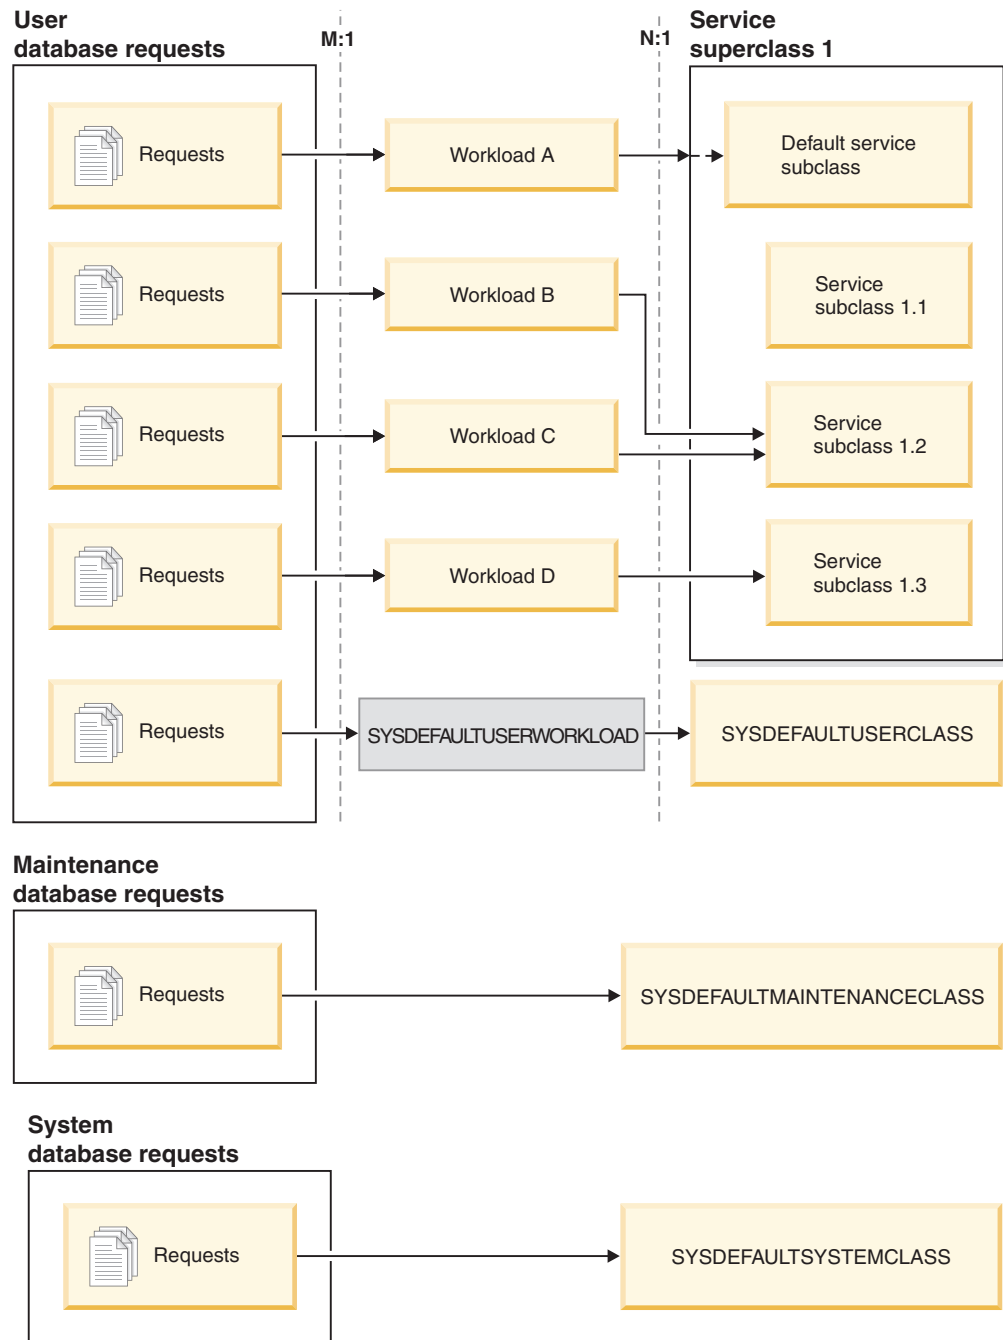


Figure 15. Database connections being mapped to a service superclass

Connections that do not map to a user-defined workload definition are mapped to the default user workload definition, `SYSDEFAULTUSERWORKLOAD`. By default, connections from the default workload definition (`SYSDEFAULTUSERWORKLOAD`) are mapped to the `SYSDEFAULTUSERCLASS` service superclass, which is the default service superclass for user requests. You can alter the `SYSDEFAULTUSERWORKLOAD` workload so that it maps to a different service class. Internal DB2 maintenance connections are mapped to the `SYSDEFAULTMAINTENANCECLASS`, which is the default service superclass for maintenance requests. Internal system entities and connections are mapped to `SYSDEFAULTSYSTEMCLASS`, which is the default service superclass for internal DB2 connections and threads that perform system-level tasks.

Agent priority of service classes

You can associate each DB2 service class with a relative agent priority, which controls processor priority on your data server. This priority is set for all agents that work in a service class and is relative to the agent priority of all other DB2 agents.

If you do not specify the agent priority value for a service class, all agents in that service class have the same priority as all other DB2 agents.

Setting the agent priority for a DB2 service class adjusts the priority of agents only for new work that enters the service class. Non-agent threads running in the service class do not use the agent priority value that you specify. If you are integrating DB2 service classes with an operating system workload manager such as AIX Workload Manager or Linux workload management, you can use the operating system workload manager to specify the processor priority to be used for the operating system class (as processor shares), then have the DB2 service class inherit this value through the OUTBOUND CORRELATOR value of the DB2 service class. The processor priority that you specify using the operating system workload manager controls the priority for agents that run in the DB2 service class, and any service class agent priority setting is ignored.

Important: Do not use the deprecated **agentpri** database manager configuration parameter with DB2 workload manager. You can use this configuration parameter to set the absolute processor priority of all agents in a DB2 instance to a fixed value. However, if you set the absolute priority for an agent by using **agentpri**, you cannot alter the relative priority of the agent by setting the DB2 service class agent priority or by using an operating system workload manager. If you set **agentpri**, the service class agent priority and operating system workload manager have no effect on the priority of agents.

On UNIX operating systems and Linux, valid values are DEFAULT and -20 to 20 (SQLSTATE 42615). Negative values denote a higher relative priority. Positive values denote a lower relative priority.

On Windows operating systems, valid values are DEFAULT and -6 to 6 (SQLSTATE 42615). Negative values denote a lower relative priority. Positive values denote a higher relative priority.

On AIX operating systems, the instance owner must have CAP_NUMA_ATTACH and CAP_PROPAGATE capabilities to set a higher relative priority for agents in a service class using AGENT PRIORITY. To grant these capabilities, logon as root and run the following command:

```
chuser capabilities=CAP_NUMA_ATTACH,CAP_PROPAGATE
```

Prefetch priority of service classes

Prefetchers retrieve data from disk and store this data in buffer pools so that it can be quickly accessed by agents. In DB2 workload manager, each service superclass and subclass can be assigned to have a different prefetch priority.

Agents send read-ahead requests to the database prefetch queue. The prefetchers take these read-ahead requests from the queue, then retrieve the data into the buffer pools. When an agent requires specific data, it first checks the buffer pools to see if the data is available. If not, the agent retrieves the data from disk. Prefetchers perform expensive disk I/O operations, which frees agents to perform computational work in parallel.

Any connection routed to a service class has its prefetch requests processed according to the prefetch priority assigned for the service class. Each service class can be associated with one of the three prefetch priorities: high, medium, or low. You specify the prefetch priority of a service class with the PREFETCH PRIORITY keyword on either the CREATE or ALTER SERVICE CLASS statement.

Specifying DEFAULT for a service superclass sets a medium prefetch priority for the service superclass. You can specify a different prefetch priority for any service subclass in the service superclass, but if you use the default prefetch priority for the service subclass, the service subclass inherits its prefetch priority setting from its service superclass.

High-priority prefetch requests are processed before medium-priority prefetch requests, which, in turn, are processed before low-priority prefetch requests. Prefetch priority affects the order in which prefetch requests are processed, but it does not affect the speed at which they are processed.

Buffer pool priority of service classes

Setting the buffer pool priority of service classes allows you to influence the proportion of pages in the buffer pool that may be occupied by activities in a given service class, which can improve the throughput and performance of activities in that service class.

You can associate each DB2 service class with a relative buffer pool priority, which controls how likely pages fetched into the buffer pool by activities in the service class are to be swapped out. Increasing the buffer pool priority potentially increases the proportion of pages in use by agents of a particular service class.

If you do not specify a buffer pool priority, or if you specify BUFFERPOOL PRIORITY DEFAULT, the buffer pool priority assigned to the service class is DEFAULT. For service superclasses, DEFAULT maps to a value of LOW; for service subclasses, DEFAULT maps to the value of the buffer pool priority of the parent service superclass. All default subclasses have a buffer pool priority of DEFAULT, which cannot be changed.

When upgrading from an earlier DB2 version, the buffer pool priority of existing service classes is set to DEFAULT.

Realizing the benefits of setting buffer pool priority

You are more likely to realize a performance advantage with setting the buffer pool priority for a service class if there is a reasonable amount of contention on the buffer pool. Buffer pool contention demonstrated by an overall hit ratio of 85% or less is likely to see the most benefit. If the overall hit ratio exceeds 90%, there is likely not substantial buffer pool contention to begin with, and setting buffer pool priority will yield less or little benefit in most cases. What benefits you realize are dependent on the type of workload your data server runs.

For some workloads, setting buffer pool priority is more effective if you also turn on proactive page cleaning. This is because buffer pool priority settings are effective only for non-dirty pages and proactive page cleaning is more aggressive about writing out dirty pages to disk. Note that you should turn on proactive page cleaning only if it yields a performance benefit.

If you use asynchronous page cleaning (also known as classic page cleaning), setting the **chnpggs_thresh** database configuration parameter to a lower value will likely yield the same effect of making your buffer pool priority settings more effective, because a low value for this parameter also ensures that there are enough clean pages in the buffer pool.

It is possible that the positive effects of setting buffer pool priority can be surpassed by the effects of prefetching, with or without setting prefetch priority, if there is a reasonable amount of prefetching taking place. For example, if you define a service class with high buffer pool priority where there is only little prefetching, the effective advantage of this buffer pool priority setting might be small when compared to a service class with low buffer pool priority but where activities perform a significant amount of prefetching. Due to the benefits of prefetching, the activities in the service class with low buffer pool priority might even outperform the activities in the high buffer pool priority service class. However, setting buffer pool priority can still supplement your workload management strategy under these circumstances, and you should use it.

States of connections and activities in a service class

Service classes collect connection statistics for each service class. You can see which connections and activities are in a service class, and the state of either the connection or activity.

States of a connection

Following are the possible states of a connection in a service class:

CONNECTED

The connection successfully connected to the database but is not yet associated with its workload and service superclass.

DECOUPLED

The connection does not have a coordinator agent assigned (concentrator case).

DISCONNECTPEND

The connection is disconnecting from the database.

FORCED

The connection has been forced.

INTERRUPTED

The connection has been interrupted.

MAPPED

The connection is mapped to a workload and has joined a service superclass. The connection can now submit activities for execution.

QUEUED

The connection coordinator agent is queued by Query Patroller or a DB2 workload manager queuing threshold. In a database partitioning feature (DPF) environment, this state may indicate that the coordinator agent has made an RPC to the catalog partition to obtain threshold tickets and has not yet received a response.

TRANSIENT

The connection is attempting to join a service class that has reached its connections threshold. The connection is queued to join the service class. When the service class is not violating its connections threshold, the

connection will join the service class. A connection in the transient state cannot submit activities for execution.

TERMINATING

The connection received a connect reset from the client or is being terminated because of a force or an error condition.

UOWEXEC

The connection is processing a request.

UOWWAIT

The connection is waiting for a request from the client.

States of an activity

Following are the possible states of an activity in a service class:

CANCEL_PENDING

If you cancel an activity that has no agent actively working on a request for the activity, the activity is placed in the CANCEL_PENDING state and is cancelled on the next request that is received.

EXECUTING

The activity is executing.

IDLE There is no agent actively processing a request for the activity.

INITIALIZING

The activity was created and is being prepared for execution.

QP_CANCEL_PENDING

The same as the CANCEL_PENDING state, but the activity was cancelled by Query Patroller rather than by the WLM_CANCEL_ACTIVITY procedure.

QP_QUEUED

The activity is queued by Query Patroller.

QUEUED

The activity cannot be executed because of a concurrency threshold at the database or service class level. The activity is queued until it is permitted to execute.

Note: On the AIX operating system, if a queued activity receives SQL4297N, ensure that the DB2 client and data server have the following APAR installed:

- For AIX 5.3, IY89429
- For AIX 5.2, IY89387

TERMINATING

The activity is being terminated.

UNKNOWN

The state of the activity is unknown.

System-level entities not tracked by service classes

Service classes are used for monitoring and controlling objects at the database level. However, not all DB2 entities work directly in a database.

Because service classes work in a database and are stored in the catalog tables of the database, entities that do not work in a database cannot be tracked by service

classes. Instance-level entities, such as the system controller and the health monitor daemons, work at the instance level and are not directly associated with any database. Agents that perform instance attachments and gateway connections are not tracked by service classes either. Because instance attachment agents and gateway agents do not work in a database, they are not tracked by service classes.

The following list is a partial list of entities that do not work within a database and are not tracked by service classes:

- DB2 system controllers (db2sysc)
- IPC listeners (db2ipccm)
- TCP listeners (db2tcpcom)
- FCM daemons (db2fcms, db2fcmr)
- DB2 resynchronization agents (db2resync)
- Idle agents (agents with no database association)
- Instance attachment agents
- Gateway agents
- All other instance-level EDUs

Creating a service class

You create service superclasses and service subclasses under them using the DDL statement `CREATE SERVICE CLASS`.

To create a service class, you require `WLMADM` or `DBADM` authority.

Also see the following topics for other prerequisites:

- “DDL statements for DB2 workload manager” on page 16
- Naming rules

To create a service class:

1. Specify one or more of the following properties for the service class on the `CREATE SERVICE CLASS` statement:
 - Specify the name of the service class:

Note: Once set, the name of a service class cannot be changed.

 - If you are creating a service superclass, the name must be unique among all service superclasses in the database.
When a service superclass is created, its associated default service subclass is automatically created. Only after you have created a service superclass can you create other service subclasses under it.
 - If you are creating a service subclass, the name must be unique among all service subclasses in the service superclass. A service subclass cannot have the same name as its service superclass.
 - If you are creating a service subclass, specify the name of the parent service superclass. After a service subclass is created under a service superclass, it cannot be associated with a different service superclass.
 - Specify the agent priority for the service class. When the agent priority is set to `DEFAULT`, agents in the service class are assigned the same priority that the operating system assigns all DB2 threads. If you set the `AGENT PRIORITY` parameter to a value other than `DEFAULT`, the agent threads are set to a priority equal to the default priority, plus the value set when the next

activity begins. For example, if the default priority is 20 and you set agent priority to -10, the resulting agent priority is set to $20 + (-10) = 10$.

In the SYSCAT.SERVICECLASSES catalog view, an agent priority of DEFAULT is represented as -32768.

On Linux and UNIX, the valid values are -20 to 20 (a negative value indicates a higher relative priority). On Windows-based platforms, the valid values are -6 to 6 (a negative value indicates a lower relative priority)

- Specify the buffer pool priority for the service class which affects how likely pages fetched by activities in the service class are to be swapped out. For service superclasses, the DEFAULT value internally maps to LOW. Service subclasses set to DEFAULT inherit the buffer pool priority from their parent superclasses.
- Specify the prefetch priority. You can specify the priority with which agents in the service class can submit their prefetch requests. Depending on the value specified, the prefetch requests are routed to the high, medium, or low priority prefetch queues. The default prefetch priority is medium.
- Specify the outbound correlator string if you want to associate the DB2 service class with an AIX class or a Linux class. A null value indicates no operating system workload manager association.

If the outbound correlator is set, all threads in the DB2 service class are associated with the operating system workload manager using the outbound correlator when the next activity begins.

If the outbound correlator is set to NONE for a service subclass and the outbound correlator is specified for the associated service superclass, the service subclass inherits the outbound correlator specified for its service superclass.

- Specify the activity data to collect. When activity data collection is enabled, information about an activity is sent from the coordinator partition to the applicable event monitor at the end of the activity. You can write data to the event monitor that includes information about the statement that was run, its compilation environment, and any applicable input data values. You can also specify that no activity data is collected. By default, no activity data is collected.
- Collected aggregate activity information. The aggregate activity information used for the service class only changes after the CREATE SERVICE CLASS operation is committed.
- The type of request metrics to collect for requests submitted by a connection that is associated with the service superclass you specify. By default, basic metrics are always collected for activities associated with the workload.
- The histogram templates that the service subclass should use as templates for its histograms. The histogram templates specified are reflected in the SYSCAT.HISTOGRAMTEMPLATEUSE view. For more information on histograms and histogram templates, see “Histograms in workload management” on page 180.
- Specify whether the service class is enabled or disabled.
 - If a service class is created as enabled (the default), connections and activities can be mapped to the service class. If a service class is created as disabled, new connections and activities mapped to it are rejected.
 - If you create a service superclass as disabled, all service subclasses that you associate with this service superclass behave as though they are disabled, even though they may be displayed as being enabled when you query the SYSCAT.SERVICECLASSES view.

2. Commit your changes. When you commit your changes the service class is added to the SYSCAT.SERVICECLASSES view.

Altering a service class

If you want to change a service class definition, use the ALTER SERVICE CLASS statement.

To alter a service class, you require SQLADM, WLMADM or DBADM authority. To specify any clause other than a COLLECT clause, the authorization ID must include WLMADM or DBADM authority.

See “DDL statements for DB2 workload manager” on page 16 for more information about prerequisites.

Activities that have already acquired resources and are running are not affected by the ALTER statement. These activities will hold their resources and run until completion. However, if a subagent request is sent to a remote database partition during the ALTER SERVICE CLASS operation, the service class definition seen by the coordinator agent and the subagent can differ. Consider the following example in which the prefetch priority for the service class is initially set to MEDIUM:

Table 28. Differences between the views of a coordinator agent and subagent of an altered service class

Event order	Connection 1	Connection 2
1	Coordinating agent sends a request to remote partition (prefetch priority of service class was previously set to MEDIUM)	
2		ALTER SERVICE CLASS issued; set prefetch priority to HIGH
3		COMMIT is issued (the altered service class property is committed at the catalog partition and loaded to memory at all database partitions)
4	Remote subagent receives the request. At this time, the subagent sees the new prefetch priority of HIGH for the service class definition	

This situation described in the previous table is temporary, and only affects connections that issue subagent requests during the ALTER SERVICE CLASS operation. All new connections will see the updated service class definition with the prefetch priority of HIGH.

To alter a service class:

1. Specify one or more of the following properties for the service class on the ALTER SERVICE CLASS statement:
 - Specify whether the service class is enabled or disabled. If you change a service class from enabled to disabled, existing connections or activities

remain with the service class and continue to use previously allocated resources until complete. You can disable a service class if the work coming to the service class is overwhelming the system, or you want to reject all work coming to the service class.

When a service superclass is disabled the following happens:

- a. The service superclass is disabled.
- b. Its service subclasses are disabled.

The service subclasses are only disabled while their service superclass is disabled. When the service superclass is enabled, the service subclasses return to their previous states as defined in the catalog table.

When a service subclass is disabled, its service superclass is not affected, nor other service subclasses associated with the service superclass.

You cannot explicitly disable a default service subclass. To prevent new requests from running under a default service subclass, you must disable the associated service superclass.

- Specify the agent priority for the service class. When the agent priority is set to DEFAULT, agents in the service class are assigned the same priority that the operating system assigns all DB2 threads. If you set the AGENT PRIORITY parameter to a value other than DEFAULT, the agent threads are set to a priority equal to the default priority, plus the value set when the next activity begins. For example, if the default priority is 20 and you set agent priority to -10, the resulting agent priority is set to $20 + (-10) = 10$.

In the SYSCAT.SERVICECLASSES catalog view, an agent priority of DEFAULT is represented as -32768.

On Linux and UNIX, the valid values are -20 to 20 (a negative value indicates a higher relative priority). On Windows-based platforms, the valid values are -6 to 6 (a negative value indicates a lower relative priority)

- Specify the prefetch priority. You can specify the priority with which agents in the service class can submit their prefetch requests. Depending on the value specified, the prefetch requests are routed to the high, medium, or low priority prefetch queues. The default prefetch priority is medium. If the prefetch priority is altered after a prefetch request is submitted, the request will not change its priority.
- Specify the buffer pool priority for the service class which affects how likely pages fetched by activities in the service class are to be swapped out. For service superclasses, the DEFAULT value internally maps to LOW. Service subclasses set to DEFAULT inherit the buffer pool priority from their parent superclasses.
- Specify the outbound correlator string if you want to associate the DB2 service class with an AIX class or a Linux class. A null value indicates no operating system workload manager association.

If the outbound correlator is changed from a non-null value to a null value, all threads in the DB2 service class will disassociate with the operating system workload manager when the next activity begins.

If the outbound correlator is set to NONE for a service subclass and the outbound correlator is specified for the associated service superclass, the service subclass inherits the outbound correlator specified for its service superclass.

If a service superclass uses an outbound correlator, the agent priority of the service superclass must be set to default.

If a service subclass uses an outbound correlator (either explicitly as part of the service subclass definition or implicitly through inheritance from the service superclass), the agent priority of the service subclass must be set to default.

- Specify the activity data to collect. When activity data collection is enabled, information about an activity is sent from the coordinator partition to the applicable event monitor at the end of the activity. You can write data to the event monitor that includes information about the statement that was run, its compilation environment, and any applicable input data values. You can also specify that no activity data is collected. By default, no activity data is collected.
 - Collected aggregate activity information. The aggregate activity information used for the service class only changes after the ALTER SERVICE CLASS operation is committed.
 - The monitoring request metrics collection level for requests submitted by connections mapped to a subclass under the specified service superclass. Note that the effective collection setting for requests running under a service superclass is the combination of both the service class collection level and the **mon_req_metrics** database configuration parameter.
 - Whether to alter the histogram templates used by a service subclass that has enabled aggregate activity data collection using COLLECT AGGREGATE ACTIVITY DATA or aggregate request data collection using COLLECT AGGREGATE REQUEST DATA. Updating the histogram templates used by a service subclass will update the corresponding rows in the SYSCAT.HISTOGRAMTEMPLATEUSE view which displays the histogram templates referenced by a service class or work action. For more information on histograms and histogram templates, see “Histograms in workload management” on page 180.
2. Commit your changes. When you commit your changes the service class is updated in the SYSCAT.SERVICECLASSES view.

Dropping a service class

You drop service classes using the DDL statement DROP SERVICE CLASS.

To drop a service class, you require WLMADM or DBADM authority.

See “DDL statements for DB2 workload manager” on page 16 for more information about prerequisites.

You cannot drop the default service superclasses (SYSDEFAULTUSERCLASS, SYSDEFAULTMAINTENANCECLASS, and SYSDEFAULTSYSTEMCLASS) or their associated service subclasses. The only way to drop the default service superclasses and their associated service subclasses is to drop the database.

A service class you defined cannot be dropped if any of the following conditions apply:

- It is enabled
- It contains user defined service subclasses
- It is referenced by any workload, work action or threshold
- It is still referenced by a workload occurrence
- Any connection or activity is currently mapped to the service class
- If the service class is set as the target of a REMAP ACTIVITY action.

To drop a service class:

1. Disable the service class by using the ALTER SERVICE CLASS statement. If you are dropping a service superclass, this action disables all service subclasses associated with the service superclass. Disabling a service class prevents any additional activities from being associated with it. After disabling the service class, issue a COMMIT statement.

Activities already running under the service class will continue to run. You can list agents that are currently mapped to a service class using the WLM_GET_SERVICE_CLASS_AGENTS_V97 table function. If you do not want these activities to complete, you can use the application identifier returned by the table function and use the FORCE APPLICATION command to force these applications off the database.

2. Use the DROP WORKLOAD statement to drop all workloads associated with the service class. Issue a COMMIT statement after dropping each workload.
3. Drop all applicable work actions that are associated with the service class you want to drop:
 - If you are dropping a service superclass, and a work action set is associated with it, drop that work action set using the DROP WORK ACTION SET statement. Issue a commit statement after dropping the work action set.
 - If you are dropping a service subclass and a work action maps to that service subclass, drop the work action using the DROP WORK ACTION keyword of the ALTER WORK ACTION SET statement. Alternatively, you can also drop the work action set that contains the work action that maps to the service subclass by using the DROP WORK ACTION SET statement. Issue a COMMIT statement after dropping each work action, or after dropping the work action set.
4. Use the DROP THRESHOLD statement to drop all thresholds associated with the service class you want to drop. Issue a COMMIT statement after dropping each threshold.
5. Depending on the object you are dropping, do the following:
 - If you are dropping a service subclass, use the DROP SERVICE CLASS statement to drop the service subclass.
 - If you are dropping a service superclass, use the DROP SERVICE CLASS statement to drop all service subclasses associated with the service superclass, and issue a COMMIT statement after dropping each service subclass. Then issue the DROP SERVICE CLASS statement to drop the service superclass.

Note: You cannot manually drop the default service subclass for the service superclass. The default service subclass for a service superclass is dropped when the service superclass is dropped.

6. Commit your changes. When you commit your changes the service class is removed from the SYSCAT.SERVICECLASSES view.

Example: Using service classes

The following example shows how to use service classes to control database workload.

This example occurs in the fictitious International Beer Emporium. International Beer Emporium is a medium-sized business made up of five major departments: Sales, Accounting, Engineering, Testing and Production. All five departments share the same product catalog database.

Initial implementation of a DB2 workload manager solution

The product catalog database runs well most of the time. However, sometimes users complain that their applications cannot connect to the database because the maximum number of connections has been exceeded. After upgrading to DB2 Version 9.7, Bob, the database administrator, decides to try service classes. Bob wants to know the usage patterns of the product catalog database by each of the five departments and figure out why his database runs out of connections occasionally. Following are the steps Bob follows to set up the service classes:

1. First, Bob creates service superclasses for each of the departments (the default service subclass is also automatically created for each service superclass):
 - SALES is created for the Sales department:

```
CREATE SERVICE CLASS SALES
```
 - ACCOUNTING is created for the Accounting department:

```
CREATE SERVICE CLASS ACCOUNTING
```
 - ENGINEERING is created for the Engineering department:

```
CREATE SERVICE CLASS ENGINEERING
```
 - TESTING is created for the Testing department:

```
CREATE SERVICE CLASS TESTING
```
 - PRODUCTION is created for the Production department:

```
CREATE SERVICE CLASS PRODUCTION
```
2. Bob creates session user groups with appropriate authorization IDs for each of the departments:
 - A session user group is created with the authorization ID SALESGRP. This group includes the authorization IDs of all users in the Sales department.
 - A session user group is created with the authorization ID ACCTNGRP. This group includes the authorization IDs of all users in the Accounting department.
 - A session user group is created with the authorization ID ENGINGRP. This group includes the authorization IDs of all users in the Engineering department.
 - A session user group is created with the authorization ID TESTGRP. This group includes the authorization IDs of all users in the Testing department.
 - A session user group is created with the authorization ID PRODGRP. This group includes the authorization IDs of all users in the Production department.
3. Bob creates workloads to map connections from each group to the associated service class:
 - Workload WL_SALES is created with its session user group set to SALESGRP. WL_SALES maps its connections to the service superclass SALES:

```
CREATE WORKLOAD WL_SALES SESSION_USER GROUP ('SALESGRP')  
SERVICE CLASS SALES
```
 - Workload WL_ACCOUNTING is created with its session user group set to ACCTNGRP. WL_ACCOUNTING maps its connections to the service superclass ACCOUNTING:

```
CREATE WORKLOAD WL_ACCOUNTING SESSION_USER GROUP ('ACCTNGRP')  
SERVICE CLASS ACCOUNTING
```
 - Workload WL_ENGINEERING is created with its session user group set to ENGINGRP. WL_ENGINEERING maps its connections to service class ENGINEERING:


```
CREATE WORKLOAD WL_ENGINEERING SESSION_USER GROUP ('ENGINGRP')
SERVICE CLASS ENGINEERING
```

- Workload WL_TEST is created with its session user group set to TESTGRP. WL_TEST maps its connections to service class TESTING:

```
CREATE WORKLOAD WL_TEST SESSION_USER GROUP ('TESTGRP')
SERVICE CLASS TESTING
```

- Workload WL_PRODUCTION is created with its session user group set to PRODGRP. WL_PRODUCTION maps its connections to service class PRODUCTION:

```
CREATE WORKLOAD WL_PRODUCTION SESSION_USER GROUP ('PRODGRP')
SERVICE CLASS PRODUCTION
```

Bob uses the default service class and workload settings. He wants to observe the database usage patterns before placing any controls on the service classes. The resulting service superclass definitions are as follows:

Table 29. Service class definitions

Service class
SALES
ACCOUNTING
ENGINEERING
TESTING
PRODUCTION
SYSDEFAULTUSERCLASS
SYSDEFAULTMAINTENANCECLASS
SYSDEFAULTSYSTEMCLASS

With a DB2 workload manager solution implemented as described above, work from each department is routed to its own service superclass. Work from departments not specifically accounted for is mapped to the SYSDEFAULTUSERCLASS default service superclass. Using this configuration, Bob can monitor the work in each of the service classes to determine the database usage pattern of the departments.

First refinement of the DB2 workload manager implementation

Following the most recent connection spike, Bob queries service superclass statistics using the WLM_GET_SERVICE_SUPERCLASS_STATS table function and examines the connection high-water mark value for each service superclass. Bob discovers that the connection high-water mark for all departments except Testing is close to 100. However, the statistic for the Testing department shows that at one time, the test team established over 800 connections

Once a month, the Testing department performs its monthly intensive product testing. At this time, the department establishes up to 1000 concurrent connections. Because the database manager configuration parameter **max_connections** is set to 1000, the Testing department uses most of the available connections to the database. When the system has 1000 connections, all subsequent connections are rejected.

Because of memory constraints on the system, the **max_connections** and **maxagents** configuration values cannot be increased on the data server to permit more connections.

To prevent the Testing department from using all the connections, Bob decides to limit the number of connections from the Testing department and ensure that each of the other four departments can obtain sufficient connections to the database to meet their business objectives.

The other four departments ordinarily do not require more than 150 concurrent connections each. In addition, Bob also notices that the default user, default maintenance, and default system service superclasses rarely contain any connections, so he decides that 100 connections should be sufficient for these default service superclasses. After 700 connections (600 for the four departments and 100 for the default classes) are allocated from the **max_connections** pool of 1 000 available connections, 300 connections are available for the Testing department. By limiting the Testing department to a maximum of 300 connections, users from other departments should not have their connection requests rejected.

To limit the Testing group to a maximum of 300 concurrent connections, Bob creates a MAXSERVICECLASSCONNECTIONS threshold of 300 for the TESTING service class.

```
CREATE THRESHOLD MAXSERVICECLASSCONNECTIONS FOR SERVICE CLASS TESTING ACTIVITIES
ENFORCEMENT DATABASE PARTITION
WHEN TOTALSCPARTITIONCONNECTIONS > 300 STOP EXECUTION
```

After implementing this change, the DB2 workload manager configuration is as follows:

Table 30. Configuration after adding threshold for the TESTING service superclass

Service class	MAXSERVICECLASSCONNECTIONS threshold
SALES	N/A
ACCOUNTING	N/A
ENGINEERING	N/A
TESTING	300
PRODUCTION	N/A
SYSDEFAULTUSERCLASS	N/A
SYSDEFAULTMAINTENANCECLASS	N/A

Because the TESTING service class can contain a maximum of only 300 concurrent connections, all connection requests above this threshold are rejected. A MAXSERVICECLASSCONNECTIONS threshold is not applied on the other service classes, so these service classes share the remaining 700 available connections to the data server. Because there is no contention for connections among these service classes, Bob does not place connection thresholds on them.

Second refinement of the DB2 workload manager implementation

Although connections from the Sales, Accounting, Engineering, and Production departments are no longer being rejected, users from these departments still complain about poor performance when the Testing department performs intensive product testing. Bob examines the queries that the Testing department runs during

its product test cycle and discovers that the queries contain complicated joins that involve large amounts of data. These queries generate considerable prefetch activity, which prevents connections from other departments having their prefetch requests processed. Bob decides to lower the prefetch priority of the connections from the Testing department and alters the TESTING service class to set its prefetch priority to LOW:

```
ALTER SERVICE CLASS TESTING PREFETCH PRIORITY LOW
```

The DB2 workload manager configuration is as follows:

Table 31. Configuration after changing prefetch priority for the TESTING service superclass

Service class	MAXSERVICECLASSCONNECTIONS threshold	Prefetch priority
SALES	N/A	DEFAULT
ACCOUNTING	N/A	DEFAULT
ENGINEERING	N/A	DEFAULT
TESTING	300	LOW
PRODUCTION	N/A	DEFAULT
SYSDEFAULTUSERCLASS	N/A	DEFAULT
SYSDEFAULTMAINTENANCECLASS	N/A	DEFAULT

Setting the prefetch priority of the TESTING service class to LOW causes prefetch requests from connections issued from the Testing department to be serviced only after all prefetch requests from the other departments are processed. This change increases the query throughput of the other departments and decreases the throughput of the Testing department during its product testing phase.

Third refinement of the DB2 workload manager implementation

After the prefetch problem is resolved, the Engineering department tells Bob that it needs a few connections for an experimental application called Brewmeister. Because the application is experimental, Bob wants to ensure that it does not consume too many database connections and that queries from the application will not compete for prefetchers when the system is busy. To accomplish these objectives, he creates a new service subclass under the ENGINEERING service superclass for the experimental application and a workload to map connections from the application to the new service subclass. Bob updates the service class and workloads as follows:

- Service subclass EXPERIMENT is created under the service superclass ENGINEERING:

```
CREATE SERVICE CLASS EXPERIMENT UNDER ENGINEERING
```
- Threshold MAXSERVICECLASSCONNECTIONS of 50 is created for the service subclass EXPERIMENT:

```
CREATE THRESHOLD MAXSERVICECLASSCONNECTIONS FOR SERVICE CLASS EXPERIMENT UNDER ENGINEERING ACTIVITIES ENFORCEMENT DATABASE WHEN TOTALDBPARTITIONCONNECTIONS > 50 STOP EXECUTION
```
- Workload WL_EXPERIMENT is created to map connections from the application BREWMEISTER to the service subclass EXPERIMENT:

```
CREATE WORKLOAD WL_EXPERIMENT APPLNAME ('BREWMEISTER') SERVICE CLASS EXPERIMENT UNDER ENGINEERING
```
- The prefetch priority for the EXPERIMENT service subclass is set to LOW:

```
ALTER SERVICE CLASS EXPERIMENT UNDER ENGINEERING PREFETCH PRIORITY LOW
```

The DB2 workload manager configuration is as follows:

Table 32. Configuration with EXPERIMENT service subclass

Service class	MAXSERVICECLASSCONNECTIONS threshold	Prefetch priority
SALES	N/A	DEFAULT
ACCOUNTING	N/A	DEFAULT
ENGINEERING	N/A	DEFAULT
EXPERIMENT	50	LOW
TESTING	300	LOW
PRODUCTION	N/A	DEFAULT
SYSDEFAULTUSERCLASS	N/A	DEFAULT
SYSDEFAULTMAINTENANCECLASS	N/A	DEFAULT

With this configuration, the BREWMEISTER application can only maintain 50 concurrent connections to the database. In addition, prefetch requests from this application are sent to the low priority prefetch queue. The Engineering department can now safely experiment with the application, knowing that it cannot accidentally overwhelm the database system.

Example: Analyzing a service class–related system slowdown

If you notice a system slowdown (for example, some applications take much longer than expected to complete) and are unsure whether the problem is related to the configuration of the service classes, you can use table function data to investigate and, if necessary, correct the problem.

First, obtain a high-level overview of what is occurring in the service classes. This high-level overview should include the average activity lifetime, the number of activities that completed normally rather than abnormally, and the high watermark for concurrent coordinator activities in the system. To obtain this information, you can create a general query with aggregation across service classes and database partitions by using the data obtained from the table function WLM_GET_SERVICE_SUBCLASS_STATS_V97. Set the first and second arguments to empty strings and the third argument to -2 (a wildcard character) to indicate that data is to be gathered for all service classes on all database partitions. Your query might resemble the following one:

```
SELECT SUBSTR(SERVICE_SUPERCLASS_NAME,1,19) AS SUPERCLASS_NAME,
       SUBSTR(SERVICE_SUBCLASS_NAME,1,18) AS SUBCLASS_NAME,
       SUBSTR(CHAR(SUM(COORD_ACT_COMPLETED_TOTAL)),1,13) AS ACTSCOMPLETED,
       SUBSTR(CHAR(SUM(COORD_ACT_ABORTED_TOTAL)),1,11) AS ACTSABORTED,
       SUBSTR(CHAR(MAX(CONCURRENT_ACT_TOP)),1,6) AS ACTSHW,
       CAST(CASE WHEN SUM(COORD_ACT_COMPLETED_TOTAL) = 0 THEN 0
                ELSE SUM(COORD_ACT_COMPLETED_TOTAL * COORD_ACT_LIFETIME_AVG)
                / SUM(COORD_ACT_COMPLETED_TOTAL) END / 1000 AS DECIMAL(9,3))
       AS ACTAVGLIFETIME
FROM TABLE(WLM_GET_SERVICE_SUBCLASS_STATS_V97 ('', '', -2)) AS SCSTATS
GROUP BY SERVICE_SUPERCLASS_NAME, SERVICE_SUBCLASS_NAME
ORDER BY SERVICE_SUPERCLASS_NAME, SERVICE_SUBCLASS_NAME
```

Assume that on previous occasions, the query reported the following results:

SUPERCLASS_NAME	SUBCLASS_NAME	ACTSCOMPLETED	ACTSABORTED	ACTSHW	ACTAVGLIFETIME
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	8	0	1	3.750
BI_APPS	SYSDEFAULTSUBCLASS	4	0	1	5.230
BATCH	SYSDEFAULTSUBCLASS	1	0	1	25.600

The data returned by this query might be sufficient to show that the slowdown is occurring in the BL_APPS service class because its average activity lifetime is significantly higher than usual. This situation could indicate that the available resources for that particular service class are becoming exhausted.

If the averages for the service classes for all database partitions do not isolate the problem, consider analyzing average values for each database partition. Aggregating the average for each database partition into a global average can hide large discrepancies between database partitions. In this situation, the assumption is that every database partition is being used as a coordinator partition. If this assumption is incorrect, the average lifetime computed at non-coordinator partitions is zero.

```
SELECT SUBSTR(SERVICE_SUPERCLASS_NAME,1,19) AS SUPERCLASS_NAME,
       SUBSTR(SERVICE_SUBCLASS_NAME,1,18) AS SUBCLASS_NAME,
       SUBSTR(CHAR(DBPARTITIONNUM),1,4) AS PART,
       CAST(COORD_ACT_LIFETIME_AVG / 1000 AS DECIMAL(9,3)) AS AVGLIFETIME
FROM TABLE(WLM_GET_SERVICE_SUBCLASS_STATS_V97(' ', ' ', -2)) AS SCSTATS
ORDER BY SUPERCLASS_NAME, SUBCLASS_NAME
```

SUPERCLASS_NAME	SUBCLASS_NAME	PART	AVGLIFETIME
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	0	3.425
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	1	2.752
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	2	8.230
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	3	0.593

In this example, database partition 2 might be receiving more work than usual because its average activity lifetimes are much higher than those of the other database partitions.

Many different situations can cause a system slowdown. Use the following principles to make the best use of the information provided by the DB2 workload manager table functions:

- Address large numbers of locking conflicts at the level of the application logic and environment (isolation level and so on).
- If the service class is running close to its threshold levels (the number of concurrent requests and so on), you might need to increase the thresholds.
- If the resources allotted to a service class are becoming exhausted and `OUTBOUT CORRELATOR` is set, the mapping to the operating system service classes might be the cause of the problem (that is, the operating system service class corresponding to the service class is not getting enough processor resources).
- Higher numbers of activities than expected might be running in the service class, which might be consuming more resources than normal. Check the number of completed activities to determine whether the amount of work being done in the service class is reasonable.
- Activities might be spending more time in queues if more activities are being submitted than expected and concurrency thresholds are defined. Check whether the average queue time for activities has increased by the same amount as the average lifetime. If they have increased by the same amount, the queues are behaving as expected; however, if the lifetime is unacceptable, consider allocating more resources to the service class and reducing the concurrency threshold.

Example: Investigating agent usage by service class

DB2 workload manager provides the WLM_GET_SERVICE_CLASS_AGENTS_V97 table function, which you can use to determine the relative distribution of agents among service classes.

Situations can arise in which a data server resource, such as an agent, is overutilized by a group of users or an application. For example, assume that a group of users is using almost all of the available agents and that a user from outside this group voices a concern about that to you.

The first step to take is to determine how many agents are working for each service class. You might use a query such as the following one:

```
SELECT SUBSTR(AGENTS.SERVICE_SUPERCLASS_NAME,1,19) AS SUPERCLASS_NAME,  
       SUBSTR(AGENTS.SERVICE_SUBCLASS_NAME,1,19) AS SUBCLASS_NAME,  
       COUNT(*) AS AGENT_COUNT  
FROM TABLE(WLM_GET_SERVICE_CLASS_AGENTS_V97(' ', ' ', CAST(NULL AS BIGINT), -2))  
AS AGENTS  
WHERE AGENT_STATE = 'ACTIVE'  
GROUP BY SERVICE_SUPERCLASS_NAME, SERVICE_SUBCLASS_NAME  
ORDER BY SERVICE_SUPERCLASS_NAME, SERVICE_SUBCLASS_NAME
```

SUPERCLASS_NAME	SUBCLASS_NAME	AGENT_COUNT
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	7
TEST	SYSDEFAULTSUBCLASS	20

If you conclude that a particular service class is using more than its fair share of agents, you can take actions to restrict the number of activities permitted for a workload or a service class. Alternatively, you can restrict the number of connections for a service class.

Control of work with thresholds

Thresholds permit you to maintain stability in the system. You create threshold objects in order to catch work that behaves abnormally, either predictively before the work begins running based on the projected impact, or reactively as it is running and consuming resources.

An example of work that can be controlled with thresholds is a query that consumes large amounts of processor time at the expense of all other work running on the system. Such a query can be controlled either before it even begins executing, based on estimated cost, or after it has begun executing and is consuming more than the permitted amount of resources.

Types of thresholds

Connection thresholds

If you want to limit how many database connections can be open at any one time, or how long they can sit idle, use a connection threshold. These thresholds limit the total number of concurrent connections to your database, and they can be used to detect connections that sit idle for too long.

Table 33. Connection thresholds

Threshold	Description
CONNECTIONIDLETIME	Controls the amount of time that a connection sits idle and is not working on behalf of user requests. Use this threshold to detect inefficient use of data server resources and application wait conditions.

Activity thresholds

If you want to limit the impact that specific activities can have on how the data server is running, activity thresholds provide you with one of the means you can use. Excess execution time, abnormally high volumes of data returned, or abnormally high amounts of resources consumed are all examples of warning flags that potentially troublesome activities could be consuming excessive resources, which you can control with activity thresholds.

Table 34. Activity thresholds

Threshold	Description
ACTIVITYTOTALTIME	Controls the amount of time that any given activity can spend from submission to completion, for both execution and queue time. Use this threshold to detect jobs that are taking an abnormally long time to complete.
CPUTIME	Controls the maximum amount of combined user and system processor time that an activity may consume on a particular database partition during the execution of the activity. Use this threshold to detect and control activities which are consuming excessive processor resources.
CPUTIMEINSC	Controls the maximum amount of combined user and system processor time that an activity may consume on a particular database partition while executing in a specific service subclass. Use this threshold to detect and control activities which are consuming excessive processor resources within the current service class.
ESTIMATEDSQLCOST	Controls DML activities that the query optimizer determines to have a large estimated cost. Use this threshold to predict potentially resource-heavy SQL before it starts executing on the system and identifying poorly written SQL.
SQLROWSREAD	Controls the maximum number of rows which can be read on any database partition by an activity. Use this threshold to detect and control activities which are reading an excessive number of rows.
SQLROWSREADINSC	Controls the maximum number of rows which can be read by an activity on a particular database partition while executing in a specific service subclass. Use this threshold to detect and control activities which are reading an excessive number of rows within the current service class.
SQLROWSRETURNED	Controls the number of rows returned when executing SQL. Use this threshold to identify when the amount of data exceeds a reasonable volume.
SQLTEMPSPACE	Controls the amount of temporary table space a given activity can consume on a partition. Use this threshold to prevent certain SQL statements from using up a disproportionate amount of temporary space, impeding the progress of other work.

The data server considers requests from utilities such as REORGCHK, IMPORT, and EXPORT to be user logic, and thus these requests are subject to any defined thresholds.

Aggregate thresholds

If you want to control the collective impact that certain activities can have on your data server, you define aggregate thresholds. Aggregate thresholds often, but not always, enforce concurrency control in cases where you need

to limit the number of certain activities running at the same time. Some aggregate thresholds have a built-in queue and are known as queuing thresholds.

Table 35. Aggregate thresholds

Threshold	Description
AGGSQLTEMPSPACE	Controls the maximum amount of system temporary table space that can be consumed in total across all activities in the service subclass. Use this threshold to detect and control activities that belong to a service subclass whose activities are consuming too much system temporary table space in aggregate across all of the activities in the service subclass.
CONCURRENTWORKLOADOCCURRENCES	Controls the number of active occurrences of a workload that can run on a coordinator partition at the same time. Use to control the spread of connections from a specific source.
CONCURRENTWORKLOADACTIVITIES	Controls the number of individual activities that can run within a workload occurrence. Use to limit work within an individual workload occurrence.
CONCURRENTDBCOORDACTIVITIES	Controls the number of concurrent activities in the domain that the threshold is associated with (database, work action, service superclass, or service subclass).
TOTALDBPARTITIONCONNECTIONS	Controls the number of database connections to a given partition that can be established at the same time. Use to prevent a given partition from becoming overloaded.
TOTALSCPARTITIONCONNECTIONS	Controls the number of database connections to a given partition for work executing within a given service class at the same time. Similar to the total database partition connections but more granular because the connection is linked to a service class.

For those aggregate thresholds that support it, concurrency control is provided through a system of execution 'tickets.' Each incoming activity must claim a ticket from the applicable concurrency threshold before it can begin executing. Once all tickets are consumed, additional activities are queued until a ticket becomes available or an error is returned, depending on how you defined the threshold. If the concurrency threshold has queuing enabled, then a ticket passes from an activity that has finished execution to another activity that is in the queue. This activity can then leave the queue and begin execution. How many tickets are available per concurrency threshold depends on how you defined the threshold. For example, if you defined a CONCURRENTDBCOORDACTIVITIES threshold to limit the number of concurrently running database activities to 10, then there are 10 execution tickets available.

For stored procedures, both activity and concurrency thresholds are applied to the stored procedure itself and its child activities. If the execution of the stored procedure is queued, none of its child activities can proceed. However, when a stored procedure starts running, child activities might be queued.

Taking action when thresholds are violated

The action that is taken dynamically when a threshold is violated depends on how you define the threshold.

Stop execution (STOP EXECUTION)

A common action when a threshold is violated is to stop the activity from

executing. In this case, an error code is returned to the submitting application indicating that the threshold was violated. Note that for TOTALDBPARTITIONCONNECTIONS and TOTALSCPARTITIONCONNECTIONS thresholds, a STOP EXECUTION action prevents a connection from being established. For CONNECTIONIDLETIME thresholds, the connection is closed. For CONCURRENTWORKLOADOCCURRENCES, a new workload occurrence is prevented from being created. For all activity-related thresholds, the activity is stopped from continuing to execute. If a THRESHOLDVIOLATIONS event monitor is active, a record is written to the event monitor indicating that the threshold was violated.

Continue execution (CONTINUE)

In some situations, stopping the execution of an activity is too harsh a response. A preferable response is to permit the activity to continue to run and to collect the relevant data for an administrator to perform future analysis to determine how to prevent this condition from happening again. In this situation, no error code is returned to the submitting application. If the action is to continue, the user receives no indication that the threshold was violated. If a THRESHOLDVIOLATIONS event monitor is active, a record is written to the event monitor. If a CONTINUE threshold action is specified for a queuing threshold, this effectively renders the size of the queue unbounded, regardless of any hard value you include.

Remap the activity (REMAP ACTIVITY TO)

When an activity violates a certain limit, you may simply wish to assign different resource controls to it but to let the activity continue executing otherwise. Such a response permits you to dynamically raise or lower the amount of resources an activity can consume throughout its lifetime. In this case, an already executing activity is permitted to continue with no indication to the user or application that the threshold was violated, although the activity now runs with different resources available to it. Remapping is available with any of the in-service-class thresholds like CPUTIMEINSC and SQLROWSREADINSC.

Collect data (COLLECT ACTIVITY DATA)

When some thresholds are violated, data is collected. By default the fact that an activity threshold was violated is recorded in an activated threshold violations event monitor. If you want more detailed information about the activity that violated the threshold, you can request that information for the activity be written to the active event monitor for activities when the activity completes execution using the COLLECT ACTIVITY DATA clause

Threshold domain and enforcement scope

Each threshold operates on a domain, which is a database object. Only activities taking place in the domain of a threshold can be affected by that threshold.

The following threshold domains exist:

- Database
- Service superclass
- Service subclass
- Work action
- Workload

Within each of these threshold domains, a threshold has a scope over which it is enforceable, such as a single workload occurrence, a database partition, or all the partitions of a database. This is the enforcement scope of the threshold. For example: Service class aggregate thresholds can have one of two enforcement scopes, database and database partition; an example of an aggregate threshold that applies only at the database partition level is the maximum number of concurrent connections for a service superclass on a partition (TOTALSCPARTITIONCONNECTIONS). Similarly, the following table shows that you can specify the processor time threshold (CPUTIME) at the database, superclass, subclass, work action or workload domain and that it is enforced per partition. That is, the upper boundary specifies the maximum amount of user and system processor time per partition that an activity may use.

Table 36. Threshold domains and enforcement scopes

Threshold domain	Thresholds with database enforcement scope	Thresholds with database partition enforcement scope	Thresholds with workload occurrence enforcement scope
Database	<ul style="list-style-type: none"> • “ACTIVITYTOTALTIME threshold” on page 96 • “CONCURRENTDBCOORDACTIVITIES threshold” on page 103 • “CONNECTIONIDLETIME threshold” on page 94 • “ESTIMATEDSQLCOST threshold” on page 98 • “SQLROWSRETURNED threshold” on page 101 	<ul style="list-style-type: none"> • “CPUTIME threshold” on page 96 • “SQLROWSREAD threshold” on page 99 • “SQLTEMPSPACE threshold” on page 102 • “TOTALDBPARTITIONCONNECTIONS threshold” on page 107 	N/A
Work action	<ul style="list-style-type: none"> • “ACTIVITYTOTALTIME threshold” on page 96 • “CONCURRENTDBCOORDACTIVITIES threshold” on page 103 • “ESTIMATEDSQLCOST threshold” on page 98 • “SQLROWSRETURNED threshold” on page 101 	<ul style="list-style-type: none"> • “CPUTIME threshold” on page 96 • “SQLROWSREAD threshold” on page 99 • “SQLTEMPSPACE threshold” on page 102 	N/A
Service superclass	<ul style="list-style-type: none"> • “ACTIVITYTOTALTIME threshold” on page 96 • “CONCURRENTDBCOORDACTIVITIES threshold” on page 103 • “CONNECTIONIDLETIME threshold” on page 94 • “ESTIMATEDSQLCOST threshold” on page 98 • “SQLROWSRETURNED threshold” on page 101 	<ul style="list-style-type: none"> • “CPUTIME threshold” on page 96 • “SQLROWSREAD threshold” on page 99 • “SQLTEMPSPACE threshold” on page 102 • “TOTALSCPARTITIONCONNECTIONS threshold” on page 108 	N/A
Service subclass	<ul style="list-style-type: none"> • “ACTIVITYTOTALTIME threshold” on page 96 • “CONCURRENTDBCOORDACTIVITIES threshold” on page 103 • “ESTIMATEDSQLCOST threshold” on page 98 • “SQLROWSRETURNED threshold” on page 101 	<ul style="list-style-type: none"> • “AGGSQLEMPSPACE threshold” on page 103 • “CPUTIME threshold” on page 96 • “CPUTIMEINSC threshold” on page 97 • “SQLROWSREAD threshold” on page 99 • “SQLROWSREADINSC threshold” on page 100 • “SQLTEMPSPACE threshold” on page 102 	N/A
Workload	<ul style="list-style-type: none"> • “ACTIVITYTOTALTIME threshold” on page 96 • “ESTIMATEDSQLCOST threshold” on page 98 • “SQLROWSRETURNED threshold” on page 101 	<ul style="list-style-type: none"> • “CONCURRENTWORKLOADOCCURRENCES threshold” on page 106 • “CPUTIME threshold” on page 96 • “SQLROWSREAD threshold” on page 99 • “SQLTEMPSPACE threshold” on page 102 	<ul style="list-style-type: none"> • “CONCURRENTWORKLOADACTIVITIES threshold” on page 105

Threshold evaluation order

Thresholds are evaluated in a specific order when you define them on a database.

The following thresholds are evaluated before all others:

- TOTALDBPARTITIONCONNECTIONS. This threshold is evaluated when a new connection is made to a database.
- CONCURRENTWORKLOADOCCURRENCES. This threshold is evaluated when a new workload occurrence is started for a workload definition that has this threshold applied to it.
- TOTALSCPARTITIONCONNECTIONS. This threshold is evaluated when a connection joins a service class (either a new connection or a transfer between service classes as a result of workload reassignment).

All other thresholds are based on recognized activities resulting from an SQL statement or the execution of a utility as the load utility and are evaluated in the following order:

1. "Predictive thresholds"
2. "Reactive thresholds" on page 94

Predictive thresholds

Predictive thresholds are checked before reactive thresholds, because they affect whether or not a database activity can start to run.

The sequence in which predictive thresholds are evaluated is as follows. If you do not define a particular threshold, its step is skipped. Also, the steps described might be combined at run time for performance reasons.

1. Check if a `CONCURRENTWORKLOADACTIVITIES` threshold exists and if so, whether it has been violated. If the threshold is violated, the corresponding action is taken. If applicable, move to the next step.
2. Check if an `ESTIMATEDSQLCOST` threshold exists and if so, whether it has been violated. If you define this threshold in more than one domain, the threshold is resolved according to the scope resolution rules (see "Activity threshold scope resolution" on page 95 for more information). The result of this operation is one value of `ESTIMATEDSQLCOST` applicable to the activity. If the threshold is violated, the corresponding action is taken.
3. Check if a `CONCURRENTDBCOORDACTIVITIES` threshold exists for the database work action set threshold domain and if so, whether it has been violated. If the threshold is violated, the corresponding action is taken.
4. Check if a `CONCURRENTDBCOORDACTIVITIES` threshold exists for the service subclass threshold domain and if so, whether it has been violated. If the threshold is violated, the corresponding action is taken.
5. Check if a `CONCURRENTDBCOORDACTIVITIES` threshold exists for the service superclass threshold domain and if so, whether it has been violated. If the threshold is violated, the corresponding action is taken.
6. Check if a `CONCURRENTDBCOORDACTIVITIES` threshold exists for the database threshold domain and if so, whether it has been violated. If the threshold is violated, the corresponding action is taken.

Concurrency threshold considerations: The evaluation order for concurrency thresholds does not follow the hierarchy used for resolving activity thresholds. An activity must pass through each defined concurrency threshold before it is permitted to execute.

For concurrency thresholds, thresholds for database-level work action sets are checked first in order to avoid work action set thresholds on particular types of work blocking work of other types, which would affect concurrency. By checking work action set concurrency thresholds first, situations like the following example are avoided.

Assume that the following thresholds are defined:

- A work action concurrency threshold for `LOAD` activities is defined with a value of 1.
- The service superclass `S1` concurrency limit is set to 10.

Also, assume that one LOAD activity is already running in the database (under any service superclass) and nine activities are already running in service superclass S1. A second new LOAD activity enters as the 10th activity. If the activity threshold scope resolution hierarchy were used during threshold evaluation, the incoming LOAD activity would not violate the service class threshold, increasing the concurrency to 10. The LOAD activity is then evaluated against the work action threshold concurrency limit, which is violated because a LOAD activity is already running in the database and the work action threshold concurrency value is only 1. The second LOAD activity is then queued.

Any new activity arriving into service superclass S1 is now queued (because the service class concurrency limit is already reached). The work action threshold queue is affecting the service class, which is undesirable because activities trying to run in the service class do not necessarily have a relationship with the work action threshold condition (for example, an insert operation trying to run in service superclass S1 should not have to wait on a LOAD activity that is queued because of a work action threshold condition). Therefore, to avoid this type of situation, the work action concurrency threshold is checked first. Because the concurrency threshold is checked first, the 10th activity in the service class (which happens to be a LOAD activity) is blocked at the work action threshold level before it can attempt to consume one spot in the service superclass S1.

Reactive thresholds

Reactive thresholds are evaluated in a discrete fashion when an activity is executing and no specific order is used to evaluate reactive thresholds. The following reactive thresholds are available to you:

- ACTIVITYTOTALTIME
- AGGSQLTEMPSPACE
- CONNECTIONIDLETIME
- CPUTIME
- CPUTIMEINSC
- SQLTEMPSPACE
- SQLROWSREAD
- SQLROWSREADINSC
- SQLROWSRETURNED

Connection thresholds

A connection threshold applies controls to individual database connections. You can use connection thresholds to limit the total number of concurrent connections to the database and how long a connection can sit idle.

CONNECTIONIDLETIME threshold

The CONNECTIONIDLETIME threshold specifies a maximum amount of time that a connection can be idle (that is, not working on a user request).

Type Connection

Definition domain

Database or service superclass

Enforcement scope

Database

Tracked work

User connections

Queuing

No

Unit Time duration expressed in minutes, hours, or days

Predictive or reactive

Reactive

If a connection remains idle for longer than the duration specified by the threshold and the threshold action is STOP EXECUTION, the connection is closed.

Activity thresholds

An activity threshold applies to an individual activity. When the resource usage of an individual activity violates the upper bound of the threshold that is tracking it, the corresponding action is triggered and applied once to the activity.

After being applied once, the threshold is deactivated for the activity and not applied again.

For example: Assume that you defined a time based threshold that triggers a CONTINUE action after an elapsed time of 5 minutes. If an activity violates this threshold, the action is applied once but not reapplied every 5 minutes.

Activity threshold scope resolution

Activity thresholds apply to individual activities and when multiple thresholds apply to the same executing activity, a decision must be made about which threshold to enforce.

Aggregate thresholds are not affected, because the same activity is permitted to contribute to multiple activity aggregates simultaneously, as occurs with concurrency thresholds, for example.

The resolution about which activity threshold to apply to an executing activity follows the rule that a value defined in a local domain overrides any value from a wider or more global domain. Following is the hierarchy of domains, from the most local to the most global:

- Workload
- Service subclass
- Service superclass
- Work action
- Database

For example: A threshold that defines a maximum execution time of 1 hour for all database queries defined in the database domain is overridden by a threshold that defines a maximum execution time of 5 hours for a service superclass set up to handle large queries, which is overridden by a maximum execution time of 10 hours for a service subclass for very large queries. Similarly, the maximum execution time of 1 hour defined in the database domain can be overridden by a value of 10 minutes in a second service superclass geared towards ensuring that shorter, important queries can complete quickly.

ACTIVITYTOTALTIME threshold

The ACTIVITYTOTALTIME threshold specifies the maximum amount of time that the data server should spend processing an activity.

Type Activity

Definition domain

Database, service superclass, service subclass, work action, workload

Enforcement scope

Database

Tracked work

Recognized coordinator and nested activities (see “Activities” on page 13)

Queuing

No

Unit Time duration expressed in minutes, hours, or days

Predictive or reactive

Reactive

In situations where the activity is queued by a queuing threshold, the total activity time includes the time spent in the queue awaiting execution. When a cursor is opened, the activity associated with the cursor lasts until the cursor is closed.

When a time threshold is applied to a stored procedure, it also applies to work happening inside the stored procedure. Consequently, when a stored procedure time threshold expires, any work happening inside the stored procedure is stopped. Hierarchies of stored procedure invocations can lead to hierarchies of time thresholds being applied to activities executing in the innermost levels of nesting. The most restrictive time threshold in the hierarchy (that is, the time threshold with the closest deadline) is always the one that applies.

The data server considers IMPORT, EXPORT, and other CLP commands to be user logic. Activities that are invoked from within IMPORT, EXPORT, and other CLP commands are subject to thresholds.

CPUTIME threshold

The CPUTIME threshold specifies the maximum amount of combined user and system processor time that an activity can use on a particular database partition while the activity is running. Use this threshold to detect and control activities that are using excessive processor resources.

Type Activity

Definition domain

Database, work action, service superclass, service subclass, and workload

Enforcement scope

Database partition

Tracked work

See the information later in this topic

Queuing

No

Unit Time

Predictive or reactive

Reactive

The amount of processor time that an activity spends running is measured from the time that the activity begins running at the partition, after any queuing by thresholds, until the time that the activity finishes running.

Activities tracked by this threshold are as follows:

- All DML activities.
- CALL activities. The processor time for a CALL activity does not include the processor time of any child activity. The processor time spent in fenced processes is also not counted toward the total processor time for the CALL activity.

Activities that are initiated by the database manager through a utility or procedure, with the exception of the ADMIN_CMD procedure, are not counted for this condition. The data server considers IMPORT, EXPORT, and other CLP commands to be user logic. Activities that are invoked from within IMPORT, EXPORT, and other CLP commands are subject to thresholds. Child activities of the LOAD command are not tracked by this threshold.

Example

The following example creates a CPUTIME threshold TH1 for the database domain with a database partition enforcement scope. This threshold stops any activity that takes longer than 30 seconds to run, which it checks for at 5-second intervals. You can use this threshold to ensure that no queries on the system use an unreasonable amount of processor time, which can negatively impact other work running on the system.

```
CREATE THRESHOLD TH1 FOR DATABASE ACTIVITIES
  ENFORCEMENT DATABASE PARTITION
  WHEN CPUTIME > 30 SECONDS CHECKING EVERY 5 SECONDS
  STOP EXECUTION;
```

CPUTIMEINSC threshold

The in-service-class CPUTIMEINSC threshold specifies the maximum amount of combined user and system processor time that an activity may use on a particular database partition while running in a specific service subclass. Use this threshold to detect and control activities that are using excessive processor resources.

Class Activity

Definition domain

Service subclass

Enforcement scope

Database partition

Tracked work

See the information later in this topic

Queuing

No

Unit Time

Predictive or reactive

Reactive

The processor time that an activity spends running is measured from the time that the activity enters the current service subclass until the time that the activity leaves the service subclass or finishes running.

This threshold differs from the CPU TIME threshold in that it controls only the amount of processor time that may be used in a specific service subclass, not the total amount of processor time used during the lifetime of the activity.

Activities tracked by this threshold are as follows:

- All DML activities.
- CALL activities. Note that the processor time for a CALL activity does not include the processor time of any child activity. The processor time spent in fenced processes is also not counted towards the total processor time for the CALL activity.

Activities that are initiated by the database manager through a utility or procedure, with the exception of the ADMIN_CMD procedure, are not counted for this condition. The data server considers IMPORT, EXPORT, and other CLP commands to be user logic. Activities that are invoked from within IMPORT, EXPORT, and other CLP commands are subject to thresholds. Child activities of the LOAD command are not tracked by this threshold.

You can use the REMAP ACTIVITY action to control activities by remapping them to a service subclass with different resource assignments.

Example

The following example creates two service subclasses, A1 and A2, under a superclass A, with a single in-service-class CPU TIME threshold that remaps activities between subclasses after 1 minute of processor time has been used during query evaluation in service subclass A1. An event monitor record is logged.

```
CREATE SERVICE CLASS A;  
CREATE SERVICE CLASS A1 UNDER A;  
CREATE SERVICE CLASS A2 UNDER A;  
  
CREATE THRESHOLD T1 FOR SERVICE CLASS A1 UNDER A  
  ACTIVITIES ENFORCEMENT DATABASE PARTITION  
  WHEN CPU TIME > 1 MINUTE CHECKING EVERY 30 SECONDS  
  REMAP ACTIVITY TO A2 LOG EVENT MONITOR RECORD;
```

ESTIMATEDSQLCOST threshold

The ESTIMATEDSQLCOST threshold specifies the maximum estimated cost that is permitted for DML activities.

Type Activity

Definition domain

Database, service superclass, service subclass, work action, and workload

Enforcement scope

Database

Tracked work

See the information later in this topic

Queuing

No

Unit Estimated SQL cost expressed in timerons

Predictive or reactive

Predictive

Activities tracked by this threshold are as follows:

- DML activities that are issued at the coordinator partition.
- Nested DML activities that are invoked from a user application. Consequently, DML activities that are issued by the data server internally, such as DML activities issued from within the DB2 utilities, SYSPROC stored procedures, and internal SQL, are unaffected by this threshold unless their cost is included in the parent activity estimate. In this situation, these activities are indirectly tracked. A trigger is an example of an indirectly tracked activity. IMPORT, EXPORT, and other CLP commands are considered to be user logic. Activities that are invoked from within IMPORT, EXPORT, and other CLP commands are subject to thresholds. For information about the activities that fall under a work class with the DML work type, see “Work identification by type of work with work classes” on page 45.

SQLROWSREAD threshold

The SQLROWSREAD threshold specifies the maximum number of rows that a DML activity may read on any database partition. Use this threshold to detect and control activities that are reading an excessive number of rows.

Class Activity

Definition domain

Database, work action, service superclass, service subclass, and workload

Enforcement scope

Database partition

Tracked work

See the information later in this topic

Queuing

No

Unit Number of rows

Predictive or reactive

Reactive

This threshold differs from the SQLROWSRETURNED threshold in that it controls the maximum number of rows read during query evaluation, not the number of rows returned to a client application from the data server.

Index accesses are not counted toward the total number of rows read. If an access plan uses only indexes during query evaluation, the SQLROWSREAD threshold will not be violated.

This threshold is evaluated at user-configurable time intervals; if the interval is greater than the amount of time it takes to exceed the number of rows read, it is possible for the number of rows read for an activity on a partition to exceed the threshold boundary, before the violation is detected.

Activities tracked by this threshold are as follows:

- Coordinator activities of type DML and corresponding subagent work such as subsection execution.
- Nested DML activities that are derived from user applications. Consequently, DML activities that are issued by DB2 logic, such as utilities, SYSPROC procedures, or internal SQL statements, are unaffected by this threshold. IMPORT, EXPORT, and other CLP commands are considered to be user logic; therefore, activities that are invoked from within IMPORT, EXPORT, and other CLP commands are subject to thresholds.

Example

The following example creates an SQLROWSREAD threshold TH1 for the database domain with a database partition enforcement scope. This threshold stops the execution of any activity that reads more than 5 000 000 rows during query evaluation, which the threshold checks for at 10-second intervals. You can use this threshold to ensure that no queries on the system read an unreasonable number of rows, which can negatively impact other work running on the system.

```
CREATE THRESHOLD TH1 FOR DATABASE ACTIVITIES
  ENFORCEMENT DATABASE PARTITION
  WHEN SQLROWSREAD > 5000000 CHECKING EVERY 10 SECONDS
  STOP EXECUTION;
```

SQLROWSREADINSC threshold

The in-service-class SQLROWSREADINSC threshold specifies the maximum number of rows that a DML activity can read on a particular database partition while running in a specific service subclass. Use this threshold to detect and control activities that are reading an excessive number of rows.

Class Activity

Definition domain

Service subclass

Enforcement scope

Database partition

Tracked work

See the information later in this topic

Queuing

No

Unit Number of rows

Predictive or reactive

Reactive

This threshold differs from the SQLROWSREAD threshold in that it controls the number of rows read only from the time that an activity enters a specific service subclass, not the total number of rows read during the lifetime of the activity. This threshold also differs from the SQLROWSRETURNED threshold in that it controls the maximum number of rows read during query evaluation in the current service subclass, not the number of rows returned to a client application from the data server.

Index accesses are not counted toward the total number of rows read. If an access plan uses only indexes during query evaluation, the SQLROWSREADINSC threshold will not be violated.

This threshold is evaluated at user-configurable time intervals; if the interval is greater than the amount of time it takes to exceed the number of rows read, it is possible for the number of rows read for an activity on a partition to exceed the threshold boundary, before the violation is detected.

Activities tracked by this threshold are as follows:

- Coordinator activities of type DML and corresponding subagent work such as subsection execution.

- Nested DML activities that are derived from user applications. Consequently, DML activities that are issued by DB2 logic, such as utilities, SYSPROC procedures, or internal SQL statements, are unaffected by this threshold. IMPORT, EXPORT, and other CLP commands are considered to be user logic; therefore, activities that are invoked from within IMPORT, EXPORT, and other CLP commands are subject to thresholds.

You can use the REMAP ACTIVITY action to control activities by remapping them to a service subclass with different resource assignments.

Example

The following example creates two service subclasses, A1 and A2, under a superclass A, with a single in-service-class SQLROWSREADINSC threshold that remaps activities between subclasses after 10 000 rows have been read in service subclass A1 during query evaluation. An event monitor record is logged.

```
CREATE SERVICE CLASS A;
CREATE SERVICE CLASS A1 UNDER A;
CREATE SERVICE CLASS A2 UNDER A;

CREATE THRESHOLD T1 FOR SERVICE CLASS A1 UNDER A
  ACTIVITIES ENFORCEMENT DATABASE PARTITION
  WHEN SQLROWSREADINSC > 10000 REMAP ACTIVITY TO A2
  LOG EVENT MONITOR RECORD;
```

SQLROWSRETURNED threshold

The SQLROWSRETURNED threshold specifies the maximum number of rows that can be returned by the data server to the client.

Type Activity

Definition domain

Database, service superclass, service subclass, work action, and workload

Enforcement scope

Database

Tracked work

See the information later in this topic

Queuing

No

Unit Number of rows

Predictive or reactive

Reactive

When multiple result sets are returned by a CALL statement, the threshold applies to each result set separately and not as an aggregate to the total number of rows returned across all result sets. For example, if you define the threshold for 20 rows and the CALL statement returns two result sets returning 15 rows and 19 rows respectively, the threshold is not triggered.

Activities tracked by this threshold are as follows:

- DML activities that are issued at the coordinator partition.
- Nested DML activities that are invoked from a user application. Consequently, DML activities that are issued by the data server internally, such as DML activities issued from within the DB2 utilities, SYSPROC stored procedures, and internal SQL, are unaffected by this threshold.

SQLTEMPSPACE threshold

The SQLTEMPSPACE threshold specifies the maximum amount of system temporary table space that can be consumed by a DML activity at any database partition. DML activities often use temporary table space for operations such as sorting and the manipulation of intermediate result sets.

Type Activity

Definition domain

Database, service superclass, service subclass, work action, workload

Enforcement scope

Database partition

Tracked work

See the information later in this topic

Queuing

No

Unit Amount of temporary table space expressed in kilobytes (KB), megabytes (MB), or gigabytes (GB)

Predictive or reactive

Reactive

Activities tracked by this threshold are as follows:

- DML activities that are issued at the coordinator partition.
- Nested DML activities that are derived from user applications. Consequently, DML activities that are issued by DB2 logic (such as utilities, SYSPROC procedures, or internal SQL) are unaffected by this threshold.

The data server considers IMPORT, EXPORT, and other CLP commands to be user logic. Activities that are invoked from within IMPORT, EXPORT, and other CLP commands are subject to thresholds.

Aggregate thresholds

An aggregate threshold places collective control over elements of work in a database. The boundary that you define using an aggregate threshold operates as a running total, to which any work tracked by the threshold contributes.

When newly instantiated work causes the upper boundary to be violated, the corresponding action is triggered. The work that caused the upper boundary to be violated is the only one affected by the triggered action.

Activity queuing

Some thresholds have a built-in queue and permit you to enforce how many activities can execute concurrently by queuing all additional activities once the concurrency limit is reached, up until the set limit for the queue is exceeded.

When an activity violates the threshold boundary of a queuing threshold, new work requests are queued automatically in a first-in, first-out fashion, until the queue reaches the size specified by the queuing boundary. When the queue is full, the upper boundary is reached, and the action specified for the threshold is applied to any newly arriving work being tracked by that threshold. For example, an action of STOP EXECUTION causes the newly arriving work to be rejected.

You can also define the upper queuing boundary as being unbounded, in which case there is no upper limit to the size of the queue. In this situation, newly arriving work is added to the queue. If you define a hard limit for the upper boundary and define an action of CONTINUE as the threshold action, all newly arriving work that violates the threshold boundary is added to the queue, and the threshold behaves as if its queuing boundary were unbounded.

AGGSQLTEMPSPACE threshold

The AGGSQLTEMPSPACE threshold specifies the maximum amount of system temporary table space that can be used in total across all concurrently running activities in a service subclass.

Class Aggregate

Definition domain

Service subclass

Enforcement scope

Database partition

Tracked work

See the information later in this topic

Queuing

No

Unit Kilobytes, megabytes, or gigabytes

Predictive or reactive

Reactive

Activities tracked by this threshold are as follows:

- DML activities that are issued at the coordinator partition.
- Nested DML activities that are derived from user applications. Consequently, DML activities that are issued by DB2 logic, such as utilities, SYSPROC procedures, or internal SQL statements, are unaffected by this threshold. IMPORT, EXPORT, and other CLP commands are considered to be user logic; therefore, activities that are invoked from within IMPORT, EXPORT, and other CLP commands are subject to thresholds.

CONCURRENTDBCOORDACTIVITIES threshold

The CONCURRENTDBCOORDACTIVITIES threshold specifies the maximum number of recognized coordinator activities that can run concurrently across all database partitions in the specified definition domain.

If an application starts more than one concurrent activity, it might have to pass this threshold more than once, potentially consuming the concurrency available for this threshold and creating a self-deadlock scenario.

Type Aggregate

Definition domain

Database, work action, service superclass, service subclass

Enforcement scope

Database

Tracked work

Recognized coordinator and nested activities (see further below and “Work identification by type of work with work classes” on page 45)

Queuing

Yes

Unit Number of concurrent database activities

Predictive or reactive

Predictive

This threshold is a generalization of the `CONCURRENTWORKLOADACTIVITIES` threshold. The `CONCURRENTWORKLOADACTIVITIES` applies only to activities running in a workload domain, but you can apply the `CONCURRENTDBCOORDACTIVITIES` threshold to a variety of domains, ranging from the entire database to a single work action. Similar to the `CONCURRENTWORKLOADACTIVITIES` threshold, the `CONCURRENTDBCOORDACTIVITIES` threshold tracks coordinator activities and any nested activities generated. Unlike the `CONCURRENTWORKLOADACTIVITIES` threshold, the `CONCURRENTDBCOORDACTIVITIES` threshold is a queuing threshold.

When creating queuing thresholds of the `CONCURRENTDBCOORDACTIVITIES` type, be aware of configurations that might lead to queue-based contention or a deadlock scenario. For example:

1. A concurrency threshold of type `CONCURRENTDBCOORDACTIVITIES` is created with a maximum concurrency value of 1 and a queue size greater than 1.
2. An application opens a cursor (or calls a stored procedure) that the DB2 data server recognizes as activity A1, which consumes the unique ticket that is available for the threshold.
3. While the activity A1 is still active, the application now issues a second SQL statement, which the data server recognizes as activity A2, and which is also subject to the concurrency threshold. Because the A1 activity is already running, the new activity A2 is queued.

The application is now in a deadlock state that cannot be resolved. It is waiting for A2 to execute but A2 is waiting for A1 to finish executing. This situation will not resolve itself without external intervention. It will not be detected and resolved by the deadlock detector. .

This example can be generalized to multiple applications and queues. You can resolve this situation by increasing the concurrency values, or cancelling certain activities if the concurrency values are correctly set. You can also use the `ACTIVITYTOTALTIME` threshold to prevent an activity from remaining queued indefinitely, which permits scenarios such as this one to resolve themselves without additional intervention.

To reduce the chance of creating inadvertent deadlock scenarios, the `CONCURRENTDBCOORDACTIVITIES` threshold affects different types of activities as follows:

- `CALL` statements are not controlled by the threshold, but all nested child activities are under threshold control. Note that both anonymous blocks and autonomous routines are classified as `CALL` statements.
- User defined functions (UDFs) are under threshold control, but child activities and requests nested within UDFs are not controlled. If an autonomous routine is called from within a user defined function, neither the autonomous routine and nor any child activities of the autonomous routine are under threshold control.

- Trigger actions that invoke CALL statements and the child activities of these CALL statements are not under threshold control. Note that the INSERT, UPDATE or DELETE statements themselves that can cause a trigger activate are under threshold control.

CONCURRENTWORKLOADACTIVITIES threshold

The CONCURRENTWORKLOADACTIVITIES threshold specifies the maximum number of coordinator and nested activities that can concurrently run in a workload occurrence.

Type Aggregate

Definition domain
Workload

Enforcement scope
Workload occurrence

Tracked work
Recognized coordinator and nested activities (see “Activities” on page 13)

Queuing
No

Unit Number of concurrent workload activities

Predictive or reactive
Predictive

This threshold applies to a single workload occurrence. If you have multiple occurrences of a workload running concurrently, the threshold applies separately to each workload occurrence. The tracked activities included all recognized coordinator activities and any nested activities that are generated as a result of the execution of the coordinator activity. For example, if a stored procedure is called and that stored procedure executes some SQL, both the CALL statement (which is the coordinator activity) and the SQL statements executed by the stored procedure (which are the nested activities) count towards the threshold total.

COMMIT, ROLLBACK, and ROLLBACK to SAVEPOINT statements are unaffected by this threshold.

Nested activity considerations

The nested activities that are tracked by this threshold must satisfy the following criteria:

- They must be a recognized coordinator activity. Nested coordinator activities that are not recognized types as described in “Work identification by type of work with work classes” on page 45 are not counted.
- They must be directly invoked from user logic, such as a user-written stored procedure issuing SQL or from the SYSPROC.ADMIN_CMD stored procedure. Nested coordinator activities that are started by the invocation of a DB2 utility or any other code in the SYSIBM, SYSFUN, or SYSPROC schemas are not counted towards the upper boundary specified by this threshold.

Example

In this example, the CONCURRENTWORKLOADACTIVITIES threshold maximum value is set to 5. The user logic causes the following sequence of operations to occur in a workload occurrence:

1. Issue a load command: the current number of workload activities is 1.
 - The load command internally issues some SQL. The current number of workload activities is 1. (SQL generated by a utility does not count against the CONCURRENTWORKLOADACTIVITIES threshold.)
 - The load command ends. The current number of workload activities is 0.
2. CALL the SYSPROC.SP1 stored procedure. The current number of workload activities is 1.
 - The SYSPROC.SP1 stored procedure generates some SQL. The current number of workload activities is 1. (SQL generated by a utility does not count against the CONCURRENTWORKLOADACTIVITIES threshold.)
 - The SYSPROC.SP1 stored procedure ends. The current number of workload activities is 0.
3. Open a cursor C1. The current number of workload activities is 1.
4. Issue a runstats command. The current number of workload activities is 1.
 - The runstats command generates some SQL. The current number of workload activities is 1.
 - The runstats command ends. The current number of workload activities is 1.
5. Close the cursor C1. The current number of workload activities is 0.
6. CALL the BOB.SP1 stored procedure. The current number of workload activities is 1.
 - The BOB.SP1 stored procedure opens three cursors. The current number of workload activities is 4.
 - The BOB.SP1 stored procedure calls the SYSPROC.SP2 stored procedure. The current number of workload activities is 5.
 - The SYSPROC.SP2 stored procedure issues some SQL. The current number of workload activities is 5.
 - The SYSPROC.SP2 stored procedure ends. The current number of workload activities is 4.
 - The BOB.SP1 stored procedure calls the BOB.SP2 stored procedure. The current number of workload activities is 5.
 - The BOB.SP2 stored procedure issues some SQL. At this point, the threshold is triggered.
 - The BOB.SP2 stored procedure ends. The current number of workload activities is 4.
 - The BOB.SP1 stored procedure ends. The current number of workload activities is 0.
7. Open a cursor C2. The current number of workload activities is 1.
8. CALL the BOB.SP2 stored procedure. The current number of workload activities is 2.

CONCURRENTWORKLOADOCCURRENCES threshold

The CONCURRENTWORKLOADOCCURRENCES threshold is an aggregate threshold that specifies the maximum number of workload occurrences that can run concurrently on the coordinator partition.

Type Aggregate

Definition domain
Workload

Enforcement scope
Database partition

Tracked work

Workload occurrences

Queuing

No

Unit Number of concurrent workload occurrences**Predictive or reactive**

Predictive

When a workload occurrence is started, if the work that it generates is sent to non-coordinator partitions, the work on these partitions does not count towards the concurrent threshold total on the coordinator partition. For example, assume that a `CONCURRENTWORKLOADOCCURRENCES` threshold is defined to permit only one occurrence of workload A on a database partition. Then assume that an application connects to database partition 1, resulting in an occurrence of workload A being started, and that this workload causes work to be sent to database partitions 1, 2, and 3. In this situation, the total number of occurrences of workload A is one on database partition 1 and zero on database partitions 2 and 3. Therefore, if another application connects to database partition 1 and another occurrence of workload A is started on database partition 1, that workload is rejected. However, new occurrences of workload A can still be started on database partitions 2 and 3.

TOTALDBPARTITIONCONNECTIONS threshold

The `TOTALDBPARTITIONCONNECTIONS` threshold specifies the maximum number of concurrent database connections on a coordinator partition for a database, that is, this threshold controls the maximum number of clients that can connect to the database on each of its database partitions.

This threshold is not enforced for users with `DBADM` authority.

Type Aggregate**Definition domain**

Database

Enforcement scope

Database partition

Tracked work

Connections

Queuing

Yes (enforced at 0)

Unit Number of concurrent connections**Predictive or reactive**

Predictive

For example, if you set the `TOTALDBPARTITIONCONNECTIONS` threshold to 10 and the database has five partitions, each partition can have up to 10 clients connected concurrently, for a total of 50 client connections across the entire database.

The `TOTALDBPARTITIONCONNECTIONS` threshold controls only coordinator connections. Connections made by subagents are not counted towards the threshold.

This threshold is useful for situations in which you want to have multiple databases in the same instance. Setting a TotalDBPartitionConnections threshold on a database partition ensures that client connections from one database cannot use all of the available connections on a database partition.

Ensure that you set the **max_connections** database manager configuration parameter high enough to support the maximum number of connections that you expect across the database. If you set a TOTALDBPARTITIONCONNECTIONS threshold for a database, you must set **max_connections** to at least the threshold value. If you want to run multiple databases on the same instance, ensure that you set **max_connections** high enough to support the maximum number of connections for all databases. The data server does not check for this condition because it is impossible to know beforehand how many of the databases will be active concurrently.

TOTALSPARTITIONCONNECTIONS threshold

The TOTALSPARTITIONCONNECTIONS threshold specifies the maximum number of concurrent database connections on a coordinator partition for a service superclass.

Type Aggregate

Definition domain

Service superclass

Enforcement scope

Database partition

Tracked work

Connections

Queuing

Yes

Unit Number of concurrent connections in service class

Predictive or reactive

Predictive

When the TOTALSPARTITIONCONNECTIONS threshold in the service class is reached, subsequent coordinator connections that join the service superclass are queued until the specified queue size is reached. By default, the queue size is zero, which means that no connections can be queued. If a connection joins the queue of a TOTALSPARTITIONCONNECTIONS threshold, the connection is considered to be in a *transient* state.

Tracked connections include both new client connections and existing client connections that switch to the service class from another service class. Connections switch service classes by associating with a different workload definition that is mapped to a different service class. Workload reevaluation occurs only at transaction boundaries, so connections can switch service classes only at transaction boundaries; however, because resources that are associated with WITH HOLD cursors are maintained across transaction boundaries, connections with open WITH HOLD cursors cannot switch service superclasses. When the connection concentrator is on, any application that is switched leaves the service class. When the application is switched in at the subsequent statement, it must rejoin the service class and consequently pass the threshold.

When the queue size threshold is reached, the threshold action is triggered. The TOTALSCPARTITIONCONNECTIONS threshold controls only coordinator connections. Connections made by subagents are not counted towards the threshold.

If you set a threshold value for TOTALDBPARTITIONCONNECTIONS, set it large enough to accommodate the threshold that you specify for TOTALSCPARTITIONCONNECTIONS. For example, if you define five service superclasses for a database and each of them has a TOTALSCPARTITIONCONNECTIONS threshold value of 10, the TOTALDBPARTITIONCONNECTIONS threshold value should be at least 50.

Creating a threshold

Create thresholds using the DDL statement CREATE THRESHOLD (or the CREATE WORK ACTION SET statement). You create a threshold to impose a limit on resource consumption.

Before you begin

To create a threshold, you require WLMADM or DBADM authority.

See the following topics for more information about prerequisites:

- “DDL statements for DB2 workload manager” on page 16
- Naming rules

To create a threshold for a work action set, use the CREATE WORK ACTION SET statement or the ALTER WORK ACTION SET statement with the ADD WORK ACTION keywords. For more information, see CREATE WORK ACTION SET statement or ALTER WORK ACTION SET statement.

Procedure

To create a threshold:

1. Issue the CREATE THRESHOLD statement, specifying one or more of the following properties for the threshold:
 - The name of the threshold.
 - The threshold domain. The threshold domain is the database object that the threshold is both attached to and operates on. The domain that applies depends on the type of threshold, see “Threshold domain and enforcement scope” on page 91 for more information.
 - The enforcement scope for the threshold. The threshold scope is the enforcement range of the threshold in its domain. The enforcement scope that applies depends on the type of threshold, see “Threshold domain and enforcement scope” on page 91 for more information.
 - Optional: Disable the threshold when it is created. By default a threshold is created as enabled. If you create the threshold as disabled and want to enable it later, use the ALTER THRESHOLD statement.
 - The threshold predicate to specify the type of threshold and the maximum value permitted. When the maximum value is violated, the action specified for the threshold is enforced. For more information on which thresholds are available to you, see “Connection thresholds” on page 94, “Activity thresholds” on page 95 and “Aggregate thresholds” on page 102.

- The actions to be taken if the maximum value for the threshold is exceeded. The actions consist of a mandatory action that affects the execution of the activity (STOP EXECUTION, CONTINUE, or REMAP ACTIVITY TO) and an optional collect activity action (COLLECT ACTIVITY DATA). The options you specify for the collect activity action determine what information is collected for the activity that caused the threshold boundary to be violated.
2. Commit your changes. When you commit your changes, the threshold is added to the SYSCAT.THRESHOLDS view.

Altering a threshold

Alter thresholds using the ALTER THRESHOLD statement. You might alter a threshold to modify the limit imposed on a specific resource.

Before you begin

To alter a threshold, you require SQLADM, WLMADM or DBADM authority. To specify any clause other than a COLLECT clause, the authorization ID must include WLMADM or DBADM authority.

See “DDL statements for DB2 workload manager” on page 16 for more information about prerequisites.

To alter a threshold for a work action set, use the ALTER WORK ACTION SET statement with the ADD WORK ACTION keywords. For more information, see ALTER WORK ACTION SET statement.

Restrictions

You cannot alter the threshold type with the ALTER THRESHOLD statement. For example: You cannot change a TOTALDBPARTITIONCONNECTIONS threshold into a TOTALSCPARTITIONCONNECTIONS threshold, for example. If you require a different threshold type, drop the existing thresholds and then create a new threshold.

Procedure

To alter a threshold:

1. Specify one or more of the following properties for the threshold on the ALTER THRESHOLD statement. You can change the following properties:
 - The boundary for the threshold predicate.
 - The actions to be taken, if the threshold boundary is violated.
 - Whether the threshold is enabled or disabled.
2. Commit your changes. When you commit your changes, the threshold is updated in the SYSCAT.THRESHOLDS view.

Dropping a threshold

Drop a threshold that you no longer require using the DDL statement DROP THRESHOLD.

Before you begin

To drop a threshold, you require WLMADM or DBADM authority.

See “DDL statements for DB2 workload manager” on page 16 for more information about prerequisites.

If you want to drop a threshold in a work action set, use the ALTER WORK ACTION SET statement. You can also drop a threshold by dropping the entire WORK ACTION SET with the DROP statement.

Procedure

To drop a threshold:

1. Do one of the following steps:
 - If the threshold is a queuing threshold, use the ALTER THRESHOLD statement to disable it.
 - If you disabled a queuing threshold by using an ALTER THRESHOLD statement, issue a COMMIT statement to commit the change.
2. Use the DROP THRESHOLD statement to drop the threshold.
3. Commit your changes. When you commit your changes the threshold is removed from the SYSCAT.THRESHOLDS view.

Example: Using thresholds

You can use thresholds for a variety of purposes. In this scenario, thresholds are used to control the number of large jobs running in order to permit different execution times for different applications, and to control the behavior of an application that is in development.

One way of setting up a DB2 workload manager solution is to divide and manage the database resources across the various departments in a company. For example, assume that the sales department runs two main reports, which consist of the monthly and yearly sales. Assume also that the human resources department runs a payroll application every other week and that the development team is working on a new type of report at the request of the management team. To define WLM execution environments for these departments, create service classes:

```
CREATE SERVICE CLASS SALES
CREATE SERVICE CLASS HUMANRESOURCES
CREATE SERVICE CLASS DEVELOPMENT
```

In this situation, you create a workload definition for each one of these applications to map the application to its applicable service superclass:

```
CREATE WORKLOAD MONTHLYSALES APPLNAME('monthlyrpt.exe') SERVICE CLASS SALES
CREATE WORKLOAD YEARLYSALES APPLNAME('yearlyrpt.exe') SERVICE CLASS SALES
CREATE WORKLOAD PAYROLL APPLNAME('payroll.exe') SERVICE CLASS HUMANRESOURCES
CREATE WORKLOAD NEWREPORT APPLNAME('dev.exe') SERVICE CLASS DEVELOPMENT
```

The database catalog therefore contains the following workload definitions:

- MonthlySales, mapping to the service superclass Sales
- YearlySales, mapping to the service superclass Sales
- Payroll, mapping to the service superclass Human Resources
- NewReport, mapping to the service superclass Development

Threshold on the number of large jobs

Because the YearlySales report is very large, you do not want to have more than one occurrence of this application running in the database at any time. You therefore create a threshold to set the maximum number of concurrent occurrences of this workload to 1:

```
CREATE THRESHOLD SINGLEYEARLYSALESRPT FOR WORKLOAD YEARLYSALES ACTIVITIES
  ENFORCEMENT DATABASE PARTITION
  WHEN CONCURRENTWORKLOADOCCURRENCES > 1
  STOP EXECUTION
```

You can achieve a similar solution by associating the YearlySales application with a service subclass YearlySalesReports (under the Sales service superclass) and setting the maximum concurrency threshold to a value of 1 for the service subclass:

```
CREATE SERVICE CLASS YEARLYSALESREPORTS UNDER SALES

ALTER WORKLOAD YEARLYSALES SERVICE CLASS YEARLYSALESREPORTS UNDER SALES

CREATE THRESHOLD SINGLEYEARLYSALESREPORT FOR SERVICE CLASS YEARLYSALESREPORTS
  UNDER SALES ACTIVITIES ENFORCEMENT DATABASE
  WHEN CONCURRENTDBCOORDACTIVITIES > 1
  STOP EXECUTION
```

In either situation, you can set the threshold action to STOP EXECUTION to prevent more than one occurrence of the workload from executing. You can also collect activity information if you want additional information about the conditions when the threshold is violated.

Threshold on activity lifetimes

Because all applications are expected to complete in an hour or less, you create a threshold with a database domain, preventing any activity from running longer than 1 hour. The only exception to this rule is the yearly report, which can take up to 5 hours to complete. Therefore, you can associate an activity total time threshold of 5 hours with the YearlySales workload. This will override the activity total time threshold applied to the yearly sales report, relaxing the time constraints. The new value of 5 hours now applies to the YearlySales workload although the global value of 1 hour applies elsewhere in the database:

```
CREATE THRESHOLD MAXDBACTIVITYTIME FOR DATABASE ACTIVITIES
  ENFORCEMENT DATABASE
  WHEN ACTIVITYTOTALTIME > 1 HOUR
  STOP EXECUTION

CREATE THRESHOLD MAXYRPTACTIVITYTIME FOR WORKLOAD YEARLYSALES
  ACTIVITIES ENFORCEMENT DATABASE
  WHEN ACTIVITYTOTALTIME > 5 HOURS
  STOP EXECUTION
```

Threshold on the number of coordinator and nested activities

The NewReport application makes heavy use of stored procedures and user-defined functions and is not fully debugged yet, so it tends to generate large numbers of activities that impact the rest of the system. After consulting with the developer, you learn that this new report is not supposed to generate more than 20 activities in total, so you define a threshold of type workload activities on the NewReport workload and set it to 20. Initially, you set the threshold action to

STOP EXECUTION and COLLECT ALL to stop any unwanted side effect of the application starting large numbers of activities and to help the developer identify any problems:

```
CREATE THRESHOLD MAXDEVACTIVITIES FOR SERVICE CLASS DEVELOPMENT ACTIVITIES
ENFORCEMENT DATABASE
WHEN CONCURRENTDBCOORDACTIVITIES > 20
COLLECT ACTIVITY DATA WITH DETAILS AND VALUES
STOP EXECUTION
```

When the application becomes more stable, it enters its optimization phase. During the phase, the developer tries to reduce the number of activities generated by the application from between 15 and 20 to 15. At this time, you alter the threshold by changing its upper boundary value to 15 and the threshold action to CONTINUE. This threshold definition helps identify and address situations in which the number of generated activities exceeds 15 but the increased stability of the application does not require that its execution be stopped.

```
ALTER THRESHOLD MAXDEVACTIVITIES
WHEN CONCURRENTDBCOORDACTIVITIES > 15
COLLECT ACTIVITY DATA ON COORDINATOR WITH DETAILS AND VALUES
CONTINUE
```

Priority aging of ongoing work

Priority aging can automatically change the priority of in-progress activities over time. You can use priority aging to control longer-running activities, so that throughput for shorter-running activities can be improved.

Changing the priority of activities by remapping

System resources are allocated and controlled by using service classes. With priority aging, the priority of an activity can be changed by moving the activity from one service class to another service class. The priority increases if the new service class has more resources, and the priority decreases if the new service class has fewer resources. Activities are moved when a threshold with a REMAP ACTIVITY action is violated, based upon predetermined maximum usage of a specific resource such as processor time or rows read. After an activity is mapped to a new service class, it continues to run with the new resource constraints applied.

A simple approach that you can use to help short queries to run faster is to define a series of service classes with successively lower levels of resource priority and threshold actions that move activities between the service subclasses. Using this setup, you can decrease, or age, the priority of longer-running work over time and perhaps improve response times for shorter-running work without having detailed knowledge of the activities running on your data server.

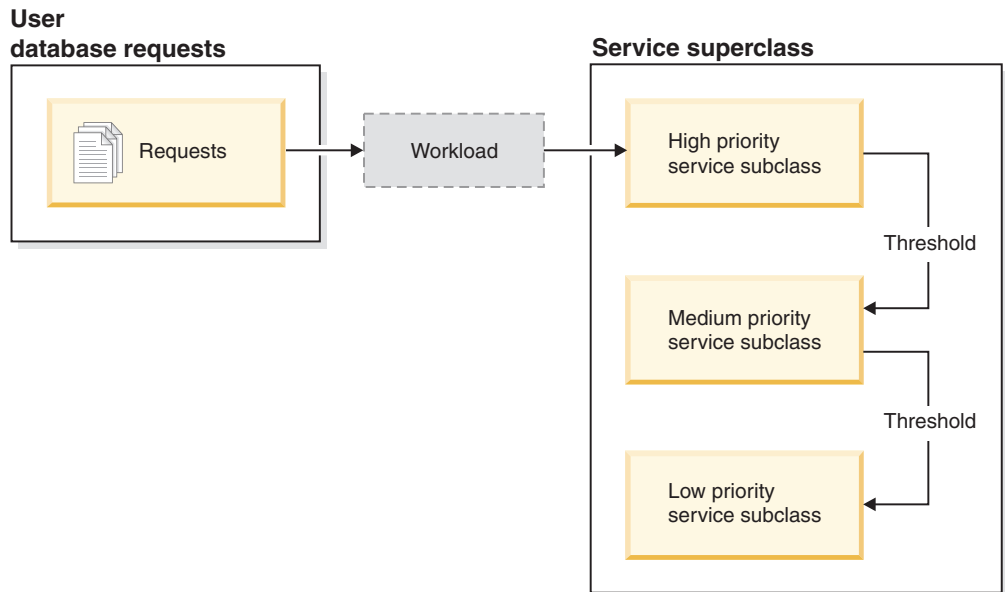


Figure 16. A simple tiered setup that shows three service classes with successively lower priority

You can create this setup by assigning a high priority for all applicable resources to one service class, medium priority to a second service class, and low priority to a third service class. As work enters the system, it is automatically placed into the first service class and begins running using the high-priority settings of this service class. If you also define thresholds for each of the service classes that limit the time or resources used during execution, work is dynamically reassigned to the next-lower service class if the threshold of the next-higher class is violated. This dynamic resource control is repeatedly applied until the work is completed or is in the lowest-priority class, where it remains until it is completed or you force it to stop running.

In-service-class thresholds

Remapping of activities is available with any of the in-service-class thresholds, which control the amount of a resource that may be used while an activity is running in a particular service subclass. Examples of resources are the amount of processor time used (CPUTIMEINSC threshold) and the number of rows read by an application (SQLROWSREADINSC threshold), per activity per partition. These thresholds differ from other activity thresholds, which control resources used throughout the entire lifetime of an activity.

Because of the control that in-service-class thresholds provide over service subclasses, you can define in-service-class thresholds only on a service subclass domain. The in-service-class thresholds provide controls similar to DB2 Governor rules, which act on processor time and rows read monitor elements.

When an in-service-class threshold is associated with a REMAP ACTIVITY action, agents working for the activity periodically check whether the threshold has been violated on each partition. If an agent detects a threshold violation on a partition, the agent triggers the REMAP ACTIVITY action for the activity on the partition and then remaps itself to the target service subclass. All other agents working for the activity on the same partition remap themselves to the target service subclass when they detect that the activity has been remapped. Only one agent detects the

threshold violation and remaps the activity, and the activity is considered remapped after that agent has detected the threshold violation and performed the remapping.

Two monitor elements provide information about activity remapping within service subclasses. The **act_remapped_in** monitor element provides a counter that records how many activities were remapped into a service subclass and is incremented each time for the target service subclass that an activity is remapped to. Similarly, the **act_remapped_out** monitor element counter is incremented each time for the source service subclass that an activity is remapped out of. An additional monitor element, **num_remaps**, counts the number of times in total that an activity has been remapped between service subclasses.

An activity can be remapped multiple times to different service subclasses, and an activity can return to its original service subclass after being remapped to another service subclass.

The in-service class thresholds are evaluated separately for an activity on each partition, without coordination. Because there is no coordination between partitions, when an activity is remapped on one partition, it is possible for the same activity to be in different service subclasses on different partitions simultaneously.

When subagent work for an activity is completed on a remote partition and further work for the same activity is sent to the same partition later, the activity restarts in the same service subclass as the agent that sent the request to the partition. If you defined an in-service-class threshold for this service subclass, the timer or counter for the activity on the remote partition restarts at zero.

Where activities are nested, parent and child activities are tracked separately. Therefore, if a child activity is using an excessive amount of resources, only this activity, not its parent or sibling activities, violates a threshold.

Using the in-service-class thresholds

On data servers where the primary resource activities have to compete for is processor time, use the `CPUTIMEINSC` threshold as your first measure of control. On data servers where queries reading many table rows result primarily in I/O contention, use `SQLROWSREADINSC`. On systems that see a combination of heavy processor and IO activity, use a combination of the `CPUTIMEINSC` and `SQLROWSREADINSC` thresholds.

You should set the agent priority of the service subclasses relative to each other, so that your data server can treat activities of different business priority differently. Note that the agent priority of the default system class should always be higher than any user defined service classes you create to avoid a negative impact on performance. The agent priority of the default maintenance class can be set lower than your user defined service classes.

How much of a given resource you permit activities to consume in a service subclass before remapping them to a different service subclass depends largely on your particular environment. To find the best value for each threshold condition, you need to monitor how activities are being processed on your data server. If the maximum amount of processor time that can be used or the maximum number of rows that can be read in a service class is set too high, activities will inappropriately start and finish in the same service subclass regardless of how

much resource each activity requires. If the maximum processor time or rows read is set too low, no activity will finish in the service class it is originally mapped to and every activity will end up being remapped to the another service class regardless of business priority. In either case, your tiered configuration will not benefit the overall throughput on your data server and activities are not treated according to their business priority effectively.

In addition to determining how much of a given resource an activity can consume, some thresholds allow you to define a check interval for how often the data server checks for threshold violations. This capability is provided for thresholds where it is too expensive to check the threshold each time a unit of the resource being controlled is consumed and determines the latency with which violations of these thresholds are detected. Both the CPUTIME and SQLROWSREAD thresholds and their in-service class counterparts CPUTIMEINSC and SQLROWSREADINSC support check intervals, which you can specify in the CHECKING EVERY clause when you create or alter one of these thresholds. On serial database instances, the check interval equals the the amount of real time that you want to elapse between checks for A threshold violation. On DPF or SMP instances, the check interval should be set to a value that is less than the amount of real time elapsed to take into account that there can be more than one agent accumulating processor time simultaneously for the activity. To calculate the approximate check interval on DPF or SMP instances, divide the amount of real time you want to elapse between checks by the degree of parallelism for the activity and use the resulting value for the CHECKING EVERY clause.

For example: In a single partition database, if you want a CPUTIMEINSC threshold to trigger a REMAP ACTIVITY action after 30 seconds of processor time have been consumed, you can set the check interval to 30 seconds and be certain that the threshold action will be triggered after no more than 30 seconds of processor time have been consumed (processor time used cannot outstrip real time elapsed). In a partitioned database environment, if you define a CPUTIMEINSC threshold that is set at 5 seconds with a check interval of 5 seconds, and an activity has 1 coordinator partition agent and 4 subagents working on its behalf, it is possible for the activity to consume 5 seconds of CPU time in just 1 second of real time, because 5 agents simultaneously accumulate 1 second of processor time each. To prevent the activity from consuming a multiple of 5 seconds of processor time, the check interval should in this case be set to 1 second.

For additional information on how to use the thresholds, see the sample tiering scripts and priority aging scenarios.

Effect of remapping on thresholds

Which thresholds continue to apply after remapping through a REMAP ACTIVITY action depends on whether the thresholds apply only to a specific service subclass or throughout the lifetime of an activity.

When you remap an activity to a new service subclass, only the in-service-class thresholds, such as CPUTIMEINSC and SQLROWSREADINSC, change. These in-service-class thresholds no longer affect an activity after it leaves the source service subclass, and they are replaced with the corresponding thresholds for the target subclass, if you defined those thresholds. All other activity thresholds from the service subclass to which the activity was originally mapped remain unchanged, and applicable threshold timers and counters are not reset. The activity is not re-evaluated against any other thresholds that you defined for the target service subclass.

For example, assume that two service subclasses with thresholds are defined as follows:

- Service subclass A with the following thresholds:
 - An `ACTIVITYTOTALTIME` lifetime threshold TH1 with a `STOP EXECUTION` action after 30 minutes have elapsed
 - An `SQLROWSREADINSC` in-service-class threshold TH2 with a `REMAP ACTIVITY` action to service subclass B after more than 2000 rows have been read
- Service subclass B with the following thresholds:
 - An `ACTIVITYTOTALTIME` lifetime threshold TH3 with a `STOP EXECUTION` action after 5 minutes have elapsed
 - An `SQLROWSREADINSC` threshold TH4 with a `STOP EXECUTION` action after more than 1000 rows have been read

When an activity enters the system in service subclass A, both thresholds TH1 and TH2 apply to the activity. If the activity reads more than 2000 rows during query evaluation, it is dynamically remapped to service subclass B. Because of the remapping of the activity to subclass B, the applicable in-service-class thresholds change, and TH4 rather than TH2 now applies to the activity. Counters for both thresholds are reset to zero, and even though the activity has read more than 2000 rows in the original service subclass, the counter for TH4 is restarted at zero; the activity must read more than 1000 rows while running in service subclass B before threshold TH4 is violated. Threshold TH1, which applies throughout the lifetime of the activity, continues to apply, even though the activity is now running in a different subclass. Threshold TH3 does not exercise any control over the remapped activity at all, because it did not apply to the first service subclass that the activity entered when it began running.

Sample priority aging scripts

Use the provided sample scripts to quickly create a tiered service class configuration on your data server. With a tiered configuration, you can address specific performance goals by decreasing the priority of longer-running queries over time, known as priority aging. You can also modify the scripts according to your own business priorities when adapting them to your environment.

The two sample scripts `wlmtiersdefault.db2` and `wlmtierstimerons.db2` are intended to demonstrate how you might use priority aging on your data server to improve overall throughput. DB2 workload manager provides you with the controls that can help with overall throughput on your data server, but to gain the full benefit of the scripts, and of priority aging in general, you will need to run your data server for an extended period of time and monitor how this work is performed, then adjust the service class and threshold settings accordingly.

The scripts are provided in the `samples/admin_scripts` directory under your installation directory.

The `wlmtiersdefault.db2` and `wlmtierstimerons.db2` scripts create three service subclasses under a common superclass with successively lower agent priority settings from high to low along with `CPUTIMEINSC` thresholds that move or remap activities in response to the consumption of processor time. The scripts differ in how activities are mapped to a service class when they first enter your data server. A third script, `wlmtiersdrop.db2`, drops the WLM objects created by the other two sample tiering scripts.

wlmtiersdefault.db2

All activities entering the data server are mapped to a high priority service subclass WLM_SHORT without differentiation between types of work. Activities will complete in the high priority service subclass whilst the highest priority is assigned to them, unless they exceed the maximum amount of processor time they are allowed to consume. Activities that consume too much processor time are first remapped by a threshold REMAP ACTIVITY action to a medium priority service subclass WLM_MEDIUM and then, if they still do not complete without exceeding the allotted processor time in that subclass, remapped to a low priority service subclass WLM_LONG, where they continue to be processed until they complete. Activities that cannot be remapped by a CPUTIMEINSC threshold are mapped directly to service subclass WLM_MEDIUM where they will remain.

wlmtierstimerons.db2

DML activities entering the data server are evaluated according to their estimated cost and mapped to one of the three service subclass. DML activities estimated to be short are mapped to a high priority service subclass WLM_SHORT, DML activities estimated to be of medium length are mapped to a WLM_MEDIUM service subclass that receives medium priority, and long DML activities are mapped to a WLM_LONG service subclass that receives the lowest priority. Non-DML activities enter the highest priority service subclass. As they are processed, activities that consume more processor time than assigned to a service subclass are successively remapped by a threshold REMAP ACTIVITY action to the next lowest priority service subclass, until they are remapped to the lowest priority service subclass where they continue to be processed until they complete. Activities that cannot be remapped by a CPUTIMEINSC threshold are mapped directly to service subclass WLM_MEDIUM where they will remain.

wlmtiersdrop.db2

This script drops all DB2 workload manager service classes, thresholds, workloads, work class sets and work action sets that are created by the scripts wlmtiersdefault.db2 and wlmtierstimerons.db2.

By default, the wlmtiersdefault.db2 and wlmtierstimerons.db2 scripts use the following service class and threshold definitions:

Table 37. Service classes with agent priorities and prefetch priority settings created by the scripts

Service class	Agent priority (on UNIX and Linux operating systems / on Windows operating systems)	Prefetch priority
WLM_SHORT (high priority)	-10 / 3	High
WLM_MEDIUM (medium priority)	0 / 0	Medium
WLM_LONG (low priority)	10 / -3	Low
Default system class	-15 / 5	High
Default maintenance class	15 / -5	Low

Table 38. Thresholds created by the scripts

Threshold	Maximum amount of processor time that can be used in the service class before remapping
WLM_TIERS_REMAP_SHORT_TO_MEDIUM	30 seconds
WLM_TIERS_REMAP_MEDIUM_TO_LONG	30 seconds

The wlmtiersdefault.db2 sample script creates the following work action set and work class set, which is used to map activities that cannot be remapped by the CPUTIMEINSC threshold directly to the WLM_MEDIUM service subclass. These activities will remain in the WLM_MEDIUM service subclass for the duration of their execution.

Table 39. Work class set created by the wlmtiersdefault.db2 sample script

Work class	Work action
WLM_DML_WC	For DML activities, mapped to service class WLM_SHORT initially. These activities can be remapped by a CPUTIMEINSC threshold.
WLM_CALL_WC	For CALL activities, mapped to service class WLM_SHORT initially. These activities can be remapped by a CPUTIMEINSC threshold.
WLM_OTHER_WC	For activities that cannot be remapped by a CPUTIMEINSC threshold, mapped to service class WLM_MEDIUM. These activities will remain in the WLM_MEDIUM service subclass.

The wlm-tierstimerons.db2 sample script also creates the following work action set and work class set, which is used to map activities according to their estimated cost:

Table 40. Work class set created by the wlm-tierstimerons.db2 sample script

Work class	Estimated cost range in timerons and work action
WLM_SHORT_DML_WC	For DML activities with an estimated cost of 0 to 999 timerons, mapped to service class WLM_SHORT initially. These activities may get remapped by a CPUTIMEINSC threshold.
WLM_MEDIUM_DML_WC	For DML activities with an estimated cost of 1000 to 99 999 timerons, mapped to service class WLM_MEDIUM initially. These activities may get remapped by a CPUTIMEINSC threshold.
WLM_LONG_DML_WC	For DML activities with an estimated cost of 100 000 to infinity timerons, mapped to service class WLM_LONG.
WLM_CALL_WC	For CALL activities, mapped to service class WLM_SHORT initially. These activities can be remapped by a CPUTIMEINSC threshold.
WLM_OTHER_WC	For activities that cannot be remapped, mapped to service class WLM_MEDIUM

Modifying the scripts for your environment

When you modify the sample scripts to adapt them to your environment, the most important setting to consider is the maximum amount of processor time that can be used in each service class. How much processor time you permit activities to consume in each service subclass depends largely on your particular environment. To find the best values, you need to monitor how activities are being processed on your data server. By default, both the `wlmtiersdefault.db2` and `wlmtierstimerons.db2` scripts will log event monitor records to the threshold violations event monitor, if one is active, with the option to turn on and enable the activity event monitor and to collect activity data (at the cost of incurring additional overhead). For `wlmtiersdefault.db2`, if the maximum amount of processor time that can be used in each service class is set too high, most activities will always start and finish in the high priority class regardless of how much actual processor time each requires. If the maximum amount of processor time is set too low, no activity will finish in the high priority service class and every activity will end up being remapped to the medium or low priority service class regardless of business priority. In either case, the script will not benefit overall throughput on your data server and activities are not treated according to their business priority effectively. The same issue is true to a lesser extent for `wlmtierstimerons.db2` where activities are differentiated initially by being mapped to service subclasses according to estimated cost. If the maximum amount of processor time that can be used in each service class is set incorrectly, activities will fail to be remapped to a more appropriate service subclass if they consume too much processor time, or are remapped too quickly despite having higher business priority.

Note that the `wlmtiersdefault.db2` and `wlmtierstimerons.db2` scripts set the agent priority of the default system class to be higher and the priority of the default maintenance class to be lower than the three user defined service classes. If you modify the agent priority of the user defined service classes, you should always set the priority of the default system class to be as high as or higher than the highest priority service subclass you create to avoid a negative impact on performance.

For more information about the specific DB2 workload manager objects created by the scripts and about how to run them, refer to the scripts.

Sample scenarios

Two examples have been included in the documentation that show you how you can adapt the sample tiering scripts on your data server to make use of priority aging.

Scenario: Controlling resource intensive business intelligence reports with priority aging

The following scenario shows how you can configure your data server to dynamically lower the priority of expensive business intelligence reports that cannot be identified before execution starts in order to maintain system performance for other queries.

The problem: There is a business intelligence report which any end user can run and which is very expensive. Anytime the report runs, it compromises the performance of the system. The front end tool used to generate the report does not set any client information that could be used to identify the report in advance which would permit you to map it to a low priority service class using a workload.

The solution: You can use the `wlmtiersdefault.db2` sample tiering scripts to configure your data server with a tiered configuration that dynamically lowers, or ages, the priority of processor intensive activities during their lifetime in order to prevent compromising data server performance for all other users. After a workload initially maps all work to a high priority service subclass, the expensive reports are detected by the `CPUTIMEINSC` in-service-class threshold based on the amount of processor time consumed. If an activity violates the `CPUTIMEINSC` threshold by using the maximum amount of allowed processor time, a `REMAP ACTIVITY` moves the activity to a lower priority service subclass. The activity can be remapped in response to processor time consumption again until it executes in the lowest priority service subclass where it will continue until it completes or you intervene manually. Other activities which do not exceed the thresholds continue to run in the high priority service subclass, where they receive higher agent priority.

An event monitor record is logged every time an activity is remapped, if you created a threshold violations event monitor. If you want to collect additional information about remapped activities to investigate further, you can add the `COLLECT ACTIVITY DATA` clause to the `ALTER THRESHOLD` statement in the `wlmtiersdefault.db2` script. Simply rerun the script for the change to take effect.

After running the workload for a period of time, you can use the `WLM_GET_SERVICE_SUBCLASS_STATS_V97` table function to see how many activities were remapped between the service subclasses:

```
SELECT substr(service_superclass_name,1,21) AS superclass,
       substr(service_subclass_name,1,21) AS subclass,
       substr(char(coord_act_completed_total),1,10) AS completed,
       substr(char(act_remapped_in),1,10) AS remapped_in,
       substr(char(act_remapped_out),1,10) AS remapped_out,
       substr(char(last_reset),1,19) AS last_reset
FROM table( WLM_GET_SERVICE_SUBCLASS_STATS_V97(
           CAST(NULL AS VARCHAR(128)),
           CAST(NULL AS VARCHAR(128)),
           -2 )
          ) AS TF_subcls_stats@

SELECT SUBSTR(WORKLOAD_NAME,1,19) AS WL_NAME,
       COORD_ACT_LIFETIME_AVG,
       COORD_ACT_LIFETIME_STDDEV
FROM TABLE(WLM_GET_WORKLOAD_STATS_V97
           (CAST(NULL AS VARCHAR(128)), -2))
 AS WLSTATS
ORDER BY WL_NAME@
```

SUPERCLASS	SUBCLASS	COMPLETED	REMAPPED_IN	REMAPPED_OUT	LAST_RESET
SYSDEFAULTSYSTEMCLASS	SYSDEFAULTSUBCLASS	0	0	0	2008-10-06-20.53.47
SYSDEFAULTMAINTENANCE	SYSDEFAULTSUBCLASS	3	0	0	2008-10-06-20.53.47
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	0	0	0	2008-10-06-20.53.47
WLM_TIERS	SYSDEFAULTSUBCLASS	0	0	0	2008-10-06-20.53.47
WLM_TIERS	WLM_SHORT	999	0	35	2008-10-06-20.53.47
WLM_TIERS	WLM_MEDIUM	19	35	16	2008-10-06-20.53.47
WLM_TIERS	WLM_LONG	16	16	0	2008-10-06-20.53.47

7 record(s) selected.

If you notice that no or only very few activities are being remapped to the lower priority service subclasses, decrease the `CPUTIMEINSC` threshold value and the check interval used by the `ALTER THRESHOLD` statements in the script to improve the mapping of activities across service class tiers according to business priority. If most or almost all activities are being remapped to the lower priority service subclasses, increase the `CPUTIMEINSC` threshold value and the check interval for the `ALTER THRESHOLD` statements to permit more activities to complete with higher priority. After your changes are complete, rerun the `wlmtiersdefault.db2` script to make them effective.

Scenario: Remapping incorrectly mapped queries through priority aging

The following scenario shows how you can configure your data server to dynamically remap, or age the priority of, activities that are consuming more processor time than originally estimated in order to maintain system performance for other queries.

The problem: You may have mapped expensive activities based on estimated SQL cost to a lower priority service subclass so that these activities do not impact the performance of less expensive, shorter activities. Such a mapping can be accomplished by defining a work action set at the service superclass level. However, if the estimated SQL cost is incorrect because of statistics that are out of date, for example, an expensive activity might be mapped incorrectly to a high priority service subclass where it begins to consume an excessive amount of resources, at the cost of all other high priority activities.

The solution: You can use the `wlmtierstimerons.db2` sample tiering script to configure your data server with a tiered configuration that evaluates incoming activities according to their estimated cost and maps them to one of three service subclasses, each with different agent priorities. If an activity consumes too much processor time, your data server dynamically lowers the priority of the activity during its lifetime by remapping it between performance tiers. This dynamic process of remapping activities to lower their priority is also referred to as priority aging.

After an activity has been mapped to its initial service class and begins executing, the `CPUTIMEINSC` in-service-class threshold is used by the script to control the amount of processor time an activity can consume. If the activity violates the threshold by using the maximum amount of allowed processor time, a `REMAP ACTIVITY` action is triggered which moves the activity to a service subclass with lower agent priority. The activity can be remapped in response to processor time consumption until it executes the lowest priority service subclass where it will continue until it completes or you intervene manually.

An event monitor record is logged every time an activity is remapped. If you want to collect additional information about remapped activities to investigate further, you can add the `COLLECT ACTIVITY DATA` clause to the `ALTER THRESHOLD` statement in the `wlmtiersdefault.db2` script. Simply rerun the script for the change to take effect.

After running the workload for a period of time, you can use the `WLM_GET_SERVICE_SUBCLASS_STATS_V97` table function to see how many activities were remapped between the service subclasses:

```
SELECT substr(service_superclass_name,1,21) AS superclass,
       substr(service_subclass_name,1,21) AS subclass,
       substr(char(coord_act_completed_total),1,10) AS completed,
       substr(char(act_remapped_in),1,10) AS remapped_in,
       substr(char(act_remapped_out),1,10) AS remapped_out,
       substr(char(last_reset),1,19) AS last_reset
FROM table( WLM_GET_SERVICE_SUBCLASS_STATS_V97(
           CAST(NULL AS VARCHAR(128)),
           CAST(NULL AS VARCHAR(128)),
           -2 )
         ) AS TF_subcls_stats@

SELECT SUBSTR(WORKLOAD_NAME,1,19) AS WL_NAME,
       COORD_ACT_LIFETIME_AVG,
       COORD_ACT_LIFETIME_STDDEV
```

```

FROM TABLE(WLM_GET_WORKLOAD_STATS_V97
             (CAST(NULL AS VARCHAR(128)), -2))
AS WLSTATS
ORDER BY WL_NAME@

```

SUPERCLASS	SUBCLASS	COMPLETED	REMAPPED_IN	REMAPPED_OUT	LAST_RESET
SYSDEFAULTSYSTEMCLASS	SYSDEFAULTSUBCLASS	0	0	0	2008-10-06-20.59.27
SYSDEFAULTMAINTENANCE	SYSDEFAULTSUBCLASS	3	0	0	2008-10-06-20.59.27
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	0	0	0	2008-10-06-20.59.27
WLM_TIERS	SYSDEFAULTSUBCLASS	0	0	0	2008-10-06-20.59.27
WLM_TIERS	WLM_SHORT	651	0	5	2008-10-06-20.59.27
WLM_TIERS	WLM_MEDIUM	36	5	7	2008-10-06-20.59.27
WLM_TIERS	WLM_LONG	16	7	0	2008-10-06-20.59.27

7 record(s) selected.

For this scenario, you should see relatively few activities being remapped between service subclasses, because activities should almost always be mapped to the appropriate service subclass initially, based on estimated cost. If you notice that activities typically are being completed only in the WLM_SHORT or the WLM_LONG service class, you can adjust the estimated cost values used by the ALTER WORK CLASS SET statement in the script to improve the mapping of activities across service class tiers, so that shorter activities are mapped to the WLM_SHORT_DML_WC work class and longer activities are mapped to the WLM_MEDIUM_DML_WC or the WLM_LONG_DML_WC work class. If you notice that most of the activities are being remapped, you can increase the threshold values used in the ALTER THRESHOLD statements to improve the initial mapping of activities to service subclasses. After your changes are complete, rerun the wlmstimerons.db2 script to make them effective.

Remapping activities between service subclasses

You enable remapping by specifying a REMAP ACTIVITY action with CREATE and ALTER THRESHOLD statements. The remap action, when it is triggered by a threshold violation, moves an activity from one service subclass to another service subclass under the same superclass.

Before you begin

In order to be able to remap to another service subclass, the target service subclass must exist under the same service superclass as the original service subclass of the activity. Either the target or original service subclass can be the default subclass of the superclass. The REMAP ACTIVITY action cannot be applied to service subclasses under the default system class, default maintenance class or default user class.

About this task

The REMAP ACTIVITY action will move an activity to a different service subclass within the same service superclass. Remapping is available with any of the in-service-class thresholds such as CPUTIMEINSC and SQLROWSREADINSC. You use this dynamic process of remapping activities to lower their priority over time, which is also known as priority aging. Lowering the priority of some activities over time can free up system resources, which can then be applied to other activities of higher business importance.

Agents working for the activity will periodically check if a threshold has been violated on each partition, without coordination between partitions. When any one agent detects an in-service-class threshold violation on a partition, this agent triggers the REMAP ACTIVITY action for the activity on the partition and then remaps itself to the target service subclass, after which the activity is considered

remapped. All other agents working for the activity on the same partition will remap to the target service subclass when they detect that the activity has been remapped.

Restrictions

The target service subclass cannot be the same as the original service subclass; you must remap to a different service subclass first before remapping to the original one.

If an activity is remapped to a disabled service subclass, the activity is treated as if it has been rejected by the disabled subclass and an error message -4714 is returned to the client.

Procedure

1. Identify those activities which you want to control through priority aging. A tiered setup consists of service subclasses under the same service superclass that have in-service-class thresholds with REMAP ACTIVITY threshold actions defined on them. You can refer to the basic priority aging example and to the sample scenarios as a starting point:
 - a. "Priority aging of ongoing work" on page 113
 - b. "Scenario: Controlling resource intensive business intelligence reports with priority aging" on page 120
 - c. "Scenario: Remapping incorrectly mapped queries through priority aging" on page 122
2. Select the different service subclasses that activities will be mapped to. This includes both the the service subclass an activity is originally mapped to when activity execution starts, and any other service subclass or classes the activity will be remapped to during its lifetime. For more information on service classes, including on how to create them, see "Resource assignment with service classes" on page 63.
3. Create or alter your thresholds for controlling activities. For more information on thresholds, see "Priority aging of ongoing work" on page 113.
 - a. Define your in-service-class thresholds so that they include a REMAP ACTIVITY action, which is triggered when the threshold is violated. Note that an in-service-class threshold applies to and is affected by an activity only while the activity is mapped to the associated service subclass; affected counters and timers are reset after remapping. Consider if a threshold violation record should be logged each time an activity remaps. These records provide information about which service classes an activity spent time executing in, which you can use for performance analysis. Note that logging threshold violation records can begin consuming substantial amounts of disk space, if remapping of activities between service subclasses is a common occurrence.
 - b. You can also define any thresholds that you want to apply to the lifetime of the activity, but note that only thresholds from the first service subclass that the activity is originally mapped to continue to apply throughout the lifetime of the activity. If you also define any thresholds on any of the service subclasses that an activity is later remapped to, they do not apply.
4. Commit your changes. When you commit your changes, your thresholds are added to the SYSCAT.THRESHOLDS view.
5. Allow your data server to execute the activities you are targeting with your in-service-class thresholds and monitor their progress during their lifetime.

Activities will stay in their original service subclass during execution as long as they do not violate an in-service-class threshold. As in-service-class thresholds are violated during activity execution, activities will trigger a REMAP ACTIVITY action that dynamically remaps them to different service subclasses. Once remapped, the activities continue execution and are now controlled by the resource constraints you have placed on the target service subclass.

6. If necessary, refine your approach to priority aging to reach your stated performance goals.

Example

The following example creates a simple three-tiered setup that lowers, or ages, the priority of ongoing activity over time. Three service subclasses under a single superclass A provide the execution environment in which all queries must run. Assume that the default user workload maps incoming queries to service subclass A1, which is a high-priority subclass intended to permit shorter running queries to execute quickly. A medium-priority service subclass A2 is intended to permit longer running queries to execute, although with more stringent resource controls. Service subclass A3 provides containment for any very large queries that take an excessive amount of processor time to complete.

Three thresholds provide control over the ongoing resource consumption of queries. A query is permitted to execute in the high-priority service subclass A1 only as long as it requires less than one minute of processor time to complete. After a minute of processor time has been consumed, threshold T1 automatically remaps the activity to subclass A2, where it can continue executing as long as it consumes less than 10 minutes of processor time. If the query still has not completed after 10 minutes of consuming processor time, threshold T2 remaps the activity to the lowest priority service subclass, A3. Queries in subclass A3 are permitted to continue indefinitely, although an event monitor record is logged and activity data with details is collected when the processor time used exceeds 1 hour.

```
CREATE SERVICE CLASS A
CREATE SERVICE CLASS A1 UNDER A
CREATE SERVICE CLASS A2 UNDER A
CREATE SERVICE CLASS A3 UNDER A

CREATE THRESHOLD T1 FOR SERVICE CLASS A1 UNDER A
  ACTIVITIES ENFORCEMENT DATABASE PARTITION
  WHEN CPUTIMEINSC > 1 MINUTE REMAP ACTIVITY TO A2

CREATE THRESHOLD T2 FOR SERVICE CLASS A2 UNDER A
  ACTIVITIES ENFORCEMENT DATABASE PARTITION
  WHEN CPUTIMEINSC > 10 MINUTES REMAP ACTIVITY TO A3

CREATE THRESHOLD T3 FOR SERVICE CLASS A3 UNDER A
  ACTIVITIES ENFORCEMENT DATABASE PARTITION
  WHEN CPUTIMEINSC > 1 HOUR LOG EVENT MONITOR RECORD
  COLLECT ACTIVITY DATA WITH DETAILS
  CONTINUE
```

Apply controls to types of activities with work action sets

Work action sets contain work actions that apply controls to activities of a certain type in either a specific service superclass or to the database as a whole.

A work action provides an action that can be applied to a work class, which represent activities of a certain type like LOAD or READ activities.

If you apply a work action set to a database, there are several types of actions that you can apply to activities that fall within a work class, such as threshold definitions, prevent execution, collect activity data, and count activity. Defining a threshold for a work action is the most powerful database work action. For example, perhaps you want to prevent SQL from reading and returning more than 100 000 rows. You can define a single work class for a work action set that identifies SQL READ statements and a work action with a threshold that would stop execution if the number of rows returned is more than 100 000. For information about the possible actions, see “Work actions and the work action set domain” on page 130.

If you define the work action set for a service superclass, the different types of actions that you can apply to activities include mapping activities to a service subclass, preventing execution, collecting activity or aggregate activity data, and counting the activities. Typically, the work action maps an activity to a service subclass and has thresholds defined on the subclass to help manage the activity.

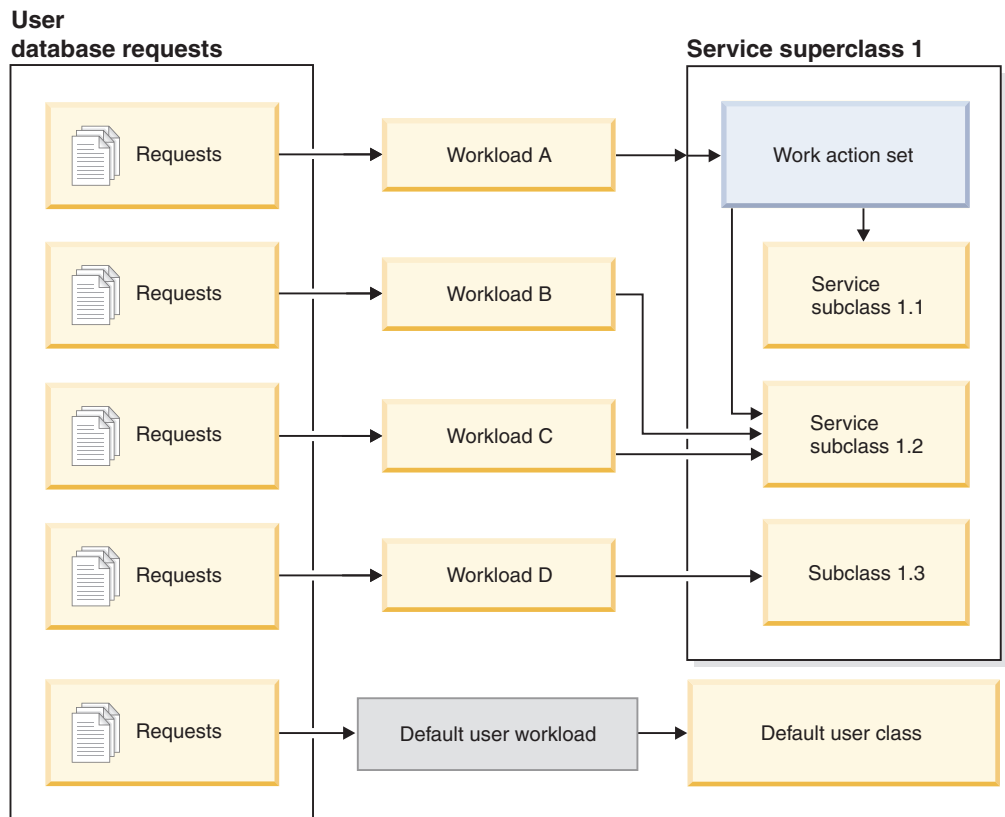


Figure 17. Work action set mapping for a service superclass

How work classes, work class sets, work actions, and work action sets work together and are associated with other DB2 objects

Work classes and work actions work together to apply specific actions to specific activity types. The best way to describe how this works is through an example.

The following diagram shows a high-level view of how work classes, work class sets, work actions, and work action sets work together and are associated with

other DB2 objects.

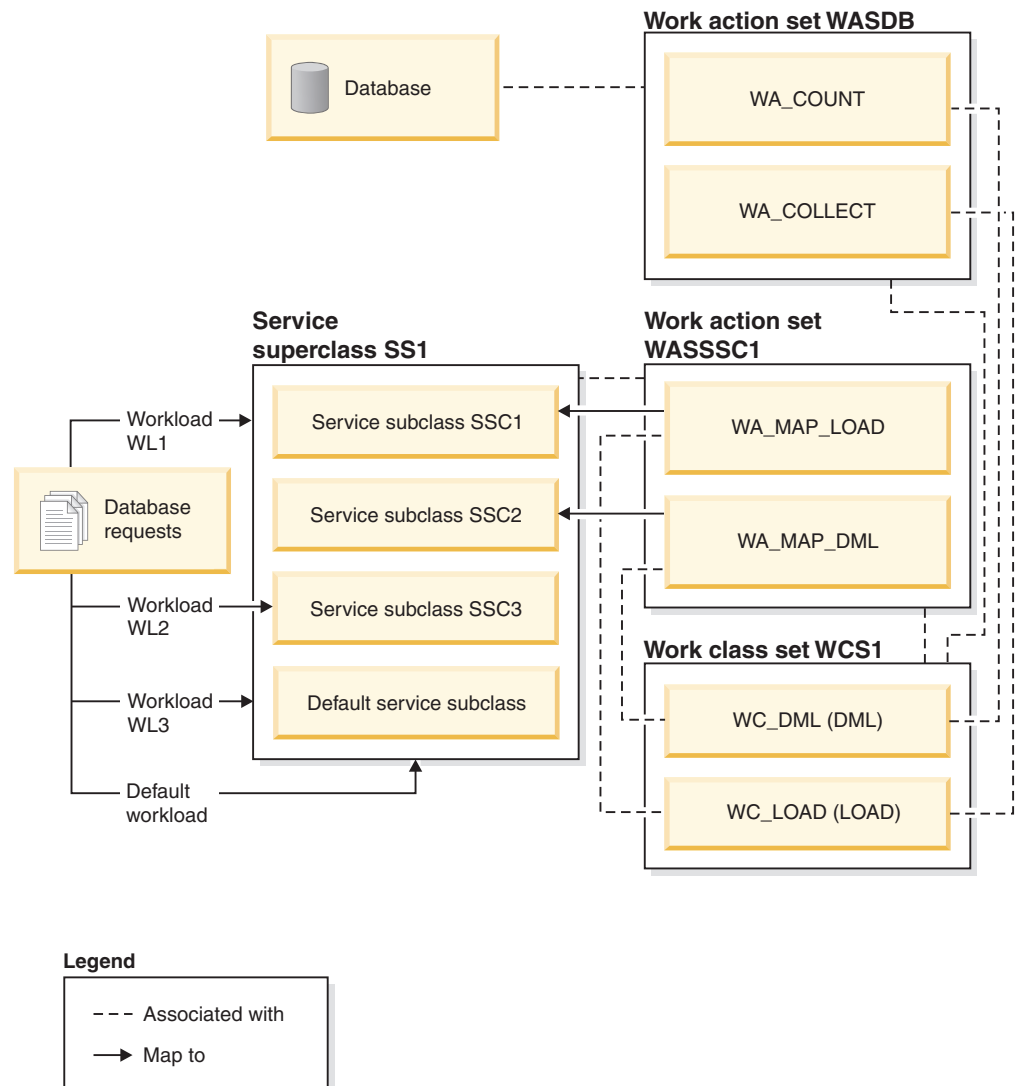


Figure 18. Overview of work action sets and work class sets

In the diagram, some database activities are mapped, through workload WL1, workload WL3, and the default user workload, SYSDEFAULTUSERWORKLOAD, to the service superclass SS1. Because work action set WASDB is applied to the database, any activities that are assigned to the default user workload, the WL1 workload, or the WL3 workload and fall under the WC_DML or WC_LOAD work classes will have the work actions in the WASDB work action set applied to them. That is, activities with the DML work type are counted, and activities with the LOAD work type have activity data collected for them and written to an active event monitor (if one is available).

The work action set WASSSC1 is applied to the service superclass SS1. Any activities that are assigned to the default user workload, the WL1 workload, or the WL3 workload and fall under the WC_DML work class and the WC_LOAD work class will also have the WA_MAP_DML and WA_MAP_LOAD work actions applied to them. That is, activities with a work type of LOAD will be mapped to

the SSC1 service subclass by the WA_MAP_LOAD work action, and activities with a work type of DML will be mapped to the SSC2 service subclass by the WA_MAP_DML work action.

Activities that are assigned to the WL2 workload are mapped directly to a service subclass (SSC3). When a workload maps activities directly to a service subclass, no work actions are applied to those activities.

Work actions and work action sets

A work action, when used in conjunction with a work class, can be used to help control specific types of activities. For example, you can apply different work actions to LOAD activities so that they are processed differently than DML. Work actions are grouped into work action sets.

Work actions

A work action consists of the following attributes:

- A user-supplied work action name, which must be unique in the work action set.
- The work class identifier the work action is to be applied to. You can define more than one work action for a work class, but each work action must perform a different action on that work class.
- The action that is to be applied to the database activity that matches the work class. The valid action type for a work action depends on whether the work action set that the work action belongs to is applied to a database or a service superclass. When a work action set is applied to a database (depending on the work classes that the work actions are associated with), the work action set applies to some or all activities that enter the database. When a work action set is applied to a service superclass (depending on the work classes that the work actions are associated with), that work action set applies to some or all activities that are run under that service superclass. For example:
 - A work action set that is applied to a database can contain threshold work actions. If an activity gets assigned to a work class that has a threshold work action defined for it, the threshold is applied to that activity.
 - A work action set that is applied to a service superclass can contain a work action that maps the activity to a service subclass in the service superclass. If an activity corresponds to a specific work class in a work class set, and the work action set has a mapping work action that is defined for that work class, that activity is mapped to the service subclass specified by the work action.

For a list of the supported actions, see “Work actions and the work action set domain” on page 130.

- An object that is the target of the specified action. Depending on the action, the object can be a service subclass that the activity is mapped to, a threshold that specifies which threshold to apply to the activity, or null if the action is to prevent execution, one of the collect actions, or count activity.
- The template describing the histogram that collects statistical information about the number of microseconds that activities associated with the work class to which this work action is assigned required to run during a specific interval. This information is only collected when the work action type is COLLECT AGGREGATE ACTIVITY DATA (either BASE or EXTENDED). For more information on histograms and histogram templates, see “Histograms in workload management” on page 180.
- Whether or not the work action is enabled.

- An automatically generated identifier that identifies the work action.

You can create a work action by using either the `WORK ACTION` keyword in the `CREATE WORK ACTION SET` statement or the `ADD` keyword in the `ALTER WORK ACTION SET` statement. You can alter a work action by using the `ALTER` keyword in the `ALTER WORK ACTION SET` statement. You can remove a work action from a work action set by using the `DROP` keyword in the `ALTER WORK ACTION SET` statement, or by dropping the entire work action set.

You can view your work actions by querying the `SYSCAT.WORKACTIONS` view.

Work action sets

A work action set consists of the following attributes:

- A work action set name that is unique in a database.
- The name of the work class set containing one or more work classes that the group of actions is to apply to.
Because the definitions of the work class sets are separate from the work action sets defined for them, you can define more than one work action set for a work class set.
- The type of object that the work action set is associated with (database or service superclass).
- The name of the service superclass that the actions and work class set apply to (for work action sets associated with a service superclass).
- Whether or not the work action set is enabled.
- User comments.
- One or more work actions (a work action set does not have to contain any work actions).
- An automatically generated ID that uniquely identifies the work action set.

You can create a work action set using the `CREATE WORK ACTION SET` statement, alter a work action set using the `ALTER WORK ACTION SET` statement, and drop a work action set using the `DROP WORK ACTION SET` statement.

You can view your work action sets by querying the `SYSCAT.WORKACTIONSETS` view.

When you create a work action set, you must specify the object that the work action set is to be applied to. The valid object types are the database or a service superclass. You must also specify which work class set the work action set is to work with. This permits you to use the work classes in the work class set to identify the types of activities that you want to apply the work actions to.

If you set up a workload to map its database activities directly to a service subclass, the work action set associated with that service superclass is never used for the activities issued by that workload. In other words, if a workload maps activities directly to a service subclass, the work action set is bypassed. None of the work actions in the work action set will be applied to the activities that are mapped directly to the service subclass.

Work actions and the work action set domain

You can define a work action set for either a database or a service superclass. The type of work actions that can be defined for a work action set depends on the type of object the work action set is defined for.

If the work action set is defined for a database, the work actions in the work action set must be one of the following actions:

- A threshold
The actual threshold is specified by the `WHEN threshold-type` keyword. Multiple threshold work actions can be applied to a single work class if all the thresholds are of different types. If this action is specified, the threshold is applied to all database activities associated with the work class.
- `PREVENT EXECUTION`
If this action is specified, all database activities that match the associated work class are not permitted to run.
- `COLLECT ACTIVITY DATA`
If this action is specified, information about the database activities corresponding to the work class for which this work action is defined are written to the active `ACTIVITIES` event monitor when the activities complete execution. See “Collecting data for individual activities” on page 195 for more information.
- `COUNT ACTIVITY`
If this action is specified, all database activity that maps to the associated work class causes the turnstile counter for that work class type to be incremented. (The turnstile counter for the work class is incremented by 1 each time an activity is associated with that work class). The `COUNT ACTIVITY` work action provides an efficient way to ensure this counter is updated. If no work action is applied to an activity corresponding to a work class, the work class activity counter is not incremented. Sometimes the only action you care about is obtaining a count of activities of a given type. See “Collecting data for individual activities” on page 195 for more information.

If the work actions in the work action set are not any of these actions, `SQL4720N` is returned.

If you are defining a work action set for a service superclass, the work actions in the work action set must be one of the following actions:

- A mapping action
You can map an activity to any service subclass in the service superclass except for the default service subclass. You specify the service subclass to map the activity to using the `MAP ACTIVITY TO SERVICE CLASS` keyword. Only one map work action in the work action set can be applied to the same work class.
- `PREVENT EXECUTION`
Behavior is the same as for the database work action.
- `COLLECT ACTIVITY DATA`
Behavior is the same as for the database work action.
- `COLLECT AGGREGATE ACTIVITY DATA`
If this action is specified, aggregate database activity data that corresponds to the work class for which this work action is defined is collected.
- `COUNT ACTIVITY`
Behavior is the same as for the database work action.

If the work actions in the work action set are not any of these actions, SQL4720N is returned.

The following figure shows an example of how the work classes in a work class set called LARGE ACTIVITIES are to be applied to both the database and a service superclass. To meet this objective, two work action sets, Database large activities and Service class large activities are created.

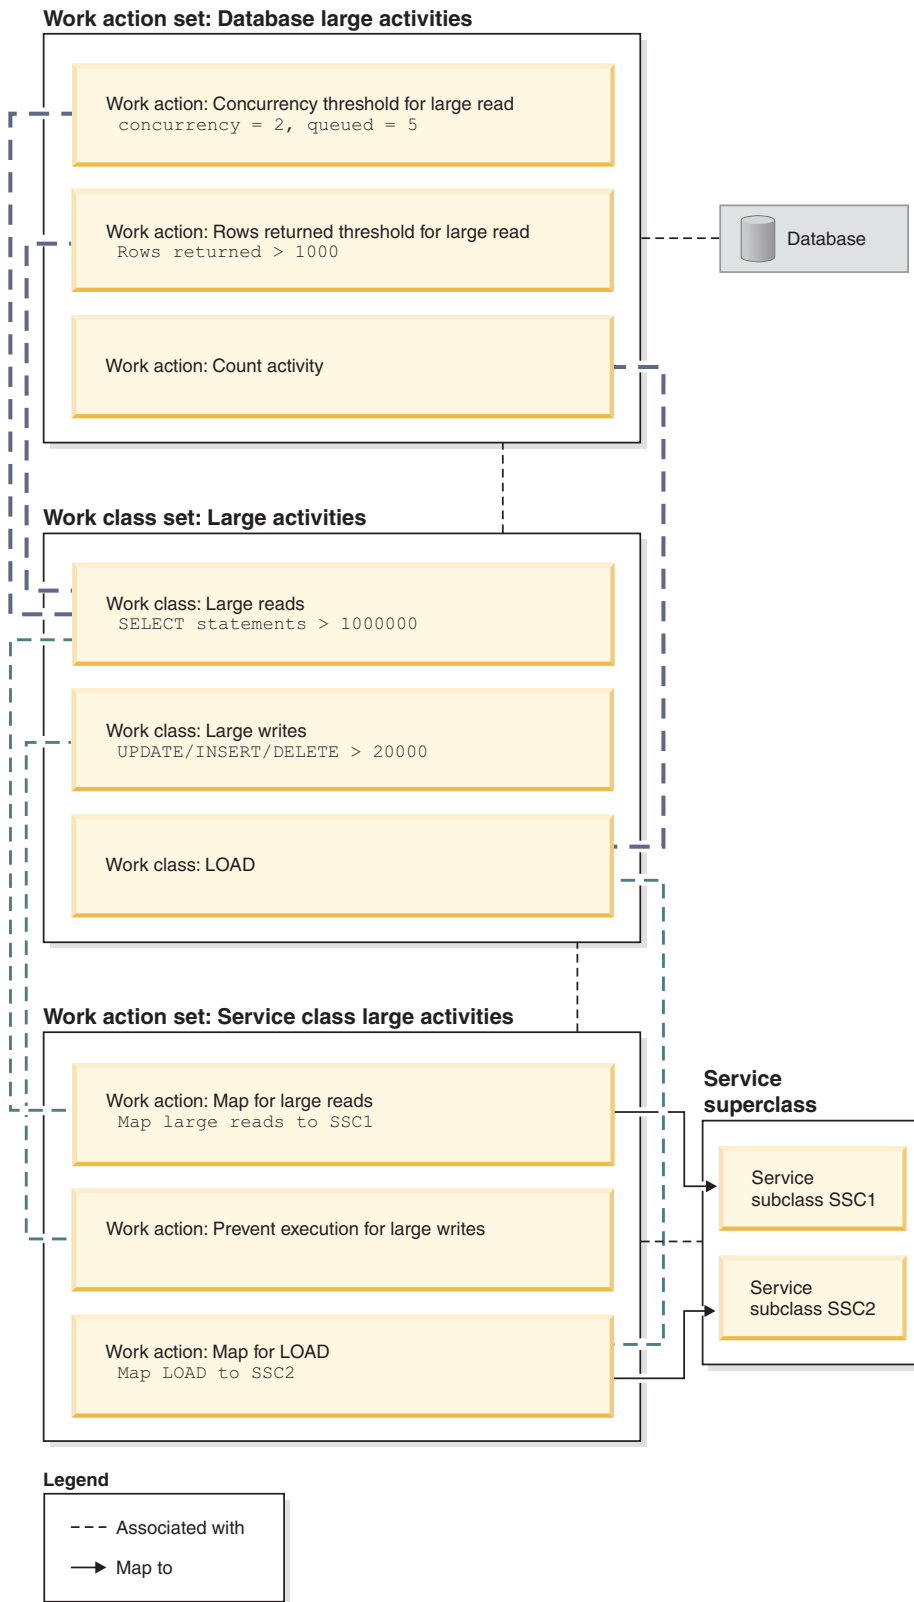


Figure 19. Example of work actions, work actions sets, work classes, and work class set

The work action sets are as follows:

- Database large activities contains:

- Concurrency threshold for large reads, which permits two large reads to run concurrently, and five large reads to be queued
- Rows returned threshold for large reads, which prevents large reads from returning more than 1000 rows
- Count activity for load, which counts the number of times the load utility runs on the database.
- Service class large activities contains:
 - Map for large reads, which maps large reads to service subclass 1
 - Map for large writes, which prevents large writes from executing.
 - Map for LOAD, which maps loads to service subclass 2

A work action set does not have to contain an action for every work class in the work class set to which the work action set is applied. In addition, a work class can have more than one work action applied to it as long as the action types are different. A work class can have more than one threshold work action applied to it as long as the threshold types are different.

Thresholds that can be used in work actions

Work action sets that you define for databases can contain work actions that specify thresholds.

The following thresholds are supported:

- Aggregate threshold:
 - CONCURRENTDBCOORDACTIVITIES
- Activity thresholds:
 - SQLTEMPSPACE
 - SQLROWSRETURNED
 - ACTIVITYTOTALTIME
 - ESTIMATEDSQLCOST
 - CPUTIME
 - SQLROWSREAD

Application of work actions to database activities

One, and only one work action set can be applied to either a database or a service superclass.

When work is submitted to the data server, it is associated with a workload, either a user-defined workload or the default workload, then mapped to a service class.

The following figure shows the process of how a work action is applied to an activity.

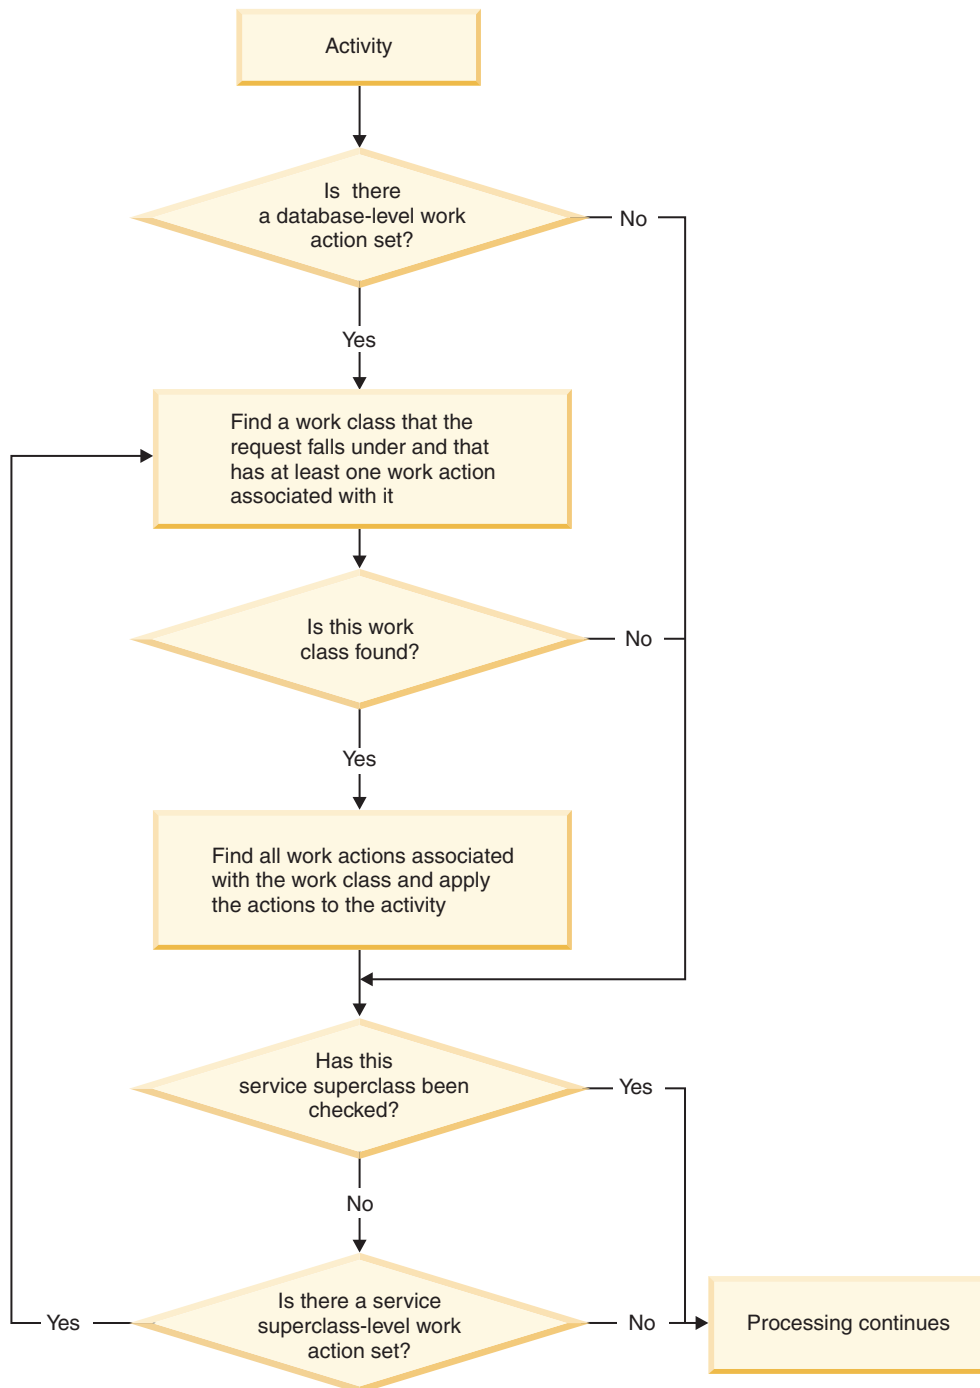


Figure 20. Application of a work action to an activity

A work action is assigned to an activity as follows:

1. When an activity is mapped to a service superclass or a service subclass, the data server checks whether an enabled database-level work action set exists.
2. If an enabled database-level work action set exists, the data server then checks whether the activity falls under any of the work classes in the work class set that the database-level work action set is associated with.
3. If the activity falls under a work class and that work class has any work actions applied to it, those work actions are applied to the activity.

4. Next, if the activity is mapped by the workload to a service superclass, the data server checks whether a work action set is applied to the service superclass.
5. If a work action set is applied to the service superclass, the data server then checks whether the activity falls under any of the work classes in the work class set that the service superclass-level work action set is associated with.
6. If the activity falls under a work class and that work class has any work actions applied to it, those work actions are applied to the activity.

Note that, if a mapping work action is applied to a stored procedure, depending on its configuration, child activities of a stored procedure can run in the same service subclass or different service subclasses than the parent activity.

In the following situations an activity is not affected by a work action set:

- Activities fall in the default system (SYSDEFAULTSYSTEMCLASS) and default maintenance (SYSDEFAULTMAINTENANCECLASS) service classes.
- Activities are assigned to the default administration workload, SYSDEFAULTADMWORKLOAD.
- Activities are inside a load operation. The load operation itself does go through work action set evaluation.
- Child activities of system stored procedures. The only exception is the SYSPROC.ADMIN_CMD stored procedure. Child activities of SYSPROC.ADMIN_CMD go through work action set evaluation.
- The work action set is disabled.
- The workload maps the activity directly to a service subclass.

Workload and work action set comparison

Depending on the type of control that you want to maintain over your database activities, you can use workloads by themselves or both workloads and work classes (when used with work actions) to map activities to service classes.

With workloads, requests are identified and assigned to a service class based on connection attributes. Workloads are the primary method for routing work to a specific DB2 service class for execution. If you want to further refine how requests are identified, you can use work classes to classify the activities based on their type and other activity attributes. For example, you can classify READ activities, WRITE activities, and LOAD activities into different work classes and have each activity type treated differently.

If you use work classes (which are grouped into work class sets), you can use work actions to exercise control over the different types of activities. For example, you can use one work action to map a specific type of activity to a service subclass and use a different work action to apply a control known as a threshold to ensure that same type of activity does not exceed certain conditions.

Work actions are grouped into work action sets. A single work action set can apply to activities in the database or to activities in a service superclass (but not both). Work class sets and work action sets work together. That is, a work class must exist for categorizing an activity as a specific type of work before a work action can be applied to it. A work class set can be associated with more than one work action set, but a work action set can be associated with only one work class set.

Figure 1 shows an example of a DB2 workload manager implementation that uses workloads and work action sets. In this figure, assume that a request is assigned to

workload WL_A based on the connection attributes of the connection that submitted the request. Workload WL_A specifies that the request is to be executed in service superclass SC_A. Assume that a work class in work class set WCS_1 matches the type of work that the request that is associated with workload WL_A is going to perform.

Now assume that an activity that does not update the catalogs (a READ activity) enters the system. The database-level work action set WAS_1 (that is associated with work class set WCS_1) contains a work action that is applied to the READ work class. The request is then mapped to service superclass SC_A (by workload WL_A). Here, the request encounters the service superclass-level work action set WAS_2, which is also associated with work class set WCS_1, and applies to activities in service superclass SC_A. This work action set contains a mapping work action, which is also applied to the READ work class so that all READ activities will be mapped to service subclass SSC_1a in service superclass SC_A.

A somewhat similar situation occurs with the request that is associated (again, based on its connection attributes) with workload WL_B. Workload WL_B maps activities to service superclass SC_B. Assume that the request is for a LOAD activity and that work class set WCS_2 contains a work class that applies to LOAD activities. Work class set WCS_2 is associated with the service superclass-level work action set WAS_3, which applies to activities in service superclass SC_B. Assume that work action set WAS_3 contains a mapping work action that is applied to the LOAD work class, so that when the LOAD activity is mapped to service superclass SC_B by workload WL_B, it will then be mapped by the work action to service subclass SSC_1b for execution.

The purpose of workload WL_C in this example is to map incoming requests directly to service subclass SSC_1b, independent of the service superclass-level work action set WAS_3 and its mapping work action. If an incoming request is associated with workload WL_C that is a LOAD activity, then this request is also mapped directly to service subclass SSC_1b for execution, and is unaffected by the mapping work action that applies to the LOAD work class.




Figure 21. Workloads and work action sets

Creating a work action set

To create a work action and a work action set, use the CREATE WORK ACTION SET statement.

To create a work action set, you require WLMADM or DBADM authority.

For additional prerequisites, see the following topics:

- “DDL statements for DB2 workload manager” on page 16
- Naming rules

When you create a work action set:

- You associate it with a work class set. The work class set must already exist.
- You also associate it with the database or a service superclass. If you are associating the work action set with a service superclass, the service class must already exist. You cannot define the work action set for the default system service class (SYSDEFAULTSYSTEMCLASS), the default user class (SYSDEFAULTUSERCLASS) or the default maintenance service class (SYSDEFAULTMAINTENANCECLASS).

To create a work action set:

1. Use the CREATE WORK ACTION SET statement with the following options:
 - Specify a name for the work action set. The name of the work action set must be unique in the database.
 - Specify the object with which the work action set is associated. You can specify a database or service superclass. If you specify that the work action set is associated with a database, none of the work actions in the work action set can be mapping work actions or collect aggregate actions. If you specify that the work action set is associated with a service superclass, none of the work actions in the work action set can be thresholds. For example, to apply the work action set to the REPORTS service superclass, you would specify:
FOR SERVICE CLASS REPORTS
To apply the work action set to the database, you would specify:
FOR DATABASE
 - Specify the work class set with which the work action set is associated. The work classes in the work class set classify the database activities that the work actions in the work action set will apply to. For example, to associate the work action set with the LARGEREADS work class set, you would specify:
USING WORK CLASS SET LARGEREADS
 - Optional: Create one or more work actions for the work action set. For instructions, see “Creating a work action” on page 140.
 - Optional: Specify whether the work action set is enabled or disabled. By default, the work action set is enabled. If the work action set is disabled, the data server does not consider this work action set (or any work actions in it) when activities are run.
2. Commit your changes. When you commit your changes the work action set is added to the SYSCAT.WORKACTIONSETS view.
A new work action set only takes effect in the database after it is committed, and does not affect any database activities currently running.

Altering a work action set

To add, alter, or drop a work action from a work action set, or to enable or disable the work action set, use the ALTER WORK ACTION SET statement.

To alter a work action set, you require SQLADM, WLMADM, or DBADM authority. To specify any clause other than a COLLECT clause, the authorization id must include WLMADM or DBADM authority.

For additional prerequisites, see the following topics:

- “DDL statements for DB2 workload manager” on page 16
- Naming rules

When you create a work action set to work with a specific work class set, you cannot change it to work with a different work class set because the work actions in the work action set have a dependency on the work classes in the work class set. If you want to change the work class set this work action set is to be applied to, you must drop and recreate the work action set.

You cannot change which object the work action set applies to because the type of work actions in the work action set depends on which object (database or service superclass) the work action set is defined for. If you want to change the which object the work action set is associated with, you must drop and recreate the work action set.

To alter a work action set:

1. If you want to add a new work action to the work action set, use the ADD keyword. For information about the parameters that you can specify when adding a work action to a work action set, see “Creating a work action” on page 140
2. If you want to alter an existing work action, use the ALTER keyword. For information about altering a work action, see “Altering a work action” on page 143.
3. If you want to drop a work action, use the DROP keyword. For information about dropping a work action from a work action set, see “Dropping a work action” on page 145.
4. You can enable a work action set that is not currently enabled, and the reverse. If you disable an enabled work action set, the data server ignores it after you commit your changes. For more information, see “Disabling a work action set.” If you enable the work action set, after you commit your changes, the work action set is applied to the next applicable activity that enters the database.

Note: Disabling a work action set does not disable the work actions within the work action set, but the work action set will no longer affect any work. If you want to drop a work action set that contains a concurrency work action threshold, you must first disable the concurrency work action before the work action set can be dropped, because concurrency thresholds must be disabled before they can be dropped.

5. Commit your changes. When you commit your changes, the work action set is updated in the SYSCAT.WORKACTIONSETS view. The SYSCAT.WORKACTIONS views is updated for any added, altered, or dropped work actions.

Disabling a work action set

To disable a work action set, use the DISABLE keyword of the CREATE WORK ACTION SET statement or the ALTER WORK ACTION SET statement.

To disable a work action set, you require WLMADM or DBADM authority.

Disabling a work action set does not disable the work actions within the work action set, but the work action set will no longer affect any work. At runtime, a disabled work action set is treated as if it does not exist.

Note: If you want to drop a work action set that contains a concurrency work action threshold, you must first disable the concurrency work action before the work action set can be dropped, because concurrency thresholds must be disabled before they can be dropped.

For example, assume that you have a work action set called READACTIVITIES that is associated with a work class set called READCLASSES, and that work action set is defined for a service superclass called READSERVICECLASS. The SMALLREAD work action set has a work action in it that remaps all SELECT statements to the service subclass SMALLREADSERVICECLASS. If the READACTIVITIES work action set is disabled, all SELECT statements are treated as though the READACTIVITIES work action set does not exist, and are mapped to the default service subclass.

To disable a work action set:

1. Use one of the following statements, depending on whether you are creating or altering a work action set:
 - To create a work action set that is disabled:

```
CREATE WORK ACTION SET work-action-set-name ... DISABLE
```
 - To disable an already existing work action set:

```
ALTER WORK ACTION SET work-action-set-name ... DISABLE
```
2. Commit your changes. When you commit your changes, the work action set is updated in the SYSCAT.WORKACTIONSETS view.

Dropping a work action set

Use the DROP WORK ACTION SET statement to drop a work action set.

To drop a work action set, you require WLMADM or DBADM authority.

Dropping a work action set drops the work action set and all work actions in it.

If the work action set contains a CONCURRENTDBCOORDACTIVITIES threshold work action, that work action must first be disabled before the work action set can be dropped.

To drop a work action set:

1. Use the DROP WORK ACTION SET statement.
2. Commit your changes. When you commit your changes the work action set is removed from the SYSCAT.WORKACTIONSETS view. In addition, all work actions that were part of the work action set are removed from the SYSCAT.WORKACTIONS view. If the work action set contains threshold work actions, the thresholds are removed from the SYSCAT.THRESHOLDS view.

Creating a work action

Use the CREATE WORK ACTION SET statement or the ALTER WORK ACTION SET statement to create a work action.

To create a work action, you require WLMADM or DBADM authority.

For additional prerequisites, see the following topics:

- “DDL statements for DB2 workload manager” on page 16
- Naming rules

When you create a work action:

- You associate a work action with a work class. The work class must already exist in the work class set that the work action set is applied to.

- If the work action is a threshold, the work action set must be defined for the database. For the list of supported thresholds for work actions, see “Thresholds that can be used in work actions” on page 133.
- If you are creating a mapping work action, the work action set must be defined for a service superclass. The service subclass being mapped to must already exist in the service superclass this work action set is being defined for. In addition, you cannot specify the default service subclass.
- Only one work action of the same type can be applied to the same work class from the same work action set. Thresholds are the exception. You can apply more than one threshold to a work class, but each threshold must be of a different type.
- If you are creating a collect aggregate activity data work action, the work action set must be defined for a service superclass.

To create a work action:

1. Use the *work-action-definition* keyword of the CREATE WORK ACTION SET statement, or the ADD *work-action-definition* keyword of the ALTER WORK ACTION SET statement. Specify one or more of the following for the work action:
 - A name for the work action. The name of the work action must be unique within the work action set.
 - The name of the work class to which this work action applies. The work class must be one of the work classes in the work class set that the work action set is associated with. For example, to apply this work action to the work class LARGEDML, you would specify:


```
ON WORK CLASS LARGEDML
```
 - The action that is to apply to activities that match the work class for this work action:
 - If the work action set is associated with a service superclass, you can specify the MAP ACTIVITY keyword so that the work action maps activities to a service subclass in the service superclass. By default, mapping work actions cause activities that are nested to be mapped to the same service subclass as its parent. A cursor that has been opened inside a routine is an example of a nested activity.

For example, if you want the work action to map to the service subclass SMALLREAD, and you want all nested activities to be mapped to the same service subclass, you would specify:

```
MAP ACTIVITY TO SMALLREAD
```

You could also specify:

```
MAP ACTIVITY WITH NESTED TO SMALLREAD
```

If you want the work action to map to the service subclass and to not map nested activities to this service subclass, you would specify:

```
MAP ACTIVITY WITHOUT NESTED TO SMALLREAD
```

If you define the work action as WITHOUT NESTED, nested activities are handled according to their activity type instead of automatically being mapped to the same service subclass as the parent activity. For example, if a CALL activity is mapped to service subclass subsc1, and the routine has an open cursor inside it, the open cursor might be mapped to a different service subclass if it falls under another work class that has another mapping work action applied to it.

- If the work action set is associated with a database, you can specify a WHEN keyword to indicate a threshold to apply to the activity, and the action to take if the activity causes the threshold to be violated. You can specify the following thresholds for a work action:
 - ACTIVITYTOTALTIME
 - ESTIMATEDSQLCOST
 - CONCURRENTDBCOORDACTIVITIES and its QUEUEDACTIVITIES keyword.
 - CPUTIME
 - SQLROWSREAD
 - SQLTEMPSPACE
 - SQLROWSRETURNED

If the threshold is violated, you can specify the following actions to be taken:

- Whether activity data is to be collected about the activity that caused the threshold to be violated. If collected, when the activity completes execution, the activity data is written to an active activities event monitor. By default, no data about the activity is collected. If you want to collect data about this activity, you can collect it from the coordinator partition, a specific database partition, or from all database partitions. You have the option of collecting this data with or without details about the statement and its compilation environment. If you want to collect details about the statement and compilation environment, you can also specify that the input data values used in the activity.
- Whether the activity that caused the threshold to be violated is to be permitted to continue running or not. By default, the activity is stopped.

For example, if you want the work action to check for DML statements that have a cost over 2 000 timerons, collect the basic data about this activity when the threshold is violated and continue to run, you would specify:

```
WHEN ESTIMATEDSQLCOST > 2000 COLLECT ACTIVITY DATA CONTINUE
```

- To prevent any activities that correspond to the work class defined for this work action from executing, you can use the PREVENT EXECUTION keyword.
- To count the number of database activities associated with the work class without incurring the additional overhead of another action (such as collecting data or mapping an activity), you can specify the COUNT ACTIVITY keyword.
- To collect activity data for activities that fall under the work class, specify the COLLECT ACTIVITY DATA keyword. If collected, when the activity completes execution, the activity data is written to an active activities event monitor. By default, no data about the activity is collected. If you want to collect data about this activity, you can collect it from the coordinator partition or from all database partitions. If you want to collect activity details such as the statement and the compilation environment information, you can do so by specifying the WITH DETAILS keyword. You can also use the AND VALUES keyword to have input data values (for those activities that have them) sent to the activities event monitor.

For example, assume that you have a work action set that is applied to a service superclass. You want to have activity data for all activities that are assigned to this work action written to the applicable event monitor, including all aggregate activity information, information about the compilation environment, and any input data values. You would specify:

```
COLLECT ACTIVITY DATA ON ALL WITH DETAILS AND VALUES
```

- To collect aggregate activity data for activities that fall under the work class, specify the COLLECT AGGREGATE ACTIVITY DATA keyword. If collected, aggregate activity data is captured and sent to the applicable event monitor. This information is collected periodically on an interval that is specified by the `wlm_collect_int` database configuration parameter.

For example, assume that you have a work action set that is applied to a service superclass. You want to have aggregate activity data for all activities that are assigned to this work action written to the applicable event monitor, including the base data, the activity data manipulation language (DML) estimated cost histogram, and the activity DML inter-arrival time histogram. You would specify

```
COLLECT AGGREGATE ACTIVITY DATA EXTENDED
```

- The histogram templates used by a COLLECT AGGREGATE ACTIVITY DATA work action to describe the histograms created for the corresponding work class. Specifying the histogram templates used by a work action adds the corresponding rows in the SYSCAT.HISTOGRAMTEMPLATEUSE, view which displays the histogram templates referenced by the service class or work action. For example, if you want to collect interarrival statistics for the default interarrival histogram template, you would specify:

```
INTERARRIVALTIME HISTOGRAM TEMPLATE SYSDEFAULTHISTOGRAM
```

For more information on histograms and histogram templates, see “Histograms in workload management” on page 180.

- Whether the work action is enabled or disabled. By default a work action is created as enabled, but you can specify whether it is enabled or disabled by using the ENABLE or DISABLE keyword. If the work action is disabled, the data server does not consider this work action when activities enter the database or service superclass (depending on the object you created the work action set for).
2. Commit your changes. When you commit your changes, the work actions is added to the SYSCAT.WORKACTIONS view. If the work action is a threshold, the threshold is added to the SYSCAT.THRESHOLDS view.

A new work action only takes effect in the database after it is committed, and does not affect any database activities currently running.

Altering a work action

If you need to alter a work action, use the ALTER WORK ACTION SET statement.

To alter a work action, you require SQLADM, WLMADM or DBADM authority. To specify any clause other than a COLLECT clause, the authorization ID must include WLMADM or DBADM authority.

See “DDL statements for DB2 workload manager” on page 16 for additional prerequisites.

To alter a work action:

1. Use the ALTER keyword of the ALTER WORK ACTION SET statement to change one or more of the following characteristics of the work action.
 - You can alter the work class to which the work action is applied. The work class must already exist in the work class set to which the work action set is applied.
 - If the work action maps to a service subclass, you can alter which service subclass the database activity is to be mapped. You can only change the

mapping to a service subclass in the same service superclass. You cannot map to the default service subclass. You can also change whether nested activities in the activity are mapped to the same service subclass or not. For example, if the work action is currently defined as WITH NESTED, you can change this to WITHOUT NESTED. This change would cause the nested activities to be handled according to their activity type instead of automatically being mapped to the same service subclass as the parent activity. For example, if a CALL statement is mapped to service subclass SUBSC1, and the routine has an open cursor inside it, the open cursor might be mapped to a different service subclass if it falls under another work class that has another mapping work action applied to it.

- You can alter the action type specified for the work action (that is, mapping, threshold, prevent execution, count activity, collect actions), but you must alter it to a valid work type. For example, if the work action is to map the activity to a service subclass, you cannot change the work action to a threshold, or the reverse. The reason is because, in this example, the work action set must have been applied to a service superclass in order to have a mapping action and threshold actions are not valid for work action sets applied to service superclasses. If you alter the type of a work action that is a threshold work action or alter the type of work action to a threshold, the following occurs:
 - If the work action was a threshold and has been changed to a non-threshold, the threshold is removed from the SYSCAT.THRESHOLDS view.
 - If the work action was not a threshold and has been changed to a threshold, a new threshold will be created in the SYSCAT.THRESHOLDS view.

Note: If the action is a threshold, you cannot alter the type of threshold to a different threshold. So, for example, if the work action was an SQLROWSRETURNED threshold, you cannot change it to a SQLTEMPSPACE threshold. In addition, you cannot change the work action type of an enabled CONCURRENTDBCOORDACTIVITIES work action threshold.

- You can alter the histogram templates used by a COLLECT AGGREGATE ACTIVITY DATA work action to describe the histograms created for the corresponding work class. Updating the histogram templates used by a work action updates the corresponding rows in the SYSCAT.HISTOGRAMTEMPLATEUSE view, which displays the histogram templates referenced by the service class or work action. For more information on histograms and histogram templates, see “Histograms in workload management” on page 180.
 - You can alter whether you want to enable or disable the work action. By default, work actions are enabled. When enabled, the data server considers the work action for application against the activity that falls under the work class for the work action. If the work action is disabled, the data server ignores it.
2. Commit your changes. When you commit your changes, the work action is updated in the SYSCAT.WORKACTIONS view.

Disabling a work action

You can disable a work action that you do not want applied to a work class. At runtime, the disabled work action is treated as if it does not exist.

To disable a work action, you require WLMADM or DBADM authority.

To disable a work action:

1. Use one of the following statements, depending on whether you are creating or altering a work action set:
 - Use the DISABLE keyword and the ADD keyword of the CREATE WORK ACTION SET statement. For example:

```
ADD WORK ACTION work-action-name ON WORK CLASS work-class-name ... DISABLE
```
 - Use the DISABLE keyword and the ALTER keyword of the ALTER WORK ACTION SET statement. For example:

```
ALTER WORK ACTION work-action-name ... DISABLE
```
2. Commit your changes. When you commit your changes, the work action is updated in the SYSCAT.WORKACTIONS view.

Dropping a work action

If you no longer require a work action, you can drop it from the work action set.

- To drop a work action, you require WLMADM or DBADM authority.
- See “DDL statements for DB2 workload manager” on page 16 for additional prerequisites.

To drop a work action:

1. Use the DROP keyword of the ALTER WORK ACTION SET statement. If you want to drop a CONCURRENTDBCOORDACTIVITIES threshold work action, you must disable the work action in one ALTER WORK ACTION SET operation, commit the change, verify that there are no queued activities, and then drop the threshold in a second ALTER WORK ACTION SET operation.
2. Commit your changes. When you commit your changes, the work action is removed from the SYSCAT.WORKACTIONS view. If the work action is a threshold work action, the threshold is also removed from the SYSCAT.THRESHOLDS view.

An altered work action set and work action only takes effect in the database after it is committed, and does not affect any database activities currently running.

Example: Using a work action set and database threshold

This example shows different approaches to using work action sets and thresholds to control the resources consumed by DB2 activities. Before creating DB2 workload manager objects, you need to understand how they are used.

Assume that you have a work class set called ALLSQL, and it contains the following work classes in this order:

1. SMALLDML, which is for all DML-type SQL statement that have an estimated cost of less than 1 000 timerons
2. MEDDML, which is for all DML-type SQL statements that have an estimated cost between 1 000 and 20 000 timerons
3. LARGEDML, which is for all DML-type SQL statements that have an estimated cost greater than 20 000 timerons
4. ALLDDL, which is for all DDL-type SQL statements
5. ALLACTIVITY, which is for all database activity

The following SQL statements create the work class set and the work classes:


```

CREATE WORK CLASS SET ALLSQL
(WORK CLASS SMALLDML WORK TYPE DML FOR TIMERONCOST FROM 0 TO 1000,
WORK CLASS MEDDML WORK TYPE DML FOR TIMERONCOST FROM 1001 TO 20000,
WORK CLASS LARGEDML WORK TYPE DML FOR TIMERONCOST FROM 20001 TO UNBOUNDED,
WORK CLASS ALLDDL WORK TYPE DDL,
WORK CLASS ALLACTIVITY WORK TYPE ALL)

```

These work classes already have work actions, such as COUNT ACTIVITY, COLLECT, and thresholds (that are not ACTIVITYTOTALTIME thresholds) applied to them.

Assume that you want to permit large DML activities to run for no longer than 5 hours. All other SQL can take no longer than 30 minutes to run. The following two examples show possible methods for accomplishing this objective.

Method 1

One method is to set up a work action with the ACTIVITYTOTALTIME threshold specified for each work class as follows:

Table 41. ACTIVITYTOTALTIME threshold specified for each work class

Work action	Work class applied to	Threshold type and value	Actions
SMALLDMLTIMEALLOWED	SMALLDML	ACTIVITYTOTALTIME < 31 MINUTES	<ul style="list-style-type: none"> • Stop execution • Collect activity data
MEDDMLTIMEALLOWED	MEDDML	ACTIVITYTOTALTIME < 31 MINUTES	<ul style="list-style-type: none"> • Stop execution • Collect activity data
LARGEDMLTIMEALLOWED	LARGEDML	ACTIVITYTOTALTIME < 5 HOURS	<ul style="list-style-type: none"> • Stop execution • Collect activity data
ALLDDLTIMEALLOWED	ALLDDL	ACTIVITYTOTALTIME < 31 minutes	<ul style="list-style-type: none"> • Stop execution • Collect activity data
ALLACTIVITYTIMEALLOWED	ALLACTIVITY	ACTIVITYTOTALTIME < 31 minutes	<ul style="list-style-type: none"> • Stop execution • Collect activity data

The SQL statements for this method are:

```

CREATE WORK ACTION SET WASNICK FOR DATABASE USING WORK CLASS SET WCSNICK
(WORK ACTION SMALLDMLTIMEALLOWED ON WORK CLASS SMALLDML
WHEN ACTIVITYTOTALTIME > 30 MINUTES COLLECT ACTIVITY DATA STOP EXECUTION,
WORK ACTION MEDDMLTIMEALLOWED ON WORK CLASS MEDDML
WHEN ACTIVITYTOTALTIME > 30 MINUTES COLLECT ACTIVITY DATA STOP EXECUTION,
WORK ACTION LARGEDMLTIMEALLOWED ON WORK CLASS LARGEDML
WHEN ACTIVITYTOTALTIME > 5 HOURS COLLECT ACTIVITY DATA STOP EXECUTION,
WORK ACTION ALLDDLTIMEALLOWED ON WORK CLASS ALLDDL
WHEN ACTIVITYTOTALTIME > 30 MINUTES COLLECT ACTIVITY DATA STOP EXECUTION,
WORK ACTION ALLACTIVITYTIMEALLOWED ON WORK CLASS ALLACTIVITY
WHEN ACTIVITYTOTALTIME > 30 MINUTES COLLECT ACTIVITY DATA STOP EXECUTION)

```

Method 2

Another method might be to use only one work class, LARGEDML, then create a work action set for the database that has one work action, LARGEDMLTIMEALLOWED, applied to the work class.

Table 42. LARGEDMLTIMEALLOWED work action applied to the LARGEDML work class

Work action	Work class applied to	Threshold type and value	Action
LARGEDMLTIMEALLOWED	LARGEDML	ACTIVITYTOTALTIME < 5 HOURS	<ul style="list-style-type: none"> • Stop execution • Collect activity data

You would then apply an ACTIVITYTOTALTIME threshold of less than 31 MINUTES to the database. Using this method, only those activities that correspond to the LARGEDML work class have the 5 hour threshold applied to them. Other activities will have the ACTIVITYTOTALTIME database time threshold of less than 31 minutes applied to them.

The SQL statements for this method are:

```
CREATE WORK ACTION SET WASNICK FOR DATABASE USING WORK CLASS SET WCSNICK
(WORK ACTION LARGEDMLTIMEALLOWED ON WORK CLASS LARGEDML
WHEN ACTIVITYTOTALTIME > 5 HOURS COLLECT ACTIVITY DATA STOP EXECUTION)
```

```
CREATE THRESHOLD THTEST FOR DATABASE ACTIVITIES ENFORCEMENT DATABASE
WHEN ACTIVITYTOTAL TIME > 30 MINUTES COLLECT ACTIVITY DATA STOP EXECUTION
```

Example: Using work action sets to determine the types of work being run

Using work class sets, work classes, work action sets, work actions, and some of the DB2 workload manager monitoring features, you can determine the different types of work running on your system, and the distribution of the work.

To accomplish this task, first create a work class set that contains work classes for the different types of work you are interested in. For example, if you want to know how many READ activities, WRITE activities, DDL activities, and LOAD activities are running on your system, you would create a work class set, ACTIVITYTYPES, as in the following example:

```
CREATE WORK CLASS SET ACTIVITYTYPES
(WORK CLASS READWC WORK TYPE READ,
WORK CLASS WRITEWC WORK TYPE WRITE,
WORK CLASS DDLWC WORK TYPE DDL,
WORK CLASS LOADWC WORK TYPE LOAD)
```

Next, you would create a database-level work action set, COUNTACTIONS, to apply to the ACTIVITYTYPES work class set. The work action set would contain a COUNT ACTIVITY work action for each work class in the ACTIVITYTYPES work class set, as in the following example:

```
CREATE WORK ACTION SET COUNTACTIONS FOR DATABASE USING WORK CLASS SET ACTIVITYTYPES
(WORK ACTION COUNTREAD ON WORK CLASS READWC COUNT ACTIVITY,
WORK ACTION COUNTWRITE ON WORK CLASS WRITEWC COUNT ACTIVITY,
WORK ACTION COUNTDDL ON WORK CLASS DDLWC COUNT ACTIVITY,
WORK ACTION COUNTLOAD ON WORK CLASS LOADWC COUNT ACTIVITY)
```

After a sufficient amount of time has passed, you can determine the number of each type of activity that has run by using the WLM_GET_WORK_ACTION_SET_STATS table function:

```
SELECT SUBSTR(CHAR(DBPARTITIONNUM),1,4) AS PART,
LAST_RESET,
SUBSTR(WORK_CLASS_NAME,1,15) AS WORK_CLASS_NAME,
SUBSTR(CHAR(ACT_TOTAL),1,14) AS TOTAL_ACTS
FROM TABLE(WLM_GET_WORK_ACTION_SET_STATS(CAST(NULL AS VARCHAR(128))), -2))
AS WASSTATS WHERE WORK_ACTION_SET_NAME = 'COUNTACTIONS'
ORDER BY WORK_CLASS_NAME, PART
```

Chapter 4. Monitoring and intervention

The third stage of workload management is monitoring, which must be performed on an ongoing basis.

The primary purpose of monitoring is to validate the health and efficiency of your system and the individual workloads running on it. Using table functions, you can access real-time operational data such as a list of running workload occurrences and the activities running in a service class or average response times. Using event monitors you can capture detailed activity information and aggregate activity statistics for historical analysis.

Looking at aggregate information should usually be the first step when you build a monitoring strategy. Aggregates give a good picture of overall data server activity and are also cheaper because you do not have to collect information on every activity in which you might be interested. You can collect more detailed information as you understand the scope of your monitoring needs.

Typical monitoring tasks you can perform are:

- Analyzing the workload on your system to help design your initial DB2 workload manager configuration.
- Tracking and investigating the behavior of your system by obtaining types of operational information that permit you to:
 - Analyze system performance degradation
 - Diagnose activities that are taking too long to complete
 - Investigate agent contention
 - Isolate poorly performing queries

Information is available for activities, service classes, workloads, work classes, threshold queues, and threshold violations.

- Exercising control over the execution environment by canceling queued activities that you expect will cause problems or cancel running activities that you have diagnosed as negatively impacting the system.

Real-time monitoring with table functions

Real-time monitoring data includes information about work currently running on the system, statistics, and metrics for work that has been performed on the system that can help you to determine usage patterns and resource allocation and identify problem areas. You use DB2 table functions to obtain this operational information.

Table functions with names that begin with *WLM_* are DB2 workload manager table functions. These table functions provide access to a set of data relevant to managing your workload, such as workload management statistics, as a virtual DB2 table against which you can issue a *SELECT* statement. This enables you to write applications to query data and analyze it as if it were in a physical table on the data server. The DB2 workload manager table functions are qualified with the *SYSPROC* schema name.

Table functions with names that begin with *MON_* are monitoring metrics functions. Monitoring metrics provide monitoring data about the health of and query performance on your DB2 data server, which can then be used as input to a

3rd party tool or in conjunction with additional scripting you provide to analyze the metrics returned. Only those monitoring metrics functions that are relevant for DB2 workload manager are included here. The monitor metrics table functions are similar to the workload manager statistics table functions. Both return elements describing work that has taken place on the system. The key differences between these monitoring metrics table functions and the DB2 workload manager table functions are:

- The DB2 workload manager table functions provide data that is more statistical in nature, such as computed values like averages, high watermarks, standard deviations, etc. In contrast, the monitoring metrics table functions provide a much more complete set of raw monitoring data.
- The data reported by the DB2 statistics functions is reset when data is sent to a statistics event monitor or when the WLM_COLLECT_STATS procedure is invoked. This resetting of data is necessary to make values such as high watermarks meaningful over a specific collection interval. Data reported by the monitoring metrics functions is also captured by a statistics event monitor, but is never reset. The data reported by monitoring interfaces accumulates from the time a database is activated until the time it is deactivated.

Some table functions return sets of information about the work that is currently running on a system:

Table 43. Table functions that show you the work currently running on the system

Objects for which information is collected	Functions and information returned
Workload occurrences	<p>The WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97 table function returns a list of workload occurrences, across database partitions, that are assigned to a service class. For each occurrence, there is information about the current state and the connection attributes used to assign the workload to the service class and activity statistics indicating activity volume and success rates. For an example of how to use this table function, see “Example: Investigating agent usage by service class” on page 88.</p> <p>The deprecated WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES table function is also available.</p>
Workload occurrence activities	<p>The WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97 table function returns a list of current activities associated with a workload occurrence. For each activity, information is available about the current state of the activity (for example, executing or queued), the type of activity (for example, LOAD, READ, or DDL), and the time at which the activity started. For examples of how to use this table function, see “Example: Aggregating data using DB2 workload manager table functions” on page 159 and “Scenario: Identifying activities that are taking too long to complete” on page 267.</p> <p>The deprecated WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES table function is also available.</p>

Table 43. Table functions that show you the work currently running on the system (continued)

Objects for which information is collected	Functions and information returned
Service class agents	<p>The WLM_GET_SERVICE_CLASS_AGENTS_V97 table function returns a list of database agents associated with a service class or an application handle. Information returned also shows the current state of the agent, the action that the agent is performing, and the status of that action. For an example of how to use this table function, see “Example: Investigating agent usage by service class” on page 88.</p> <p>The deprecated WLM_GET_SERVICE_CLASS_AGENTS table function is also available.</p>
Activities	<p>The MON_GET_ACTIVITY_DETAILS table function returns metrics and other detailed information as an XML document about a specific activity identified by its application handle, unit of work ID, and activity ID. One detail returned is the activity type; depending on that type, a set of additional data is returned. For example, for SQL activities, cost estimates and information about the statement text, package data, and rows returned or modified are provided. Details about the isolation level and processor resource are also available, amongst others.</p> <p>The deprecated WLM_GET_ACTIVITY_DETAILS table function is also available. For an example that shows you how to use this table function, see “Example: Monitoring current system behavior at different levels using DB2 workload manager table functions” on page 155.</p>

Some table functions return monitoring data for all requests executed on the system aggregated by service subclass and workload objects:

Table 44. Table functions that show you monitoring data aggregated by DB2 workload manager objects

Objects for which data is aggregated	Functions and information returned
Workloads	<p>Both the MON_GET_WORKLOAD table function and the MON_GET_WORKLOAD_DETAILS table function return metrics for one or more workloads. The metrics returned by this function represent the accumulation of all metrics of all workload occurrences that use the same workload definition.</p> <p>The MON_GET_WORKLOAD table function returns the most commonly used metrics in a column-based format and is an efficient method of retrieving base metrics.</p> <p>The MON_GET_WORKLOAD_DETAILS table function returns the entire set of available metrics in an XML document format, which provides maximum flexibility for formatting output. The XML based output can be parsed directly by an XML parser, or it can be converted to relational format by the XMLTABLE function.</p>

Table 44. Table functions that show you monitoring data aggregated by DB2 workload manager objects (continued)

Objects for which data is aggregated	Functions and information returned
Service subclasses	<p>Both the MON_GET_SERVICE_SUBCLASS table function and the MON_GET_SERVICE_SUBCLASS_DETAILS table function return metrics for one or more service subclasses. The metrics returned by the table functions represent the accumulation of all metrics for requests that have executed under the indicated service subclass.</p> <p>The MON_GET_SERVICE_SUBCLASS table function returns the most commonly used metrics in a column based format and is an efficient method of retrieving base metrics.</p> <p>The MON_GET_SERVICE_SUBCLASS_DETAILS table function returns the entire set of available metrics in an XML document format, which provides maximum flexibility for formatting output. The XML-based output can be parsed directly by an XML parser, or it can be converted to relational format by the XMLTABLE function.</p>

Statistical information

General statistical information is also available for a number of different objects. You can use this statistical information for a number of different purposes, such as for verifying that changes to your DB2 workload manager configuration have had the expected effect. If you create a new work class to classify READ activities, for example, you can verify that READ activities are being classified under the new work class correctly. You can also use table functions to quickly recognize certain problems with the system. For example, you can use table functions to determine an acceptable value for the average activity lifetime and recognize when this value exceeds its usual range, possibly indicating a problem that requires further investigation.

The following table lists the statistics that you can obtain by using table functions. All statistics table functions return the statistics that accumulated since the last time that you reset the statistics.

Table 45. Table functions that show you statistical information

Objects for which statistics are returned	Functions and statistics returned
Service superclasses	<p>The WLM_GET_SERVICE_SUPERCLASS_STATS table function shows summary statistics across database partitions at the service superclass level: namely, high-water marks for concurrent connections, which are useful when determining peak workload activity.</p>

Table 45. Table functions that show you statistical information (continued)

Objects for which statistics are returned	Functions and statistics returned
Service subclasses	<p>The WLM_GET_SERVICE_SUBCLASS_STATS_V97 table function shows summary statistics across database partitions at the service subclass level (all activities run in service subclasses). Statistics include numbers of completed activities and average execution times. This information is useful when you are looking at general system health and distribution of activities across service classes and database partitions. For examples of how to use this table function, see “Example: Obtaining point-in-time statistics from service classes” on page 158, “Example: Aggregating data using DB2 workload manager table functions” on page 159, “Example: Analyzing a service class–related system slowdown” on page 86, and “Scenario: Investigating a workload-related system slowdown” on page 266.</p> <p>The deprecated WLM_GET_SERVICE_SUBCLASS_STATS table function is also available.</p>
Workloads	<p>The WLM_GET_WORKLOAD_STATS_V97 table function shows summary statistics across database partitions at the workload level. These include high-water marks for concurrent workload occurrences and numbers of completed activities. This information is useful when you are monitoring general system health or drilling down to identify problem areas. For an example of how to use this table function, see “Scenario: Investigating a workload-related system slowdown” on page 266.</p> <p>The deprecated WLM_GET_WORKLOAD_STATS table function is also available.</p>
Work action sets	<p>The WLM_GET_WORK_ACTION_SET_STATS table function shows summary statistics across database partitions at the work action set level: namely, the number of activities in each work class that had the corresponding work actions applied to them. This information is useful for understanding the effectiveness of a work action set and understanding the types of activities running on the system. For an example of how to use this table function, see “Example: Analyzing workloads by activity type” on page 57.</p>
Threshold queues	<p>The WLM_GET_QUEUE_STATS table function shows summary statistics across database partitions for the queues used for thresholds. Statistics include the current and total numbers of queued activities and total time spent in a queue. This information is useful when you are querying current queued activity or validating that you defined a threshold correctly. Excessive queuing might indicate that a threshold is too restrictive, and very little queuing might indicate that a threshold is not restrictive enough or not needed.</p>

Statistics are useful only if the time period during which they are collected is meaningful. Collecting statistics over a very long time, and for any length of time using the WLM_COLLECT_STATS stored procedure, might be less useful if it becomes difficult to identify changes to trends or problem areas because there is too much old data. Thus, you can reset statistics at any time.

Because of the default workload and default user service classes, monitoring capabilities exist from the moment that you install the DB2 data server. These can

help you to start identifying sources of activities that you can use to create workloads and the service classes to which you can assign them.

Example: Using DB2 workload manager table functions

A large amount of data is available through DB2 workload manager real-time monitoring. The example in this topic shows how you might start using the information.

In this situation, only the default workload and service class are in place. Use this example to understand how you can use the table functions to understand what, exactly, is running on the data server. Follow these steps:

1. Use the Service Superclass Statistics table function to show all of the service superclasses. After you install or upgrade to DB2 9.5 or later, three default superclasses are defined: one for maintenance activities, one for system activities, and one for user activities. SYSDEFAULTUSERCLASS is the service class of interest.

```
SELECT VARCHAR(SERVICE_SUPERCLASS_NAME,30) AS SUPERCLASS
       FROM TABLE(WLM_GET_SERVICE_SUBCLASS_STATS_V97('','-1)) AS T
```

```
SUPERCLASS
-----
SYSDEFAULTSYSTEMCLASS
SYSDEFAULTMAINTENANCECLASS
SYSDEFAULTUSERCLASS
```

3 record(s) selected.

2. Use the Service Subclass Statistics table function to show statistics for all the service subclasses of the SYSDEFAULTUSERCLASS superclass. For each service subclass you can see the current volume of requests that are being processed, the number of activities that have completed execution, and the overall distribution of activities across database partitions (possibly indicating a problem if the distribution is uneven). You can optionally obtain additional statistics including the average lifetime for activities, the average amount of time activities spend queued, and so on. You can obtain optional statistics for a service subclass by specifying the COLLECT AGGREGATE ACTIVITY DATA keyword on the ALTER SERVICE CLASS statement to enable aggregate activity statistics collection.

```
SELECT VARCHAR(SERVICE_SUPERCLASS_NAME, 20) AS SUPERCLASS,
       VARCHAR(SERVICE_SUBCLASS_NAME, 20) AS SUBCLASS,
       COORD_ACT_COMPLETED_TOTAL,
       COORD_ACT_ABORTED_TOTAL,
       COORD_ACT_REJECTED_TOTAL,
       CONCURRENT_ACT_TOP
FROM TABLE(WLM_GET_SERVICE_SUBCLASS_STATS_V97(
 'SYSDEFAULTUSERCLASS', 'SYSDEFAULTSUBCLASS', -1))
AS T
```

SUPERCLASS	SUBCLASS	COORD_ACT_COMPLETED_TOTAL	COORD_ACT_ABORTED_TOTAL	COORD_ACT_REJECTED_TOTAL	CONCURRENT_ACT_TOP
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	2	0	0	1

1 record(s) selected.

3. For a given service subclass, use the Workload Occurrence Information table function to list the occurrences of a workload that are mapped to the service subclass. The table function displays all of the connection attributes, which you can use to identify the source of the activities. This information can be quite useful in determining custom workload definitions in the future. For example, perhaps a specific workload occurrence listed here has a large volume of work from an application as shown by the activities completed counter.

```

SELECT APPLICATION_HANDLE,
       VARCHAR(WORKLOAD_NAME, 30) AS WORKLOAD,
       VARCHAR(SESSION_AUTH_ID, 20) AS SESSION_AUTH_ID,
       VARCHAR(APPLICATION_NAME, 20) AS APPL_NAME
FROM TABLE(WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97(
 'SYSDEFAULTUSERCLASS', 'SYSDEFAULTSUBCLASS', -1))
AS T

```

APPLICATION_HANDLE	WORKLOAD	SESSION_AUTH_ID	APPL_NAME
431	SYSDEFAULTUSERWORKLOAD	SWALKTY	db2bp

1 record(s) selected.

- a. For that application, use the Workload Occurrence Activities Information table function to show the current activities across database partitions that were created from the application's connection. You can use this information for a number of purposes, including identifying activities that might be causing problems on the data server.

```

SELECT APPLICATION_HANDLE,
       LOCAL_START_TIME,
       UOW_ID,
       ACTIVITY_ID,
       ACTIVITY_TYPE
FROM TABLE(WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97(431,-1)) AS T

```

APPLICATION_HANDLE	LOCAL_START_TIME	UOW_ID	ACTIVITY_ID	ACTIVITY_TYPE
431	2008-06-17-12.49.46.854259	11	1	READ_DML

1 record(s) selected

- b. For each activity, retrieve more detailed information by using the Activity Details table function. The data might show that some SQL statements are returning huge numbers of rows, that some activities have been idle for a long time, or that some queries are running that have an extremely large estimated cost. In situations such as these, it might make sense to define some thresholds to identify and prevent potentially damaging behavior in the future.

```

SELECT VARCHAR(NAME, 20) AS NAME,
       VARCHAR(VALUE, 40) AS VALUE
FROM TABLE(WLM_GET_ACTIVITY_DETAILS(431,11,1,-1))
AS T WHERE NAME IN ('UOW_ID', 'ACTIVITY_ID', 'STMT_TEXT')

```

NAME	VALUE
UOW_ID	1
ACTIVITY_ID	1
STMT_TEXT	select * from syscat.tables

3 record(s) selected.

Example: Monitoring current system behavior at different levels using DB2 workload manager table functions

DB2 workload manager provides a number of table functions that you can use to obtain data about your workload management configuration.

Installing DB2 Version 9.5 or later creates a set of default workloads and service classes. Before deciding how to implement your own DB2 workload manager solution, you can use the table functions to observe work being performed in the system in terms of the default workload occurrences, service classes, and activities.

You can start by obtaining the list of workload occurrences in a service class. To do this, use the WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97

table function. In the following example, an empty string is passed for *service_superclass_name* and *service_subclass_name*, and -2 (a wildcard character) is passed for *dbpartitionnum*:

```
SELECT SUBSTR(SERVICE_SUPERCLASS_NAME,1,19) AS SUPERCLASS_NAME,
       SUBSTR(SERVICE_SUBCLASS_NAME,1,18) AS SUBCLASS_NAME,
       SUBSTR(CHAR(DBPARTITIONNUM),1,4) AS PART,
       SUBSTR(CHAR(COORD_PARTITION_NUM),1,4) AS COORDPART,
       SUBSTR(CHAR(APPLICATION_HANDLE),1,7) AS APPHNDL,
       SUBSTR(CHAR(WORKLOAD_NAME),1,22) AS WORKLOAD_NAME,
       SUBSTR(CHAR(WORKLOAD_OCCURRENCE_ID),1,6) AS WLO_ID
FROM TABLE(WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97('', '', -2)) AS SCINFO
ORDER BY SUPERCLASS_NAME, SUBCLASS_NAME, PART, APPHNDL, WORKLOAD_NAME, WLO_ID
```

Assume that the system has four database partitions and that there are two applications performing activities on the database when you issue the query. The results would resemble the following ones:

SUPERCLASS_NAME	SUBCLASS_NAME	PART	COORDPART	APPHNDL	WORKLOAD_NAME	WLO_ID
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	0	0	1	SYSDEFAULTUSERWORKLOAD	1
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	0	0	2	SYSDEFAULTUSERWORKLOAD	2

The results indicate that both workload occurrences were assigned to the SYSDEFAULTUSERWORKLOAD workload. The results also show that both workload occurrences were assigned to the SYSDEFAULTSUBCLASS service subclass in the SYSDEFAULTUSERCLASS service superclass and that both workload occurrences are from the same coordinator partition (partition 0).

Next, you can also use the WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97 table function again to determine the connection attributes of the two workload occurrences:

```
SELECT SUBSTR(CHAR(APPLICATION_HANDLE),1,7) AS APPHNDL,
       SUBSTR(CHAR(WORKLOAD_NAME),1,22) AS WORKLOAD_NAME,
       SUBSTR(CHAR(WORKLOAD_OCCURRENCE_ID),1,6) AS WLO_ID,
       SUBSTR(CHAR(SYSTEM_AUTH_ID),1,9) AS SYSAUTHID,
       SUBSTR(CHAR(APPLICATION_NAME),1,15) AS APPLNAME
FROM TABLE(WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97('', '', 0)) AS SCINFO
ORDER BY APPHNDL, WORKLOAD_NAME, WLO_ID
```

APPHNDL	WORKLOAD_NAME	WLO_ID	SYSAUTHID	APPLNAME
1	SYSDEFAULTUSERWORKLOAD	1	LYNN	accountspay
2	SYSDEFAULTUSERWORKLOAD	2	KATE	businessobjects

Then, you can use the WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97 table function to show the current activities of one of the workload occurrences:

```
SELECT SUBSTR(CHAR(COORD_PARTITION_NUM),1,5) AS COORD,
       SUBSTR(CHAR(DBPARTITIONNUM),1,4) AS PART,
       SUBSTR(CHAR(UOW_ID),1,5) AS UOWID,
       SUBSTR(CHAR(ACTIVITY_ID),1,5) AS ACTID,
       SUBSTR(CHAR(PARENT_UOW_ID),1,8) AS PARUOWID,
       SUBSTR(CHAR(PARENT_ACTIVITY_ID),1,8) AS PARACTID,
       SUBSTR(CHAR(ACTIVITY_TYPE),1,9) AS ACTTYPE,
       SUBSTR(CHAR(NESTING_LEVEL),1,7) AS NESTING
FROM TABLE(WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97(1, -2)) AS WLOACTS
ORDER BY PART, UOWID, ACTID
```

COORD	PART	UOWID	ACTID	PARUOWID	PARACTID	ACTTYPE	NESTING
0	0	1	3	-	-	CALL	0
0	0	1	5	1	3	READ_DML	1

```

0    1    1    5    -    -    READ_DML 1
0    2    1    5    -    -    READ_DML 1
0    3    1    5    -    -    READ_DML 1

```

The query results show that workload occurrence 1 is running two activities. One activity is a stored procedure (indicated by the activity type of CALL), and the other activity is a DML activity that performs a read (for example, a SELECT statement). The DML activity is nested in the stored procedure call. You can tell that the DML activity is nested because the parent unit of work identifier and parent activity identifier of the DML activity match the unit of work identifier and the activity identifier of the CALL activity. You can also tell that the DML activity is executing on database partitions 0, 1, 2, and 3. The parent identifier information is available only on the coordinator partition.

You can obtain more information about an individual activity that is currently running by using the MON_GET_ACTIVITY_DETAILS table function. This table function returns an XML document where the elements in the document describe the activity. In this example, the XMLTABLE function is used to return a result table from the XML output.

```

SELECT D.APP_HANDLE,
       D.MEMBER,
       D.COORD_MEMBER,
       D.LOCAL_START_TIME,
       D.UOW_ID,
       D.ACTIVITY_ID,
       D.PARENT_UOW_ID,
       D.PARENT_ACTIVITY_ID,
       D.ACTIVITY_TYPE,
       D.NESTING_LEVEL,
       D.INVOCATION_ID,
       D.ROUTINE_ID
FROM TABLE(MON_GET_ACTIVITY_DETAILS(65592, 1, 1, -2)) AS ACTDETAILS,
XMLTABLE (XMLNAMESPACES( DEFAULT 'http://www.ibm.com/xmlns/prod/db2/mon'),
          '$details/db2_activity_details' PASSING XMLPARSE(DOCUMENT
          ACTDETAILS.DETAILS) as "details"
COLUMNS "APP_HANDLE"          BIGINT    PATH 'application_handle',
         "MEMBER"              BIGINT    PATH 'member',
         "COORD_MEMBER"       BIGINT    PATH 'coord_member',
         "LOCAL_START_TIME"    VARCHAR(26) PATH 'local_start_time',
         "UOW_ID"              BIGINT    PATH 'uow_id',
         "ACTIVITY_ID"         BIGINT    PATH 'activity_id',
         "PARENT_UOW_ID"       BIGINT    PATH 'parent_uow_id',
         "PARENT_ACTIVITY_ID" BIGINT    PATH 'parent_activity_id',
         "ACTIVITY_TYPE"       VARCHAR(10) PATH 'activity_type',
         "NESTING_LEVEL"       BIGINT    PATH 'nesting_level',
         "INVOCATION_ID"       BIGINT    PATH 'invocation_id',
         "ROUTINE_ID"          BIGINT    PATH 'routine_id'
) AS D;

```

```

APP_HANDLE      MEMBER      COORD_MEMBER      LOCAL_START_TIME
UOW_ID          ACTIVITY_ID  PARENT_UOW_ID
PARENT_ACTIVITY_ID  ACTIVITY_TYPE NESTING_LEVEL      INVOCATION_ID
ROUTINE_ID
-----
-----
          65592          1          1
2009-04-07-18.39.42.549197
1          -          - READ_DML          0          0          0
          65592          0
2009-04-07-18.39.42.552763
1          -          - READ_DML          0          0          0

```

2 record(s) selected.

The table functions mentioned previously provide a high-level description of work that is running in the system. The information that these table functions provide regarding the status of the work is limited to an activity state such as

EXECUTING. If you want to probe further to discover what exactly is occurring in a service class at a point in time, you can run the WLM_GET_SERVICE_CLASS_AGENTS_V97 table function.

In the following example, WLM_GET_SERVICE_CLASS_AGENTS_V97 is called by passing 1 for *application_handle* and -2 (a wildcard character) for *dbpartitionnum*:

```
SELECT SUBSTR(CHAR(APPLICATION_HANDLE),1,7) AS APPHANDLE,
       SUBSTR(CHAR(DBPARTITIONNUM),1,4) AS PART,
       SUBSTR(CHAR(AGENT_TID),1,9) AS AGENT_TID,
       SUBSTR(AGENT_TYPE,1,11) AS AGENTTYPE,
       SUBSTR(AGENT_STATE,1,10) AS AGENTSTATE,
       SUBSTR(REQUEST_TYPE,1,14) AS REQTYPE,
       SUBSTR(CHAR(UOW_ID),1,6) AS UOW_ID,
       SUBSTR(CHAR(ACTIVITY_ID),1,6) AS ACT_ID
FROM TABLE(WLM_GET_SERVICE_CLASS_AGENTS_V97(' ', ' ', 1, -2)) AS SCDETAILS
ORDER BY APPHANDLE, PART, AGENT_TID
```

APPHANDLE	PART	AGENT_TID	AGENTTYPE	AGENTSTATE	REQTYPE	UOW_ID	ACT_ID
1	0	3	COORDINATOR	ACTIVE	FETCH	1	5
1	0	4	PDBSUBAGENT	ACTIVE	SUBSECTION:1	1	5
1	1	2	PDBSUBAGENT	ACTIVE	SUBSECTION:2	1	5

The results show a coordinator agent and a subagent on database partition 0 and a subagent on database partition 1 operating on behalf of an activity with a unit of work identifier of 1 and an activity identifier of 5. The coordinator agent information indicates that the request is a fetch request.

Example: Obtaining point-in-time statistics from service classes

Every activity is mapped to a service class before being executed. You can monitor the system by using the service class statistics table functions and querying all of the service classes on all of the database partitions to obtain point-in-time statistics.

You can use the following statement to obtain service class statistics, such as the average activity lifetime. Passing an empty string for an argument for the WLM_GET_SERVICE_SUBCLASS_STATS_V97 table function means that the result is not to be restricted by that argument. The value of the last argument, *dbpartitionnum*, is -2 (a wildcard character), which means that data from all database partitions is to be returned.

Note: Lifetime information is only returned for those service classes that are defined with COLLECT AGGREGATE ACTIVITY DATA.

```
SELECT SUBSTR(SERVICE_SUPERCLASS_NAME,1,19) AS SUPERCLASS_NAME,
       SUBSTR(SERVICE_SUBCLASS_NAME,1,18) AS SUBCLASS_NAME,
       SUBSTR(CHAR(DBPARTITIONNUM),1,4) AS PART,
       CAST(COORD_ACT_LIFETIME_AVG / 1000 AS DECIMAL(9,3)) AS AVGLIFETIME,
       CAST(COORD_ACT_LIFETIME_STDDEV / 1000 AS DECIMAL(9,3)) AS STDDEVLIFETIME,
       SUBSTR(CAST(LAST_RESET AS VARCHAR(30)),1,16) AS LAST_RESET
FROM TABLE(WLM_GET_SERVICE_SUBCLASS_STATS_V97(' ', ' ', -2)) AS SCSTATS
ORDER BY SUPERCLASS_NAME, SUBCLASS_NAME, PART
```

SUPERCLASS_NAME	SUBCLASS_NAME	PART	AVGLIFETIME	STDDEVLIFETIME	LAST_RESET
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	0	691.242	34.322	2006-07-24-11.44
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	1	644.740	22.124	2006-07-24-11.44
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	2	612.431	43.347	2006-07-24-11.44
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	3	593.451	28.329	2006-07-24-11.44

You can also use the WLM_GET_SERVICE_SUBCLASS_STATS_V97 table function to obtain the high watermark for the concurrency of coordinator activities that run in the service class on each database partition:

```
SELECT SUBSTR(SERVICE_SUPERCLASS_NAME,1,19) AS SUPERCLASS_NAME,
       SUBSTR(SERVICE_SUBCLASS_NAME,1,18) AS SUBCLASS_NAME,
       SUBSTR(CHAR(DBPARTITIONNUM),1,4) AS PART,
       CONCURRENT_ACT_TOP AS ACTHIGHWATERMARK
FROM TABLE(WLM_GET_SERVICE_SUBCLASS_STATS_V97(' ', ' ', -2)) AS SCSTATS
ORDER BY SUPERCLASS_NAME, SUBCLASS_NAME, PART
```

SUPERCLASS_NAME	SUBCLASS_NAME	PART	ACTHIGHWATERMARK
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	0	10
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	1	0
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	2	0
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	3	0

By reviewing the average lifetime and number of completed activities, you can use the output of the WLM_GET_SERVICE_SUBCLASS_STATS_V97 table function to obtain a rolled-up view of the workload on each database partition in the database. Significant variations in the high watermarks and averages returned by a table function might indicate a change in the workload on the system.

Example: Aggregating data using DB2 workload manager table functions

You can perform various aggregations on table data in a DB2 workload manager configuration to monitor the system and identify potential problems.

The following are examples of data aggregation that you can perform to identify problems.

Identifying increases in average query lifetimes because queries are spending too much time in the queue

You can identify a situation in which the average query lifetime increases because queries are spending too much time in the queue by showing the average time in the queue for coordinator activities for each service class, across the whole system.

Following is an example that shows the percentage of time that the average query spends queued for coordinator activities for each service class, summed across all database partitions:

```
SELECT SUBSTR(SERVICE_SUPERCLASS_NAME,1,19) AS SUPERCLASS_NAME,
       SUBSTR(SERVICE_SUBCLASS_NAME,1,18) AS SUBCLASS_NAME,
       CASE WHEN (SUM(COORD_ACT_COMPLETED_TOTAL) = 0) THEN
         0
       ELSE
         SUM(COORD_ACT_QUEUE_TIME_AVG * COORD_ACT_COMPLETED_TOTAL) /
         SUM(COORD_ACT_COMPLETED_TOTAL)
       END AS AVG_QUEUE_TIME,
       CASE WHEN (SUM(COORD_ACT_COMPLETED_TOTAL) = 0) THEN
         0
       ELSE
         SUM(COORD_ACT_LIFETIME_AVG * COORD_ACT_COMPLETED_TOTAL) /
         SUM(COORD_ACT_COMPLETED_TOTAL)
       END AS AVG_LIFE_TIME,
       CASE WHEN (SUM(COORD_ACT_COMPLETED_TOTAL) = 0) THEN
         0
       ELSE CASE WHEN
         (CAST(SUM(COORD_ACT_LIFETIME_AVG * COORD_ACT_COMPLETED_TOTAL) /
          SUM(COORD_ACT_COMPLETED_TOTAL) AS INTEGER) = 0) THEN
```



```

0
ELSE
100 * (SUM(COORD_ACT_QUEUE_TIME_AVG * COORD_ACT_COMPLETED_TOTAL) /
SUM(COORD_ACT_COMPLETED_TOTAL)) /
(SUM(COORD_ACT_LIFETIME_AVG * COORD_ACT_COMPLETED_TOTAL) /
SUM(COORD_ACT_COMPLETED_TOTAL))
END
END AS PERCENT_TIME_QUEUED
FROM TABLE(WLM_GET_SERVICE_SUBCLASS_STATS_V97(' ', ' ', -2)) AS STATS
GROUP BY SERVICE_SUPERCLASS_NAME, SERVICE_SUBCLASS_NAME
ORDER BY SERVICE_SUPERCLASS_NAME, SERVICE_SUBCLASS_NAME

```

SUPERCLASS_NAME	SUBCLASS_NAME	AVG_QUEUE_TIME	AVG_LIFE_TIME	PERCENT_TIME_QUEUED
SYSDEFAULTMAINTENAN	SYSDEFAULTSUBCLASS	+0.000000000000000E+000	+0.000000000000000E+000	+0.000000000000000E+000
SYSDEFAULTSYSTEMCLA	SYSDEFAULTSUBCLASS	+0.000000000000000E+000	+0.000000000000000E+000	+0.000000000000000E+000
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	+2.328601000000000E+005	+8.234214240000000E+005	+2.828000000000000E-001

The results show that the percentage of time that the average activity spends in the queue is about 28%. If previous experience with the system and workload indicates that this is too high or too low, making adjustments to your thresholds can have an impact on the percentage of time spent queuing.

Identifying sudden increases in the number of queries running in a workload

Assume that you have a workload called WL1. You can identify a situation in which a large number of queries are running in the workload by showing the total number of executing non-nested coordinator activities for the workload across the whole system:

```

SELECT SUBSTR(WORKLOAD_NAME,1,22) AS WLNAME,
COUNT(*) AS TOTAL_EXE_ACT
FROM TABLE(WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97(' ', ' ', -2)) AS APPS,
TABLE(WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97(APPS.APPLICATION_HANDLE, -2)) AS APPACTS
WHERE WORKLOAD_NAME = 'WL1' AND
APPS.DBPARTITIONNUM = APPS.COORD_PARTITION_NUM AND
ACTIVITY_STATE = 'EXECUTING' AND
NESTING_LEVEL = 0
GROUP BY WORKLOAD_NAME

```

WLNAME	TOTAL_EXE_ACT
WL1	5

Historical monitoring with WLM event monitors

DB2 workload manager uses event monitors to capture information that might be of use in the future or for historical analysis.

Three event monitors are available for you to use. Each event monitor serves a different purpose:

Activity event monitor

This monitor captures information about individual activities in a service class, workload, or work class or activities that violated a threshold. The amount of data that is captured for each activity is configurable and should be considered when you determine the amount of disk space and the length of time required to keep the monitor data. A common use for activity data is to use it as input to tools such as db2advise or to use access plans (from the explain utility) to help determine table, column, and index usage for a set of queries.

You can collect information about an activity by specifying COLLECT ACTIVITY DATA for the service class, workload, or work action to which such an activity belongs or a threshold that might be violated by such an

activity. The information is collected when the activity completes, regardless of whether the activity completes successfully.

Note that if an activities event monitor is active when the database deactivates, any backlogged activity records in the queue are discarded. To ensure that you obtain all activities event monitor records and that none are discarded, explicitly deactivate the activities event monitor first before deactivating the database. When an activities event monitor is explicitly deactivated, all backlogged activity records in the queue are processed before the event monitor deactivates.

Threshold violations event monitor

This monitor captures information when a threshold is violated. It indicates what threshold was violated, the activity that caused the violation, and what action was taken when it occurred.

If you specify `COLLECT ACTIVITY DATA` for the threshold and an activities event monitor is created and active, information is also collected about activities that violate the threshold, but this information is collected when the activity ends (either successfully or unsuccessfully).

You can obtain details about a threshold by querying the `SYSCAT.THRESHOLDS` view.

Statistics event monitor

This monitor serves as a low-overhead alternative to capturing detailed activity information by collecting aggregate data (for example, the number of activities completed and average execution time). Aggregate data includes histograms for a number of activity measurements including lifetime, queue time, execution time and estimated cost. You can use histograms to understand the distribution of values, identify outliers, and compute additional statistics such as averages and standard deviations. For example, histograms can help you understand the variation in lifetime that users experience. The average life time alone does not reflect what a user experiences if there is a high degree of variability. See “Collecting workload management statistics using a statistics event monitor” on page 188 for a description of how to send statistics to the event monitor.

The following figure shows the different monitoring options available to access workload information: table functions to access real-time statistics, and activity details and historical information captured as efficient aggregates or as details about individual activities:

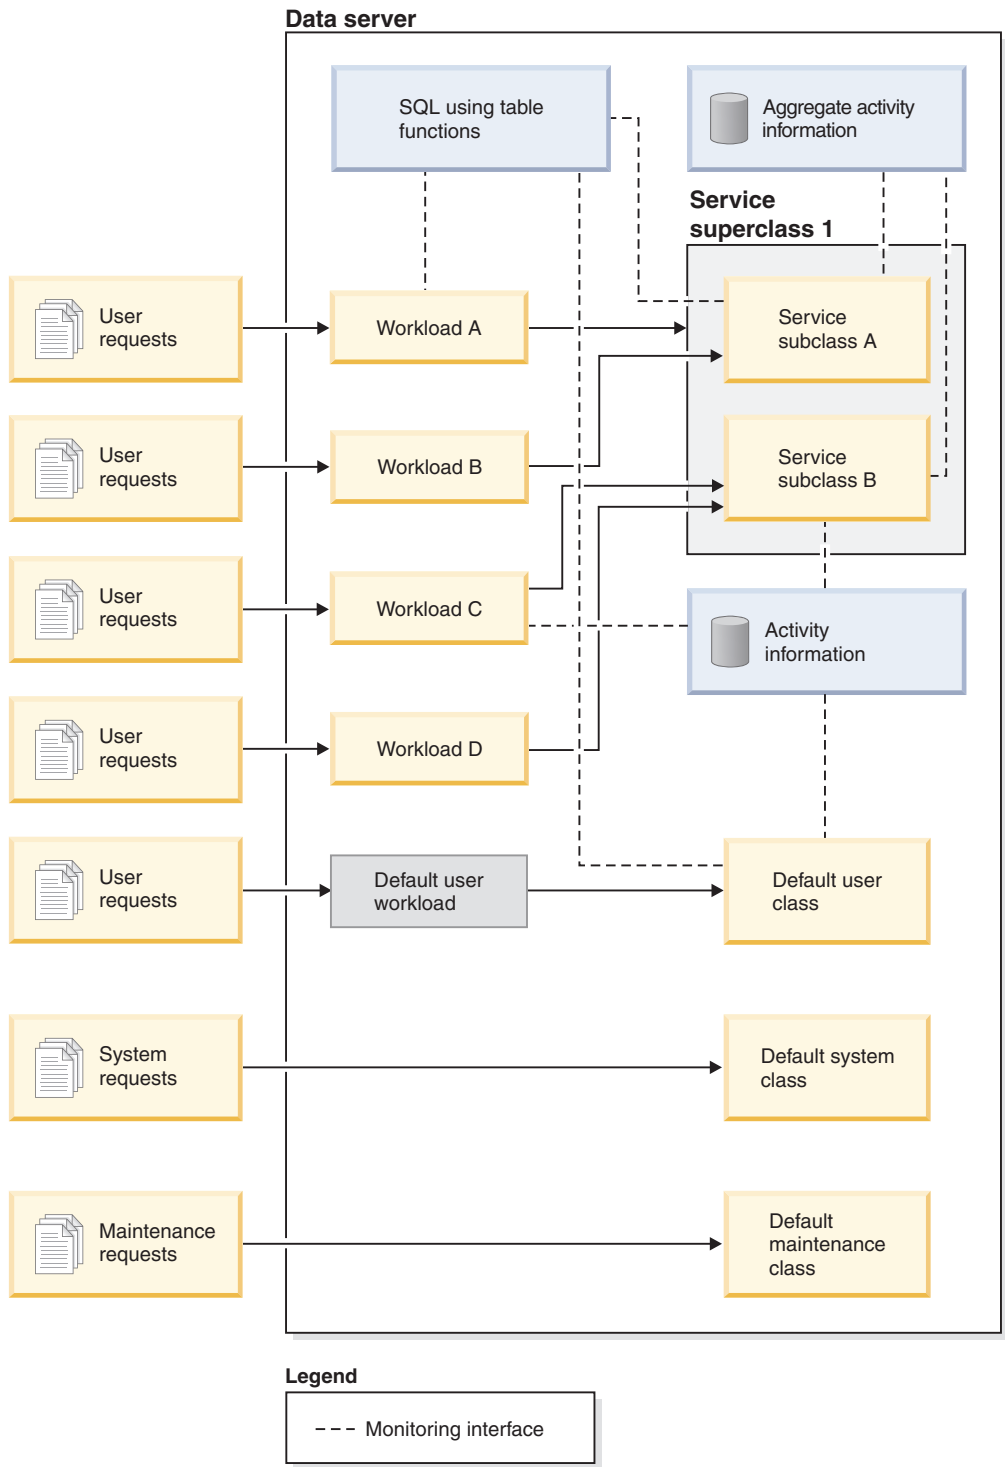


Figure 22. Workload management with monitoring

Unlike statement, connection, and transaction event monitors, the activity, statistics, and threshold violations event monitors do not have event conditions (that is, conditions specified on the WHERE keyword of the CREATE EVENT MONITOR statement). Instead, these event monitors rely on the attributes of service classes, workloads, work classes, and thresholds to determine whether these objects send their activity information or aggregate information to these monitors.

Typically, event monitors write data to either tables or files. You need to prune these tables or files periodically because they are not automatically pruned.

You can use the `wlmevmon.ddl` script in the `sqllib/misc` directory to create and enable three event monitors called `DB2ACTIVITIES`, `DB2STATISTICS`, and `DB2THRESHOLDVIOLATIONS`. If necessary, modify the script to change the table space or other parameters.

Example

Example: Identify queries with a large estimated cost using the statistics event monitor: You suspect that your database workload occasionally includes large, expensive queries, possibly due to the poor optimization of the queries themselves. You want to identify these queries so that you can prevent them from consuming excessive resources on your system, with a long-term goal of perhaps rewriting some of the queries to improve performance. The statistics event monitor provides you with a low-overhead way to measure the estimated cost of your queries which you can then use to determine what the maximum acceptable estimated cost for a query on your data server should be. A query that is poorly optimized is typically distinguished by a large estimated cost that is many times larger than the estimated cost of most other queries.

To get started, you need to create and activate a statistics event monitor and to start collecting extended aggregate activity data for the service class where the queries run:

```
CREATE EVENT MONITOR DB2STATISTICS
  FOR STATISTICS WRITE TO TABLE

SET EVENT MONITOR DB2STATISTICS STATE 1
```

In this example, all queries run in the `SYSDEFAULTSUBCLASS` subclass of the `SYSDEFAULTUSERCLASS` service class, which you can alter to collect the required data:

```
ALTER SERVICE CLASS SYSDEFAULTSUBCLASS UNDER SYSDEFAULTUSERCLASS
  COLLECT AGGREGATE ACTIVITY DATA EXTENDED
```

A full day of work might provide a reasonable approximation of the range of queries your data server typically processes. At the end of the day, you can copy the statistics collected from memory to the statistics event monitor by running the `WLM_COLLECT_STATS` stored procedure:

```
CALL WLM_COLLECT_STATS()
```

Included with the different statistics written to the event monitor tables are the estimated cost statistics of queries. To see them, you can query the service class statistics table `SCSTATS_DB2STATISTICS`:

```
SELECT STATISTICS_TIMESTAMP,
       COORD_ACT_EST_COST_AVG,
       COST_ESTIMATE_TOP
FROM SCSTATS_DB2STATISTICS
WHERE SERVICE_SUPERCLASS_NAME = 'SYSDEFAULTUSERCLASS'
  AND SERVICE_SUBCLASS_NAME = 'SYSDEFAULTSUBCLASS'
```

STATISTICS_TIMESTAMP	COORD_ACT_EST_COST_AVG	COST_ESTIMATE_TOP
2008-09-03-09.49.04.455979	169440	13246445

1 record(s) selected.

The output shows that the average query has an estimated cost in the range of hundreds of thousands of timerons, and that the largest queries have estimated costs larger than ten million timerons. You can confirm that queries of ten million or more timerons are outliers by looking at the estimated cost histogram, which was generated at the same time that the average and high watermarks shown in the output were written to the event monitor table. You can look at the histogram by querying the HISTOGRAMBIN_DB2STATISTICS table as follows:

```

SELECT STATISTICS_TIMESTAMP,
       TOP,
       NUMBER_IN_BIN
FROM HISTOGRAMBIN_DB2STATISTICS HIST,
     SYSCAT.SERVICECLASSES SC
WHERE HIST.SERVICE_CLASS_ID = SC.SERVICECLASSID
      AND SC.PARENTSERVICECLASSNAME = 'SYSDEFAULTUSERCLASS'
      AND SC.SERVICECLASSNAME = 'SYSDEFAULTSUBCLASS'
      AND HISTOGRAM_TYPE = 'COORDACTESTCOST'"

```

STATISTICS_TIMESTAMP	TOP	NUMBER_IN_BIN
2008-09-03-09.49.04.455979	1	0
2008-09-03-09.49.04.455979	2	0
2008-09-03-09.49.04.455979	3	0
2008-09-03-09.49.04.455979	5	0
2008-09-03-09.49.04.455979	8	0
2008-09-03-09.49.04.455979	12	1
2008-09-03-09.49.04.455979	19	0
2008-09-03-09.49.04.455979	29	0
2008-09-03-09.49.04.455979	44	2
2008-09-03-09.49.04.455979	68	5
2008-09-03-09.49.04.455979	103	22
2008-09-03-09.49.04.455979	158	14
2008-09-03-09.49.04.455979	241	54
2008-09-03-09.49.04.455979	369	2
2008-09-03-09.49.04.455979	562	142
2008-09-03-09.49.04.455979	858	21
2008-09-03-09.49.04.455979	1309	123
2008-09-03-09.49.04.455979	1997	512
2008-09-03-09.49.04.455979	3046	643
2008-09-03-09.49.04.455979	4647	201
2008-09-03-09.49.04.455979	7089	875
2008-09-03-09.49.04.455979	10813	1445
2008-09-03-09.49.04.455979	16493	5386
2008-09-03-09.49.04.455979	25157	2409
2008-09-03-09.49.04.455979	38373	8940
2008-09-03-09.49.04.455979	58532	9820
2008-09-03-09.49.04.455979	89280	2149
2008-09-03-09.49.04.455979	136181	798
2008-09-03-09.49.04.455979	207720	2411
2008-09-03-09.49.04.455979	316840	14989
2008-09-03-09.49.04.455979	483283	9831
2008-09-03-09.49.04.455979	737162	1451
2008-09-03-09.49.04.455979	1124409	213
2008-09-03-09.49.04.455979	1715085	24
2008-09-03-09.49.04.455979	2616055	1
2008-09-03-09.49.04.455979	3990325	0
2008-09-03-09.49.04.455979	6086529	0
2008-09-03-09.49.04.455979	9283913	0
2008-09-03-09.49.04.455979	14160950	3
2008-09-03-09.49.04.455979	21600000	0
2008-09-03-09.49.04.455979	-1	0

In the histogram, the value in the number_in_bin column for queries whose top is greater than 2616055 is zero until top reaches 14160950, where the number_in_bin becomes 3. These three queries are outliers and can be controlled with an ESTIMATEDSQLCOST threshold set to trigger if the estimated cost of a query

exceeds 10 million timerons which you can use to prevent such activities from executing and to monitor them more closely.

Example: Using the threshold violations event monitor: To control activities of a certain estimated cost, you want to define an ESTIMATEDSQLCOST threshold on your workload that applies only to that subset of your total workload exceeding a certain estimated cost. Having looked at the estimated cost histogram, you determined that activities with an estimated cost in the range of 0 to under 3 million timerons occur frequently and that activities with an estimated cost over 10 million timerons occur rarely (perhaps only a few times a day and perhaps always due to some flaw in the query, such as the use of a Cartesian join).

To verify that a threshold of 10 million timerons is effective in stopping those few activities a day that should not be allowed to run, you can create and activate a threshold event monitor:

```
CREATE THRESHOLD TH1
  FOR DATABASE ACTIVITIES
  ENFORCEMENT DATABASE
  WHEN ESTIMATEDSQLCOST > 10000000
  STOP EXECUTION

CREATE EVENT MONITOR DB2THRESHOLDVIOLATIONS
  FOR THRESHOLD VIOLATIONS
  WRITE TO TABLE

SET EVENT MONITOR DB2THRESHOLDVIOLATIONS STATE 1
```

After the end of the day, you can see what threshold violations occurred by querying the threshold violations table:

```
SELECT THRESHOLDID,
       SUBSTR(THRESHOLD_PREDICATE, 1, 20) PREDICATE,
       TIME_OF_VIOLATION,
       THRESHOLD_MAXVALUE,
       THRESHOLD_ACTION
FROM THRESHOLDVIOLATIONS_DB2THRESHOLDVIOLATIONS
ORDER BY TIME_OF_VIOLATION, THRESHOLDID
```

THRESHOLDID	PREDICATE	TIME_OF_VIOLATION	THRESHOLD_MAXVALUE	THRESHOLD_ACTION
1	EstimatedSQLCost	2008-09-02-22.39.10.000000	10000000	Stop

1 record(s) selected.

Example: Using the activity event monitor

The previous example showed how you can collect threshold information in an event monitor table to confirm that activities with a large estimated cost are being prevented from executing by a threshold. After seeing these threshold violations, you want to determine what the SQL statement texts producing these large queries are, so that you can use the explain facility to determine if an index is needed on the tables being queried.

Collecting this additional information requires creating and activating an activity event monitor and altering the threshold to turn on activity collection with details:

```
CREATE EVENT MONITOR DB2ACTIVITIES
  FOR ACTIVITIES WRITE TO TABLE

SET EVENT MONITOR DB2ACTIVITIES STATE 1
```

```
ALTER THRESHOLD TH1
  WHEN EXCEEDED
  COLLECT ACTIVITY DATA WITH DETAILS
```

When you query the threshold violations table again after another business day has passed, you can perform a join with the `ACTIVITYSTMT_DB2ACTIVITIES` table to see the SQL statement text of any activity that violated the threshold:

```
SELECT THRESHOLDID,
       SUBSTR(THRESHOLD_PREDICATE, 1, 20) PREDICATE,
       TIME_OF_VIOLATION,
       SUBSTR(STMTEXT,1,70) STMTEXT
FROM THRESHOLDVIOLATIONS_DB2THRESHOLDVIOLATIONS TV,
     ACTIVITYSTMT_DB2ACTIVITIES A
WHERE TV.APPL_ID = A.APPL_ID
      AND TV.UOW_ID = A.UOW_ID
      AND TV.ACTIVITY_ID = A.ACTIVITY_ID
```

THRESHOLDID	PREDICATE	TIME_OF_VIOLATION	STMTEXT
1	EstimatedSQLCost	2008-09-02-23.04.49.000000	select count(*) from syscat.tables,syscat.tables,syscat.tables

1 record(s) selected.

Available monitoring data

Monitoring data is available from workloads, service subclasses and service superclasses, work classes, and threshold queues. You can use this data to diagnose and correct problems and for performance tuning.

Workload monitoring data

The following figure shows the monitoring information that is available for workloads. You can collect workload statistics and information about activities that run in the workloads using event monitors. For workloads, you can also obtain aggregate activity statistics. You can access workload statistics and information about workload occurrences in real time using table functions.

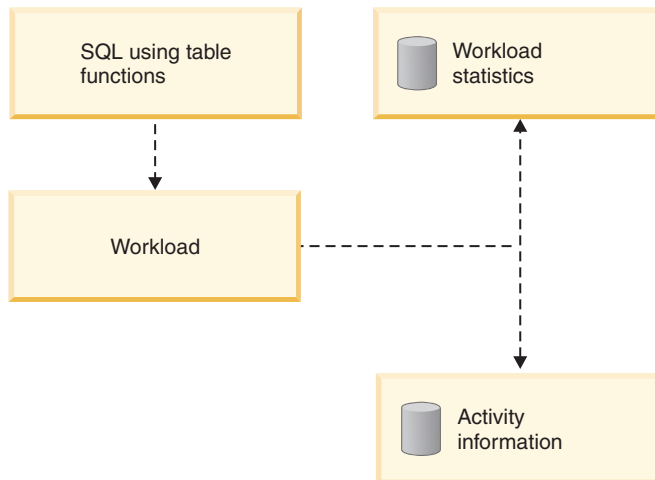


Figure 23. Monitoring data that is available for workloads

Service class monitoring data

The following figure shows the monitoring information that is available for service classes. You can collect statistics for service subclasses and service superclasses. For

service subclasses, you can also obtain aggregate activity and request statistics, and information about activities that run in the service subclass. You can access service superclass and service subclass statistics and information about agents running in a particular service class in real time using table functions.

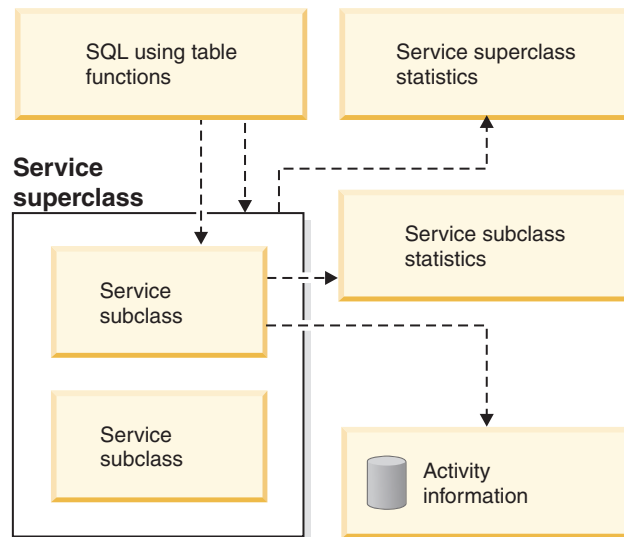


Figure 24. Monitoring data that is available for service classes

Work class monitoring data

The following figure shows the monitoring information that is available for work classes. You can collect work class statistics and information about activities that are associated with a particular work class. You can access work class statistics in real time using table functions.

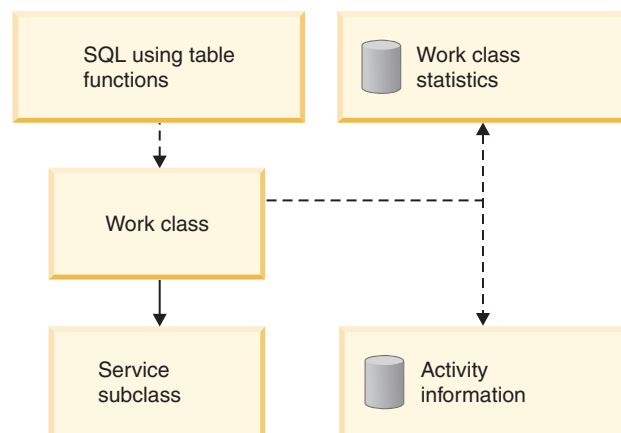


Figure 25. Monitoring data that is available for work classes

Threshold monitoring data

The following figure shows the monitoring information that is available for thresholds. You can obtain information about threshold violations, the activities that caused the threshold violations, and queuing statistics (for queuing thresholds). You can access queuing threshold statistics in real time using table

functions.

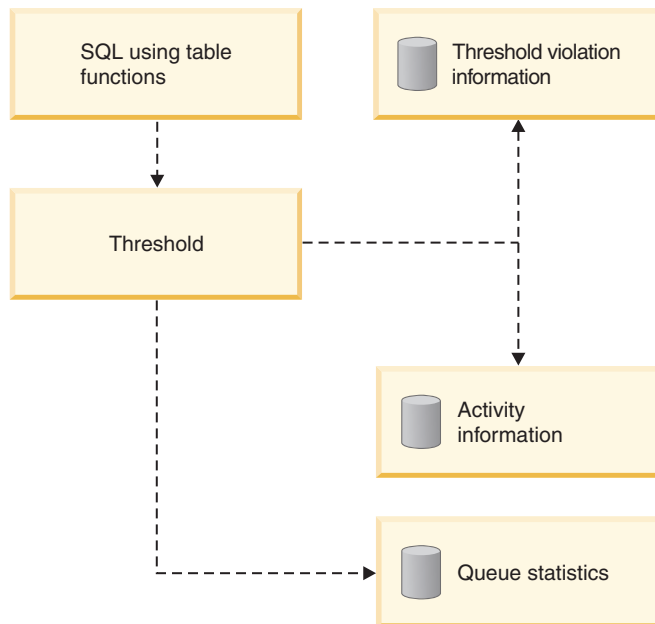


Figure 26. Monitoring data that is available for thresholds

DB2 workload manager stored procedures

You can use stored procedures for canceling an activity, capturing details about an activity, resetting the statistics on DB2 workload manager objects, and setting client information at the data server.

The following stored procedures are available for use with DB2 workload manager:

WLM_CANCEL_ACTIVITY(*application_handle, uow_id, activity_id*)

Use this stored procedure to cancel a running or queued activity. You identify the activity by its application handle, unit of work identifier, and activity identifier. You can cancel any type of activity. The application with the cancelled activity receives the error SQL4725N.

WLM_CAPTURE_ACTIVITY_IN_PROGRESS(*application_handle, uow_id, activity_id*)

Use this stored procedure to send information about an individual activity that is currently executing to the activities event monitor. This stored procedure sends the information immediately, rather than waiting until the activity completes.

WLM_COLLECT_STATS()

Use this stored procedure to collect and reset statistics for DB2 workload manager objects. All statistics tracked for service classes, workloads, threshold queues, and work action sets are sent to the active statistics event monitor (if one exists) and reset. If there is no active statistics event monitor, the statistics are only reset, but not collected.

WLM_SET_CLIENT_INFO(*client_userid, client_wrkstnname, client_applname, client_acctstr, client_workload*)

Use this procedure to set the client information attributes used at the data server to record the identity of the application or end-user currently using the connection. In cases where middleware exists between applications or

users and your data server, use the `WLM_SET_CLIENT_INFO` procedure to set distinguishing connection attributes explicitly.

Statistics for DB2 workload manager objects

Statistics are maintained for DB2 workload manager objects including service classes, work classes, workloads, and threshold queues. These statistics reside in memory and can be viewed in real-time using DB2 workload manager statistics table functions, or the statistics can be collected and sent to a statistics event monitor where they can be viewed later for historical analysis.

Note that you can also obtain monitoring metrics through the statistics event monitor. These are not discussed in this topic, which covers only those statistics that are specific to DB2 workload manager.

When statistics are sent to the event monitor, the values in memory are reset to prevent duplicate data from being collected on subsequent collection intervals. Because the DB2 workload manager statistics table functions report the current in-memory values, following a collection they report the reset values. The DB2 workload manager table functions report only a subset of the statistics. To view the full set of statistics, you must collect the statistics and send them to a statistics event monitor.

Aggregate activity data statistics collection

The following statistics are maintained on the given objects on each database partition, regardless of the value of the `COLLECT AGGREGATE ACTIVITY DATA` option specified for those objects when they are created or altered.

Table 46. Statistics collected for database objects regardless of COLLECT AGGREGATE ACTIVITY DATA setting

Database object	Statistic	Description
Service subclasses	Concurrent activity top (concurrent_act_top)	Use this activity concurrency high watermark to determine the highest concurrency of activities (including nested activities) reached on a database partition for a service class in the time interval for which the statistic is collected.
	Coordinator activities completed total (coord_act_completed_total)	Use this statistic to determine how much work is being performed in a service class.
	Coordinator activities aborted total (coord_act_aborted_total)	Use this statistic, which measures the unsuccessful completion of activities, to determine how healthy the system is. Activities can be aborted because of cancellation, errors, or reactive thresholds.
	Coordinator activities rejected total (coord_act_rejected_total)	Use this rejected non-nested coordinator activity count, which measures the rejection of activities, to obtain an indication of the usefulness of the rejection policy. Activities are counted as rejected when they violate a predictive threshold that has an action of STOP EXECUTION or when they are prevented from executing by a work action.
	The number of active requests (num_requests_active)	Use this statistic to determine the number of requests that are currently executing in a service class.
	The number of activities mapped in and the number of activities mapped out (act_remapped_in and act_remapped_out)	Use these statistics to determine the number of activities that are remapped into or out of a service subclass as part of the priority aging of ongoing activities.
Service superclasses	Concurrent connection top (concurrent_connection_top)	Use this coordinator connection concurrency high watermark to tune a connection concurrency threshold.

Table 46. Statistics collected for database objects regardless of COLLECT AGGREGATE ACTIVITY DATA setting (continued)

Database object	Statistic	Description
Workloads	Concurrent workload occurrences top (concurrent_wlo_top)	Use this workload occurrence high watermark to identify the maximum number of concurrent workload occurrences and to help set or tune a workload occurrence concurrency threshold if the number of concurrently executing workload occurrences is too high (that is, too many applications that are associated with the same workload definition are running on the system at the same time).
	Concurrent workload occurrences activity top (concurrent_wlo_act_top)	Use this element to know the highest number of concurrent activities reached on a partition for any occurrence of this workload in the time interval collected.
	Coordinator activities completed total (coord_act_completed_total)	Use this statistic, which measures the rate of successful completion of activities, to obtain an indication of the health of the system.
	Coordinator activities aborted total (coord_act_aborted_total)	Use this statistic, which measures the unsuccessful completion of activities, to determine how healthy the system is. Activities can be aborted due to cancellation, errors, or reactive thresholds.
	Coordinator activities rejected total (coord_act_rejected_total)	Use this statistic, which measures the rate of rejection of activities, to determine the usefulness of a rejection policy. Activities are counted as rejected when they violate a predictive threshold that has an action of STOP EXECUTION or when they are prevented from executing by a work action.
	Workload occurrences completed total (wlo_completed_total)	Use this statistic to determine how many occurrences of a workload complete in a specific period of time.
	Activities total (act_total)	Use this statistic to determine the effectiveness of the work action set and determine the relative percentages of the types of activities on the system.
Work class (through a work action)	Queue assignments total (queue_assignments_total)	Use this statistic to determine whether excessive queuing is occurring, or whether the right number of activities are being queued (that is, whether the concurrency threshold is too restrictive or not restrictive enough).

Table 46. Statistics collected for database objects regardless of COLLECT AGGREGATE ACTIVITY DATA setting (continued)

Database object	Statistic	Description
Threshold queues	Queue size top (queue_size_top)	Use this statistic to help determine the maximum queue size and to identify whether the queue size is sufficient.
	Queue time total (queue_time_total)	Use this statistic to determine how much time activities are spending in the queue and whether that time is excessive.

When you set the value of the COLLECT AGGREGATE ACTIVITY DATA option to BASE for a service subclass, workload, or a work class (through a work action), some of the following statistics are also collected, or the corresponding histograms are generated for each database partition. Use the averages to quickly understand where activities are spending most of their time (for example, queued or executing) and the response time (lifetime). You can also use the averages to tune the histogram templates. That is you can compare a true average with the average computed from a histogram, and if the average from the histogram deviates from the true average, consider altering the histogram template for the corresponding histogram, using a set of bin values that are more appropriate for your data.

Table 47. Statistics or histograms collected when COLLECT AGGREGATE ACTIVITY DATA is set to BASE

Statistic or histogram	Description
Average request execution time (request_exec_time_avg)	Use this statistic to determine the arithmetic mean of the execution times for requests associated with a service class.
Average coordinator activity lifetime (coord_act_lifetime_avg)	Use this statistic to determine the arithmetic mean of the lifetime for non-nested coordinator activities associated with a service class, workload or a work class.
Average coordinator activity execution time (coord_act_exec_time_avg)	Use this statistic to determine the arithmetic mean of execution time for non-nested coordinator activities associated with a service class, workload or a work class.
Average coordinator activity queue time (coord_act_queue_time_avg)	Use this statistic to determine the arithmetic mean of the queue time for non-nested coordinator activities associated with a service class, workload or a work class.
Cost estimate top (cost_estimate_top)	Use this statistic to tune estimated cost thresholds.
Actual rows returned top (rows_returned_top)	Use the information to tune the actual rows returned thresholds.

Table 47. Statistics or histograms collected when COLLECT AGGREGATE ACTIVITY DATA is set to BASE (continued)

Statistic or histogram	Description
Aggregate temporary table space top (agg_temp_tablespace_top)	<p>Use this statistic to tune aggregate system temporary table space usage.</p> <p>This statistic is monitored only if you define a threshold for aggregate temporary table space usage. For any given service subclass, this statistic is monitored whenever you define a AGGSQLTEMPSPACE threshold on the service subclass itself, or when you define a similar threshold on any service subclass within the same superclass.</p>
Temporary table space top (temp_tablespace_top)	<p>Use this statistic to tune temporary table space usage thresholds.</p> <p>This statistic is monitored only if you define a threshold for temporary table space usage. For any given service subclass, this threshold is also monitored whenever you define a AGGSQLTEMPSPACE threshold on the service subclass itself, or when you define a similar threshold on any service subclass within the same superclass.</p>
Coordinator activity lifetime (CoordActLifetime) histogram	<p>Use this histogram to obtain a view of overall system performance.</p> <p>This histogram collects the time duration between the activity arrival and end time for non-nested coordinator activities.</p> <p>If the activity is a routine that leaves a cursor open after it ends, the lifetime histogram does not count the lifetime of the cursor toward the lifetime of the routine that is the parent of the cursor.</p>

Table 47. Statistics or histograms collected when COLLECT AGGREGATE ACTIVITY DATA is set to BASE (continued)

Statistic or histogram	Description
Coordinator activity execution time (CoordActExecTime) histogram	<p>Use this histogram to measure the impact of changes to the system that affect execution time.</p> <p>This histogram collects the execution time for non-nested coordinator activities.</p> <p>The execution time is calculated as follows:</p> <ul style="list-style-type: none"> • For cursors, the execution time is the combined time for the open cursor request, any fetches, and the close cursor request. Time when the cursor is idle is not counted towards the execution time. • For routines, the execution time is from the start to the end of the routine invocation. If any cursors are left open by the routine after it ends, the lifetimes of these cursors are not counted towards the routine execution time. • For all other activities, the execution time is the difference between the activity lifetime and the time that the activity spends queued.
Coordinator activity queue time (CoordActQueueTime) histogram	<p>Use this histogram to measure the impact of queuing thresholds on activities.</p> <p>This histogram collects the amount of time that non-nested coordinator activities spend queued.</p>

When you set the value of the COLLECT AGGREGATE ACTIVITY DATA option to EXTENDED for a service subclass, workload or a work class, the following system statistics are collected or histograms are generated for each database partition for the corresponding service class or work class (through a work action). Use the averages to quickly understand the average rate of arrival of activities (arrival rate is the inverse of inter-arrival time) and the expense of activities (estimated cost). You can also use the averages to tune the histogram templates. That is you can compare a true average with the average computed from a histogram, and if the average from the histogram deviates from the true average, consider altering the histogram template for the corresponding histogram, using a set of bin values that are more appropriate for your data. EXTENDED statistics are useful for more detailed performance modelling. Also see “Workload management performance modelling” on page 199.

Table 48. Statistics or histograms collected when COLLECT AGGREGATE ACTIVITY DATA is set to EXTENDED

Statistic or histogram	Description
Coordinator activity estimated cost average (coord_act_est_cost_avg)	Use this statistic to determine the arithmetic mean of the estimated costs of coordinator DML activities at nesting level 0 that are associated with this service subclass, workload or work class since the last statistics reset.

Table 48. Statistics or histograms collected when COLLECT AGGREGATE ACTIVITY DATA is set to EXTENDED (continued)

Statistic or histogram	Description
Non-nested coordinator activity inter-arrival time average (coord_act_interarrival_time_avg)	Use this statistic to determine the arithmetic mean of the time between the arrival of one coordinator activity at nesting level 0 that is associated with this service class, workload or work class and the next coordinator activity to arrive. The average is computed since the last statistics reset.
Coordinator activity estimated cost (CoordActEstCost) histogram.	Use this histogram to obtain an approximate service time distribution. This histogram collects the estimated cost for non-nested coordinator activities. This data is useful for modelling your system or for inputting into performance-modelling applications.
Coordinator activity inter-arrival time (CoordActInterArrivalTime) histogram.	Use this histogram to obtain the inter-arrival time distribution for non-nested coordinator activities. This histogram collects the inter-arrival time for non-nested coordinator activities. This data is useful for modelling your system or for inputting into performance-modeling applications.

The following table provides a reference for which activity statistics are collected for each DB2 workload manager object and includes all aggregate statistics available to you from both table functions and event monitors. Some statistics are always collected for some objects. Other statistics are only collected when a particular COLLECT AGGREGATE option is specified. For aggregate activity statistics, if COLLECT AGGREGATE ACTIVITY DATA EXTENDED is specified, all the BASE aggregate activity statistics are also collected.

Table 49. Aggregate activity statistics collection for DB2 workload manager objects

Object type	Activity statistics always collected by default	Activity statistics collected when you specify COLLECT AGGREGATE ACTIVITY DATA BASE	Activity statistics collected when you specify COLLECT AGGREGATE ACTIVITY DATA EXTENDED
Service subclass	act_remapped_in act_remapped_out concurrent_act_top coord_act_completed_total coord_act_rejected_total coord_act_aborted_total	agg_temp_tablespace_top coord_act_exec_time_avg coord_act_lifetime_avg coord_act_lifetime_top coord_act_queue_time_avg coord_act_lifetime_stddev coord_act_exec_time_stddev coord_act_queue_time_stddev CoordActLifetime histogram CoordActExecTime histogram CoordActQueueTime histogram cost_estimate_top rows_returned_top temp_tablespace_top	coord_act_est_cost_avg coord_act_interarrival_time_avg CoordActEstCost histogram CoordActInterArrivalTime histogram

Table 49. Aggregate activity statistics collection for DB2 workload manager objects (continued)

Object type	Activity statistics always collected by default	Activity statistics collected when you specify COLLECT AGGREGATE ACTIVITY DATA BASE	Activity statistics collected when you specify COLLECT AGGREGATE ACTIVITY DATA EXTENDED
Service superclass	concurrent_connection_top	N/A	N/A
Workload	concurrent_wlo_act_top concurrent_wlo_top coord_act_aborted_total coord_act_completed_total coord_act_rejected_total wlo_completed_total	coord_act_exec_time_avg coord_act_lifetime_top coord_act_lifetime_avg coord_act_queue_time_avg coord_act_lifetime_stddev coord_act_exec_time_stddev coord_act_queue_time_stddev CoordActLifetime histogram CoordActExecTime histogram CoordActQueueTime histogram cost_estimate_top rows_returned_top temp_tablespace_top	coord_act_est_cost_avg coord_act_interarrival_time_avg CoordActEstCost histogram CoordActInterArrivalTime histogram
Work class (through a work action)	act_total	agg_temp_tablespace_top coord_act_lifetime_top coord_act_lifetime_avg coord_act_exec_time_avg coord_act_queue_time_avg CoordActLifetime histogram CoordActExecTime histogram CoordActQueueTime histogram cost_estimate_top rows_returned_top temp_tablespace_top	coord_act_est_cost_avg coord_act_interarrival_time_avg CoordActEstCost histogram CoordActInterArrivalTime histogram
Threshold	N/A	N/A	N/A
Threshold queue	queue_assignments_total queue_size_top queue_time_total	N/A	N/A

Aggregate request data statistics collection

When you set the value of the COLLECT AGGREGATE REQUEST DATA option for a service subclass to BASE, the following statistics are maintained for the service subclass.

Table 50. Statistics or histograms collected when COLLECT AGGREGATE REQUEST DATA is set to BASE

Statistic or histogram	Description
Request execution time average (request_exec_time_avg)	Use this statistic to quickly understand the average amount of time that is spent processing each request on a database partition and to help tune the histogram template for the corresponding request execution time histogram.

Table 50. Statistics or histograms collected when COLLECT AGGREGATE REQUEST DATA is set to BASE (continued)

Statistic or histogram	Description
Request execution time (ReqExecTime) histogram	<p>Use this histogram to understand where work is being performed and whether the distribution of work across partitions is uniform.</p> <p>This histogram indicates the volume of work executing in a service class and the distribution of this work across database partitions. The execution time for requests is collected in a histogram for each database partition and for all requests.</p> <p>This histogram includes requests on the coordinator partition, and any subrequests on both coordinator and non-coordinator partitions (like RPC requests or SMP subagent requests). Requests included may or may not be associated with an activity: Both PREPARE and OPEN requests are included in this histogram, for example, but while OPEN requests are always associated with a cursor activity, PREPARE requests are not part of any activity.</p> <p>The request execution time approximates the effort spent by agents working in a service class. For example, coordinator activity counts might show that most user activities originate on one database partition, but as part of processing the activities, the coordinator agent might be sending subrequests to another database partition that performs most of the work.</p> <p>The request execution time histogram can be useful in determining the size of requests sent to a database partition, that is, whether the work that is sent to the database partition consists of mostly small requests or mostly large requests or whether there is no specific distribution.</p> <p>Request execution time histograms should not be used for activity response time analysis, because activities may be composed of a number of requests and subrequests, because there is no one-to-one mapping between request and activity execution time, and because not all requests are associated with activities.</p>

The following table provides a reference for which request statistics are collected for each DB2 workload manager object and includes all aggregate statistics available to you from both table functions and event monitors. Some statistics are always collected for some objects. Other statistics are only collected when the COLLECT AGGREGATE REQUEST DATA option is specified.

Table 51. Aggregate request statistics collection for DB2 workload manager objects

Object type	Request statistics always collected by default	Request statistics collected when you specify COLLECT AGGREGATE REQUEST DATA BASE
Service subclass	num_requests_active	request_exec_time_avg request_exec_time_stddev request_exec_time_total ReqExecTime histogram
Service superclass	N/A	N/A
Workload	N/A	N/A
Work class (through a work action)	N/A	N/A
Threshold	N/A	N/A
Threshold queue	N/A	N/A

Statistics collection and monitoring with priority aging

How you collect statistics and how you monitor are both affected by the dynamic remapping of activities between service subclasses, also known as priority aging. As a rule, the activity affects the aggregate statistics of the service subclass that it finishes running in (with exceptions), and some statistics of service subclasses that the activity passes through during its lifetime.

Statistics affected by remapped activities

As one exception to the rule, the activity interarrival time, estimated cost, and queue time are all associated with the subclass in which an activity starts running, rather than with the subclass in which the activity finishes running. Because a remapped activity affects the statistics collection of both subclasses, a different number of activities can be counted in an interarrival time, an estimated cost, or a queue-time histogram than in a lifetime or execution-time histogram.

For example, consider an activity that starts running in service subclass A and later is remapped to service subclass B, in which it finishes running. The estimated cost of this activity is associated with service subclass A, but its lifetime is associated with service subclass B. As a result, for subclass A, the estimated cost histogram has one more element counted in it than the lifetime histogram has counted in it, and for service subclass B, the lifetime histogram has one more element counted in it than the estimated cost histogram has counted in it.

As a second exception to the rule, the monitor element **concurrent_act_top** can be updated in and attributed to any subclass that an activity passes through. In addition to being incremented when an activity begins and decremented when an activity ends, the monitor element is incremented when an activity is mapped to the subclass and is decremented when an activity is mapped out of the subclass (mapped to a different subclass).

Statistics about activity remapping

You can use two monitor elements to count the number of activities entering or leaving a service subclass because of a remapping action: **act_remapped_in** and **act_remapped_out**. The **act_remapped_in** and **act_remapped_out** monitor elements count the number of activities for any given subclass at any partition that were

mapped into or out of the subclass since the last reset. You can use these monitor elements to validate that the remapping of activities between service subclasses is occurring as expected.

To determine the source and destination service subclasses targeted by a remapping action, you can refer to the threshold violation event monitor record, which includes a destination service class ID (`destination_service_class_id`). You can also determine the source service class by using the threshold violation record.

Monitoring with activity remapping

Remapping activities to different subclasses affects how you monitor these activities. To ensure that all statistics are collected for an activity that starts in one service class and finishes in another because of remapping, turn on aggregate activity data collection for both the service subclass in which the activity starts running and the service subclass in which the activity finishes running when you create or alter the service classes. If you turn on aggregate activity data collection for only the service subclass in which the activity started, the activity contributes only to queue time statistics and, in the case of extended statistics, to the estimated cost and interarrival time statistics. If you turn on aggregate activity data collection for only the service subclass in which the activity finishes running, the activity contributes only to lifetime and execution time statistics, regardless of whether the option is `COLLECT AGGREGATE DATA BASE` or `COLLECT AGGREGATE DATA EXTENDED` when you issue the `CREATE SERVICE CLASS` or `ALTER SERVICE CLASS` statement.

The following tables summarize how statistics collection is affected by remapping and collection settings.

Table 52. Effect of the `COLLECT AGGREGATE DATA BASE` option on aggregate statistics collection for subclasses involved in remapping

Statistics	Starting subclass collection setting and ending subclass collection setting			
	NONE and NONE	BASE and NONE	NONE and BASE	BASE and BASE
Lifetime	Not collected	Not collected	Collected	Collected
Queue time	Not collected	Collected	Not collected	Collected
Execution time	Not collected	Not collected	Collected	Collected

Table 53. Effect of the `COLLECT AGGREGATE DATA EXTENDED` option on aggregate statistics collection for subclasses involved in remapping

Statistics	Starting subclass collection setting and ending subclass collection setting			
	NONE and NONE	EXTENDED and NONE	NONE and EXTENDED	EXTENDED and EXTENDED
Lifetime	Not collected	Not collected	Collected	Collected
Queue time	Not collected	Collected	Not collected	Collected
Execution time	Not collected	Not collected	Collected	Collected
Inter-arrival time	Not collected	Collected	Not collected	Collected
Estimated cost	Not collected	Collected	Not collected	Collected

Table 54. Effect of mixing the COLLECT AGGREGATE DATA BASE and the COLLECT AGGREGATE DATA EXTENDED options on aggregate statistics collection for subclasses involved in remapping

Statistics	Starting subclass collection setting and ending subclass collection setting	
	BASE and EXTENDED	EXTENDED and BASE
Lifetime	Collected	Collected
Queue time	Collected	Collected
Execution time	Collected	Collected
Interarrival time	Not collected	Collected
Estimated cost	Not collected	Collected

Histograms in workload management

A histogram is a collection of bins, which are containers for collecting discrete ranges of data. Histograms are useful for a variety of workload analysis and performance-tuning tasks.

DB2 workload manager histograms have a fixed number of 41 bins. The 40th bin contains the highest defined value for the histogram, and the 41st bin is for values that are beyond the highest defined value. The following figure shows a histogram of activity lifetimes that are plotted using a bar chart.




Figure 27. Histogram of activity lifetimes that are plotted using a bar chart

The activity lifetime histogram corresponds to the following data. Each count represents the number of activities whose lifetimes (in milliseconds) are within the range of the low bin value to the high bin value. For example, 156 activities had a lifetime in the range of 68 milliseconds to 103 milliseconds.

Low Bin	High Bin	Count
0	1	0
1	2	0
2	3	0
3	5	0
5	8	0

8	12	0
12	19	0
19	29	10
29	44	15
44	68	45
68	103	156
103	158	65
158	241	23
241	369	0
369	562	0
562	858	0
858	1309	0
1309	1997	0
1997	3046	0
3046	4647	0
4647	7089	0
7089	10813	0
10813	16493	0
16493	25157	0
25157	38373	0
38373	58532	0
58532	89280	0
89280	136181	0
136181	207720	0
207720	316840	0
316840	483283	3
483283	737162	0
737162	1124409	0
1124409	1715085	0
1715085	2616055	0
2616055	3990325	0
3990325	6086529	0
6086529	9283913	0
9283913	14160950	0
14160950	21600000	0
21600000	Infinity	0

You can use histograms for a number of different purposes. For example, you can use them to see the distribution of values, use them to identify outlying values, or use them to compute averages and standard deviations. See “Scenario: Tuning a DB2 workload manager configuration when capacity planning information is unavailable” on page 271 and “Example: Computing averages and a standard deviation from histograms in a DB2 workload manager configuration” on page 186 for examples of how to use histograms to better understand and characterize your workload.

In a partitioned database environment, histograms are collected on each database partition. Histogram bins have the same range of values on all database partitions, with specific counts per bin per partition. You can use the bins to analyze information on a per-partition basis. You can also combine the histograms from all database partitions by adding the counts in the corresponding bins and use this single histogram to obtain a global view of the data, which you can then use for tasks such as calculating the global average and standard deviation.

Histograms are available for service subclasses, workloads, and work classes, through work actions. Histograms are collected for these objects when you specify one of the COLLECT AGGREGATE ACTIVITY DATA clauses when creating or altering the objects. For work classes, histograms are also collected if you apply a COLLECT AGGREGATE ACTIVITY DATA work action to the work class. The following histograms are available:

- Non-nested coordinator activity lifetime, when you specify AGGREGATE ACTIVITY DATA BASE or AGGREGATE ACTIVITY DATA EXTENDED for a service subclass, for a workload, or for a work action applied to a work class
- Non-nested coordinator activity execution time, when you specify AGGREGATE ACTIVITY DATA BASE or AGGREGATE ACTIVITY DATA EXTENDED for a service subclass, for a workload, or for a work action applied to a work class
- Non-nested coordinator activity queue time, when you specify AGGREGATE ACTIVITY DATA BASE or AGGREGATE ACTIVITY DATA EXTENDED for a service subclass, for a workload, or for a work action applied to a work class
- Request execution time, when you specify AGGREGATE REQUEST DATA BASE for a service subclass; this histogram does not apply to workloads or work classes
- Non-nested activity interarrival time histogram, when you specify AGGREGATE ACTIVITY DATA EXTENDED for a service subclass, for a workload, or for a work action applied to a work class
- Non-nested DML activity estimated cost, when you specify AGGREGATE ACTIVITY DATA EXTENDED for a service subclass, for a workload, or for a work action applied to a work class

All activity-related histograms contain information about activities that are completed, are canceled, or are rejected.

Histogram templates

You can optionally specify a histogram template that is used to determine what a particular histogram looks like, including the high bin value. A histogram template is a *unitless* object, meaning that there is no predefined measurement unit assigned to it.

You can create a histogram template by using the CREATE HISTOGRAM TEMPLATE statement, specifying the maximum high bin value. All other bin values are automatically defined as exponentially increasing values that approach the high bin value. For example, to create a histogram template with a high bin value of 3 000 000, issue a statement such as the following one:

```
CREATE HISTOGRAM TEMPLATE TEMPLATE1 HIGH BIN VALUE 3000000
```

This statement creates a histogram template with the following bin values:

Low Bin	High Bin
0	1
1	2
2	3
3	4
4	6
6	9
9	13
13	19
19	28
28	41
41	60
60	87
87	127
127	184
184	268
268	389
389	565
565	821
821	1192
1192	1732

1732	2514
2514	3651
3651	5300
5300	7696
7696	11173
11173	16222
16222	23553
23553	34196
34196	49649
49649	72084
72084	104657
104657	151948
151948	220609
220609	320297
320297	465030
465030	675163
675163	980250
980250	1423197
1423197	2066299
2066299	3000000
3000000	Infinity

You apply a histogram template by using the appropriate HISTOGRAM TEMPLATE keyword when creating or altering service subclasses, workloads, or work actions. If you do not specify a histogram template, the default template, SYSDEFAULTHISTOGRAM, is used. If you do not enable AGGREGATE ACTIVITY DATA collection for an object, the histogram template is ignored.

For example, to use the TEMPLATE1 histogram template for the existing activity lifetime histogram of service subclass MYSUBCLASS under the service superclass MYSUPERCLASS, issue the following statement:

```
ALTER SERVICE CLASS MYSUBCLASS UNDER MYSUPERCLASS
ACTIVITY LIFETIME HISTOGRAM TEMPLATE TEMPLATE1
```

After you commit the ALTER SERVICE CLASS statement, the activity lifetime histogram that is collected for the MYSUBCLASS service subclass has high bin values that are determined by the TEMPLATE1 histogram template instead of by the SYSDEFAULTHISTOGRAM histogram template.

If you change a service class or a workload to use a different histogram template or change a histogram template, the change does not take effect until a statistics reset occurs.

You can drop a histogram template by using the DROP HISTOGRAM TEMPLATE statement.

You can view the histogram templates by querying the SYSCAT.HISTOGRAMTEMPLATES view and view the corresponding histogram template high bin values by querying the SYSCAT.HISTOGRAMTEMPLATEBINS view. The low bin value is always 0 for the first bin; for any other bins, the low bin value is the high bin value from the preceding bin.

Example

The following example creates a table function to compute the CoordActLifetime, CoordActExecTime, CoordActQueueTime, or CoordActEstCost histogram for a service superclass as a whole by summing across the subclasses. Summing across subclasses is useful when activities are remapped to different service subclasses under the same service superclass during execution, as can occur under a priority

aging scenario where service class tiers and specialized thresholds are used to control resources for activities dynamically. This example does not apply to the CoordActInterArrivalTime histogram because the weighted averages computed do not account for the fact that the CoordActInterArrivalTime histogram of a subclass measures the time between the arrival of a query in that subclass and the next query but the CoordActInterArrivalTime histogram of a superclass measures the time between the arrival of a query in any of its subclasses and the next query.

```
CONNECT TO SAMPLE

DROP FUNCTION histsuper

CREATE FUNCTION histsuper(superclass varchar(128),
                        histogram_type varchar(24))
RETURNS TABLE (statistics_timestamp timestamp,
                bin_top integer,
                number_in_bin integer,
                graph varchar(60))

LANGUAGE SQL
READS SQL DATA
NO EXTERNAL ACTION
DETERMINISTIC
RETURN WITH HISTOGRAMS AS
  (SELECT HISTOGRAM_TYPE,
         substr(PARENTSERVICECLASSNAME,1,26) as SUPERCLASS,
         STATISTICS_TIMESTAMP,
         TOP as BIN_TOP,
         sum(NUMBER_IN_BIN) as NUMBER_IN_BIN
   FROM HISTOGRAMBIN_DB2STATISTICS H,
        SYSCAT.SERVICECLASSES S
   WHERE H.SERVICE_CLASS_ID = S.SERVICECLASSID
        AND PARENTSERVICECLASSNAME = histsuper.superclass
        AND HISTOGRAM_TYPE = histsuper.histogram_type
        AND HISTOGRAM_TYPE IN ('CoordActLifetime', 'CoordActExecTime',
'CoordActQueueTime', 'CoordActEstCost')
   GROUP BY HISTOGRAM_TYPE, PARENTSERVICECLASSNAME, STATISTICS_TIMESTAMP, TOP)
  SELECT STATISTICS_TIMESTAMP,
         BIN_TOP,
         NUMBER_IN_BIN,
         substr(repeat('#', cast(NUMBER_IN_BIN * 60 /
         (SELECT CASE WHEN MAX(NUMBER_IN_BIN) = 0 THEN 1
         ELSE MAX(NUMBER_IN_BIN) END FROM HISTOGRAMS) AS INTEGER)),1,60)
   AS GRAPH FROM HISTOGRAMS

CONNECT RESET
```

The output looks as follows:

STATISTICS_TIMESTAMP	BIN_TOP	NUMBER_IN_BIN	GRAPH
2008-11-06-14.47.08.833188	-1	0	
2008-11-06-14.47.08.833188	1	1	
2008-11-06-14.47.08.833188	2	1	
2008-11-06-14.47.08.833188	3	2	
2008-11-06-14.47.08.833188	5	4	
2008-11-06-14.47.08.833188	8	7	
2008-11-06-14.47.08.833188	12	15	
2008-11-06-14.47.08.833188	19	29	#
2008-11-06-14.47.08.833188	29	41	#
2008-11-06-14.47.08.833188	44	67	##
2008-11-06-14.47.08.833188	68	112	###
2008-11-06-14.47.08.833188	103	228	####
2008-11-06-14.47.08.833188	158	335	#####
2008-11-06-14.47.08.833188	241	723	#####
2008-11-06-14.47.08.833188	369	1289	#####
2008-11-06-14.47.08.833188	562	1890	#####
2008-11-06-14.47.08.833188	858	2484	#####
2008-11-06-14.47.08.833188	1309	1943	#####
2008-11-06-14.47.08.833188	1997	478	#####
2008-11-06-14.47.08.833188	3046	221	#####
2008-11-06-14.47.08.833188	4647	29	#
2008-11-06-14.47.08.833188	7089	7	

2008-11-06-14.47.08.833188	10813	0
2008-11-06-14.47.08.833188	16493	2
2008-11-06-14.47.08.833188	25157	0
2008-11-06-14.47.08.833188	38373	1
2008-11-06-14.47.08.833188	58532	0
2008-11-06-14.47.08.833188	89280	0
2008-11-06-14.47.08.833188	136181	0
2008-11-06-14.47.08.833188	207720	0
2008-11-06-14.47.08.833188	316840	0
2008-11-06-14.47.08.833188	483283	0
2008-11-06-14.47.08.833188	737162	0
2008-11-06-14.47.08.833188	1124409	0
2008-11-06-14.47.08.833188	1715085	0
2008-11-06-14.47.08.833188	2616055	0
2008-11-06-14.47.08.833188	3990325	0
2008-11-06-14.47.08.833188	6086529	0
2008-11-06-14.47.08.833188	9283913	0
2008-11-06-14.47.08.833188	14160950	0
2008-11-06-14.47.08.833188	21600000	0

41 record(s) selected.

Creating a histogram template

Use the CREATE HISTOGRAM TEMPLATE statement to create a histogram template. Histogram templates are used by service subclasses and work actions to define the bin values for the statistics that are maintained using histograms.

To create a histogram template, you require WLMADM or DBADM authority.

See the following topics for more information about prerequisites:

- “DDL statements for DB2 workload manager” on page 16
- Naming rules

Some DB2 service subclass, work class activity, and request statistics are collected using histograms. All histograms have a set number of bins, and each bin represents a range in which activities or requests are counted. The type of units used for the bins depends on the type of histogram that you create. The histogram template describes the high value of the second-to-last bin in the histogram, which affects the values of all of the bins in the histogram. For more information on histograms, see “Histograms in workload management” on page 180.

To create a histogram template:

1. Issue the CREATE HISTOGRAM TEMPLATE statement, specifying the name of the histogram template that you want to create and a value for the HIGH BIN VALUE keyword to set the top value for the second-to-last bin.
2. Commit your changes. When you commit your changes, the histogram is added to the SYSCAT.HISTOGRAMTEMPLATES view and the bins are added to the SYSCAT.HISTOGRAMTEMPLATEBINS view.

Altering a histogram template

Use the ALTER HISTOGRAM TEMPLATE statement to alter an existing histogram template. Histogram templates are used by service subclasses and work actions to define the bin values for the statistics that are maintained using histograms.

You require WLMADM or DBADM authority to alter a histogram template.

See “DDL statements for DB2 workload manager” on page 16 for more information about prerequisites.

Some DB2 service subclass, work class activity, and request statistics are collected using histograms. All histograms have a set number of bins, and each bin

represents a range in which activities or requests are counted. The type of units used for the bins depends on the type of histogram that you create. The histogram template describes the high value of the second-to-last bin in the histogram, which affects the values of all of the bins in the histogram. For more information on histograms, see “Histograms in workload management” on page 180.

To alter a histogram template:

1. Issue the ALTER HISTOGRAM TEMPLATE statement, specifying the name of the histogram template that you want to alter and a value for the HIGH BIN VALUE parameter to alter the top value for the second-to-last bin.
2. Commit your changes. When you commit your changes the high bin value for the histogram is updated in the SYSCAT.HISTOGRAMTEMPLATEBINS view. The change does not take effect until the next time the workload management statistics are reset. See “Resetting statistics on DB2 workload manager objects” on page 190 for more information.
3. Optional: Run the WLM_COLLECT_STATS stored procedure to collect and reset the statistics so that the new histogram template is used immediately.

Dropping a histogram template

You can drop a histogram template if you no longer require it.

To drop a histogram template, you require WLMADM or DBADM authority.

See “DDL statements for DB2 workload manager” on page 16 for more information about prerequisites.

You cannot drop the SYSDEFAULTHISTOGRAM histogram template.

You cannot drop a histogram template if it is being referenced by a service subclass, work action, or workload. You can view the service subclasses and work actions that reference a histogram template by querying the SYSCAT.HISTOGRAMTEMPLATESUSE view.

To drop a histogram template:

1. Use the DROP HISTOGRAM TEMPLATE statement.
2. Commit your changes. When you commit your changes the histogram is removed from the SYSCAT.HISTOGRAMTEMPLATES view, and its bins are removed from the SYSCAT.HISTOGRAMTEMPLATEBINS view.

Example: Computing averages and a standard deviation from histograms in a DB2 workload manager configuration

One use for histograms is for obtaining the standard deviation for activity lifetimes. The example in this topic shows how bins are used for the calculation of this statistic.

A calculation of the average lifetime for each activity is a useful piece of information. However, the average alone does not accurately describe the user experience. If the variability in activity lifetime is large, the users whom you are supporting might see queries run fast at some times (which is fine) and slow at others (which might not be acceptable). When you define a goal for activity lifetimes, not only is the average lifetime of the activities important but also the standard deviation of the activity lifetime. You need to both understand and control variability to ensure that your users actually experience the observed average.

In a DB2 workload manager configuration, statistics are collected on each database partition. The following example shows how to obtain the average activity lifetime for a single database partition.

Suppose that you have a single-partition environment and histogram with the following bins. There are more bins in the real histograms, but this example is limited to eight bins to make the example simpler.

- Bin 1 - 0 to 2 seconds
- Bin 2 - 2 to 4 seconds
- Bin 3 - 4 to 8 seconds
- Bin 4 - 8 to 16 seconds
- Bin 5 - 16 to 32 seconds
- Bin 6 - 32 to 64 seconds
- Bin 7 - 64 to 128 seconds
- Bin 8 - 128 seconds to infinity

You can compute an approximation of the average by assuming that the average response time for a query that falls into a bin with the range x to y is $(x + y)/2$. You can then multiply this number by the number of queries that fell into the bin, sum across all bins, then divide the sum by the total count. For the preceding example, assume that the average response time for each bin is:

- Bin 1 average lifetime = $(0+2)/2 = 1$
- Bin 2 average lifetime = $(2+4)/2 = 3$
- Bin 3 average lifetime = $(4+8)/2 = 6$
- Bin 4 average lifetime = $(8+16)/2 = 12$
- Bin 5 average lifetime = $(16+32)/2 = 24$
- Bin 6 average lifetime = $(32+64)/2 = 48$
- Bin 7 average lifetime = $(64+128)/2 = 96$

Assume that the following histogram was collected during the measurement period:

Bin 1	Bin 2	Bin 3	Bin 4	Bin 5	Bin 6	Bin 7	Bin 8
count	count	count	count	count	count	count	count
20	30	80	10	5	3	2	0

To calculate average lifetime, bin 8 must be empty. Bin 8 only exists to let you know when you need to change the upper boundary of your range. For this reason, you must specify the upper bound for the range.

You can approximate the average lifetime for database partition 1 as follows:

$$\begin{aligned}
 \text{average lifetime} &= (20 \times 1 + 30 \times 3 + 80 \times 6 + 10 \times 12 + 5 \times 24 + 3 \times 48 + 2 \times 96) / 150 \\
 &= (20 + 90 + 480 + 120 + 120 + 144 + 192) / 150 \\
 &= 1166 / 150 \\
 &= 7.77 \text{ seconds}
 \end{aligned}$$

You can approximate the lifetime standard deviation as follows:

$$\text{Standard deviation} = [(20 \times (1 - 7.77)^2 + 30 \times (3 - 7.77)^2 + \dots) / 150]^{1/2}$$

For partitioned database environments, averages and standard deviations can be computed by first computing a combined histogram across all database partitions by adding the counts of each bin across the database partitions.

For example, assume that the database has two partitions, the histogram bin sizes are as described above, and the histogram has the following data:

Database partition	Bin 1 count	Bin 2 count	Bin 3 count	Bin 4 count	Bin 5 count	Bin 6 count	Bin 7 count	Bin 8 count
1	20	30	80	10	5	3	2	0
2	1	5	20	20	4	0	0	0

Because the bin sizes are the same across all database partitions, the overall histogram is easy to compute:

Bin 1	Bin 2	Bin 3	Bin 4	Bin 5	Bin 6	Bin 7	Bin 8
count	count	count	count	count	count	count	count
21	35	100	30	9	3	2	0

From the combined histogram, you can calculate the overall lifetime average and standard deviation in a similar way to how they were computed for a single-partition environment:

$$\begin{aligned}
 \text{Average lifetime} &= (21 \times 1 + 35 \times 3 + 100 \times 6 + 30 \times 12 + 9 \times 24 + 3 \times 48 + 2 \times 96) / 200 \\
 &= (21 + 105 + 600 + 360 + 216 + 144 + 192) / 200 \\
 &= 1638 / 200 \\
 &= 8.19 \text{ seconds}
 \end{aligned}$$

$$\text{Standard deviation} = [(21 \times (1 - 8.19)^2 + 35 \times (3 - 7.77)^2 + \dots) / 200]^{1/2}$$

Historical analysis tool

Your DB2 data server installation includes a pair of Perl scripts as a sample that generate information about which tables, indexes, and columns have or have not been accessed by performing historical analysis.

These scripts provide historical analysis functionality similar to the Query Patroller historical analysis feature by using information captured by the workload management activities event monitor. The workload management historical analysis tool was written in Perl; you can use these scripts as is or you can modify them to produce additional historical analysis reports to suit your needs.

The workload management historical analysis tool consists of two scripts, which can be found in the `samples/perl` path of your installation directory:

- `wlmhist.pl` - generates historical data
- `wlmhistrep.pl` - produces reports from the historical data.

A `DB2WlmHist.pm` file, which contains common Perl routines used by the two scripts, is included also.

Refer to the `README_WLMHIST` file found in the same file directory for more information on how to set up and run the scripts.

Collecting workload management statistics using a statistics event monitor

Statistics for DB2 workload manager objects can be sent to a statistics event monitor for historical analysis.

You can use statistics to understand the behavior of your system over time (for example, what is the average lifetime of activities, how much time do activities spend queued, what is the distribution of large compared to small activities, and so on), set thresholds (for example, find the upper boundary for concurrent activities), and detect problems (for example, detect whether the average lifetime that users are experiencing is higher than normal). See "Statistics for DB2 workload manager objects" on page 169 for a description of which statistics are collected for each DB2 workload manager object.

You can automatically send workload management statistics to an event monitor on a fixed interval of time, or you can manually send statistics to an event monitor at any point in time.

To automatically collect workload management statistics on a fixed time interval:

1. Use the CREATE EVENT MONITOR statement to create a STATISTICS event monitor. For example, you could issue the following statement:
CREATE EVENT MONITOR STATS1 FOR STATISTICS WRITE TO TABLE
2. Use the COMMIT statement to commit your changes.
3. Use the SET EVENT MONITOR STATE statement to activate the event monitor. Instead of using the SET EVENT MONITOR STATE statement, you can use the AUTOSTART default for the STATISTICS event monitor to have it activated the next time that the database is activated. If you want to define multiple STATISTICS event monitors, you should not use the AUTOSTART option.
4. Use the COMMIT statement to commit your changes.
5. Optional: Enable the collection of additional statistics. By default, only a minimal set of statistics is collected for each DB2 workload manager object. See “Statistics for DB2 workload manager objects” on page 169 for details on which statistics are collected by default for each object. Specify the collection of aggregate activity data for service subclasses, workloads, and work classes using the COLLECT AGGREGATE ACTIVITY DATA keyword on the ALTER SERVICE CLASS and ALTER WORK ACTION SET statements. Specify the collection of aggregate request data for service subclasses using the COLLECT AGGREGATE REQUEST DATA keyword on the ALTER SERVICE CLASS statement. COMMIT any changes.
6. Specify a collection interval by updating the database configuration parameter **wlm_collect_int**. The **wlm_collect_int** parameter specifies an interval of time in minutes. Every interval, the in-memory copy of the workload management statistics for all DB2 workload manager objects is written to the active statistics event monitor and the in-memory statistics are reset. In a partitioned database environment, the **wlm_collect_int** parameter must be updated on the catalog partition. This parameter can be updated dynamically. For example:

```
CONNECT TO database alias
UPDATE DATABASE CONFIGURATION USING WLM_COLLECT_INT 5 IMMEDIATE
```

After you perform the preceding steps, workload management statistics are written to the statistics event monitor every **wlm_collect_int** minutes. Each record written to the statistics event monitor has a STATISTICS_TIMESTAMP value and a LAST_WLM_RESET value. The interval of time from LAST_WLM_RESET to STATISTICS_TIMESTAMP defines the collection interval (that is, interval of time over which the statistics in that record were collected).

If the **wlm_collect_int** parameter is set to a nonzero value and there is no active statistics event monitor, the in-memory workload management statistics are still reset every **wlm_collect_int** minutes, but statistics are not collected. The data will be lost. For this reason, it is not recommended that you specify a nonzero **wlm_collect_int** value without activating a statistics event monitor.

If the **wlm_collect_int** parameter is set to 0 (the default) statistics are not sent to the statistics event monitor automatically. You can manually send statistics to the statistics event monitor for later historical analysis by using the WLM_COLLECT_STATS stored procedure. When this procedure is invoked, it performs the same actions that occur with an automatic statistics collection interval. That is, the in-memory statistics are sent to the statistics event monitor and the in-memory statistics are reset. If there is no active statistics event monitor, the in-memory values are reset, but data is not collected. If you only want to reset statistics, you can invoke the WLM_COLLECT_STATS procedure while there is no active statistics event monitor.

Manual collection of statistics does not interfere with the automatic collection of statistics. For example, assume that you have **wlm_collect_int** set to 60. Statistics are sent to the statistics event monitor every hour. Now assume that the last time the statistics were collected was 5:30 AM. You can invoke the **WLM_COLLECT_STATS** procedure at 5:55 AM, which sends the in-memory values of the statistics to the event monitor and resets the statistics. The next automatic statistics collection still occurs at 6:30 AM, one hour after the last automated collection. The collection interval is not affected by any manual collection and resetting of statistics that occurs during the interval.

Notes:

- The DB2 workload manager statistics table functions report the current values of the in-memory statistics. If you have automatic workload management statistics collection enabled, these values are reset periodically on the interval defined by the **wlm_collect_int** database configuration parameter. When looking at the statistics reported by the table functions, you should always consider the **LAST_RESET** column. This column indicates the last time the in-memory statistics were reset. If the time interval between the last reset time to the current time is not sufficiently large, there may not be enough data to draw any meaningful conclusions.
- If you are using automatic collection of workload management statistics, you need to prune your event monitor files or tables periodically. The event monitor does not automatically prune the data that is collected, and the automatic collection will fill your files or tables over time.
- When a database is deactivated, the in-memory statistics are reset. Deactivating the database does not send statistics to the statistics event monitor. If you do not want to lose the statistics accumulated since the last collection because of a deactivation, you should manually invoke the **WLM_COLLECT_STATS** procedure before deactivating the database.
- The **WLM_COLLECT_STATS** procedure resets statistics differently than the **RESET MONITOR** command. The **RESET MONITOR** command resets the values of snapshot monitor elements by storing their present values. After the **RESET MONITOR** command has been issued, snapshot processing reports the delta between these values and the current values. In contrast, the reset caused by the **WLM_COLLECT_STATS** procedure does not store any values, but instead resets all of the statistics counters themselves for each applicable DB2 workload manager object.

Also, with the **RESET MONITOR** command, each process (attachment) has its own private view of the monitor data. If one user performs a reset, other users are unaffected. By contrast, a reset of the workload manager statistics applies to all users.

Resetting statistics on DB2 workload manager objects

This topic describes how to reset statistics for DB2 workload manager objects.

Note that resetting statistics applies only to DB2 workload manager statistics; metrics reported by monitoring interfaces will be collected, but not reset.

Four events will reset the in-memory statistics stored for each DB2 workload manager object. (For a description of the statistics maintained for each object, see “Statistics for DB2 workload manager objects” on page 169.)

- The WLM_COLLECT_STATS stored procedure is invoked. See “Collecting workload management statistics using a statistics event monitor” on page 188 for details.
- The automatic DB2 workload manager statistics collection and reset process controlled by the **wlm_collect_int** database configuration parameter causes a collection and reset. See “Collecting workload management statistics using a statistics event monitor” on page 188 for details.
- The database is reactivated. Every time the database is activated on a database partition, the statistics for all DB2 workload manager objects on that database partition are reset.
- The object for which the statistics are maintained is modified and the change is committed. For example if a service subclass is altered, when the ALTER statement is committed, the in-memory statistics for that service subclass are reset.

You can determine the last time the statistics were reset for a given DB2 workload manager object using the statistics table functions and looking at timestamp in the LAST_RESET column. For example, to see the last time the statistics were reset for the service subclass SYSDEFAULTSUBCLASS under the SYSDEFAULTUSERCLASS service superclass, you could issue a query such as:

```
SELECT LAST_RESET
FROM TABLE(WLM_GET_SERVICE_SUBCLASS_STATS_V97( 'SYSDEFAULTUSERCLASS',
'SYSDEFAULTSUBCLASS', -2)) AS T
```

All statistics table functions return the statistics that accumulated since the last time that the statistics were reset. A statistics reset occurs when a database is activated or reactivated, when you alter a DB2 workload manager object (only the statistics for that object are reset), and when you call the WLM_COLLECT_STATS stored procedure. Statistics are also reset automatically according to the time period defined by the **wlm_collect_int** database configuration parameter, if you set this parameter to a nonzero value.

The period of time specified by **wlm_collect_int** is unaffected by a statistics reset that occurs during the interval specified by the configuration parameter. For example, if you run the WLM_COLLECT_STATS table function 5 minutes after the start of a 20-minute interval specified by **wlm_collect_int**, the interval still expires 15 minutes later. The statistics collection and reset that occur do not delay the occurrence of the next statistics collection and reset by 5 minutes.

If you change a service class or a workload to use a different histogram template or change a histogram template, the change does not take effect until a statistics reset occurs.

If you invoke the WLM_COLLECT_STATS table function to collect and reset statistics at the same time that another collection and reset is in progress (for example, if the invocation of the table function overlaps with the periodic collection and reset interval caused by **wlm_collect_int** or if another user invokes WLM_COLLECT_STATS at the same time), the collection and reset request from WLM_COLLECT_STATS is ignored, and warning SQL1632W is returned.

Monitoring metrics for DB2 workload manager

Monitoring metrics provide data about the health of and query performance on your DB2 data server, which can then be used as input to a 3rd party tool or in conjunction with additional scripting you provide to analyze the metrics returned.

Metrics are maintained for a number of DB2 database objects. These metrics reside in memory and can be viewed in real-time using DB2 monitoring metrics table functions, or the metrics can be collected and sent to an event monitor where they can be viewed later for historical analysis.

Monitoring metrics for activities

You can obtain monitoring metrics for activities using:

- The activities event monitor (DETAILS_XML column)
- The MON_GET_ACTIVITY_DETAILS table function

Monitoring metrics for activities are controlled by the **mon_act_metrics** database configuration parameter and the COLLECT ACTIVITY METRICS clause on a workload. Metrics will be collected for an activity, if the database configuration parameter is set to a value other than NONE or if the activity is submitted by a connection that is associated with a workload which has a COLLECT ACTIVITY METRICS setting other than NONE.

You can use workload-level controls to achieve better monitoring granularity, if you do not want to collect metrics for all activities. If activity metrics collection is enabled at the database level (enabled by default), then metrics are collected for all activities, regardless of the setting at the workload level.

See the monitoring documentation for more details.

System-level monitoring metrics

You can obtain system-level monitoring metrics aggregated by service classes and workloads using:

- The statistics event monitor (DETAILS_XML column in the wlstats and scstats logical groups)
- The MON_GET_SERVICE_SUBCLASS, MON_GET_SERVICE_SUBCLASS_DETAILS, MON_GET_WORKLOAD and MON_GET_WORKLOAD_DETAILS table functions

Monitoring metrics for requests to the data server, including those requests that are part of an activity, are controlled by the **mon_req_metrics** database configuration parameter and the COLLECT REQUEST METRICS clause on a service superclass. Metrics will be collected for a request, if the database configuration parameter is set to a value other than NONE or if the request is submitted by a connection that mapped to a subclass under a superclass which has a COLLECT REQUEST METRICS setting other than NONE.

You can use service superclass-level controls to achieve better monitoring granularity, if you do not want to collect metrics for all requests. If request metrics collection is enabled at the database level (enabled by default), then metrics are collected for all requests, regardless of the setting at the service superclass level.

See the monitoring documentation for more details.

Workload management table functions and snapshot monitor integration

You can use DB2 workload manager table functions with the snapshot monitor table functions when performing problem determination or performance tuning.

The DB2 workload manager table functions and the snapshot monitor table functions share the following fields. You can perform joins on these fields to derive data that you need to perform diagnostic and performance-tuning activities. Note that, unlike the snapshot table functions, the WLM table functions do not get their information from the snapshot monitor, so that the information available in the WLM table functions is not available from the snapshot monitor.

Table 55. Fields shared between the DB2 workload manager and snapshot monitor table functions

Workload manager table function field	Snapshot monitor table function field
agent_tid	agent_pid
application_handle	agent_id agent_id_holding_lock
session_auth_id	session_auth_id
dbpartitionnum	node_number
utility_id	utility_id
workload_id	workload_id

As an example of a reason to use a join between different table functions, assume that you want to obtain basic information about all of the utilities running in the BATCH service superclass. You might issue the following query:

```
SELECT SUBSTR(UTILITY_TYPE,1,4) TYPE,
       UTILITY_PRIORITY PRIORITY,
       SUBSTR(UTILITY_DESCRIPTION,1,12) AS UTILITY_DESCRIPTION,
       SUBSTR(UTILITY_DBNAME,1,8) AS DBNAME,
       UTILITY_STATE,
       SUBSTR(UTILITY_INVOKER_TYPE,1,7) INVOKER,
       SUBSTR(CHAR(WLM.DBPARTITIONNUM),1,4) PART,
       SUBSTR(CLASSES.PARENTSERVICECLASSNAME,1,19) SUPERCLASS_NAME,
       SUBSTR(CLASSES.SERVICECLASSNAME,1,18) SUBCLASS_NAME
FROM SYSIBMADM.SNAPUTIL SNAP,
     TABLE(WLM.GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97(CAST(NULL AS BIGINT), -2)) WLM,
     SYSCAT.SERVICECLASSES CLASSES
WHERE SNAP.UTILITY_ID = WLM.UTILITY_ID
      AND WLM.SERVICE_CLASS_ID = CLASSES.SERVICECLASSID
      AND CLASSES.SERVICECLASSNAME = 'SYSDEFAULTSUBCLASS'
      AND CLASSES.PARENTSERVICECLASSNAME = 'BATCH'
ORDER BY WLM.DBPARTITIONNUM
```

The output might resemble the following output:

TYPE	PRIORITY	UTILITY_DESCRIPTION	DBNAME	UTILITY_STATE	INVOKER	PART	SUPERCLASS_NAME	SUBCLASS_NAME
LOAD	- OFFLINE	LOAD	SAMPLE	EXECUTE	USER	1	BATCH	SYSDEFAULTSUBCLASS
LOAD	- OFFLINE	LOAD	SAMPLE	EXECUTE	USER	1	BATCH	SYSDEFAULTSUBCLASS
LOAD	- OFFLINE	LOAD	SAMPLE	EXECUTE	USER	1	BATCH	SYSDEFAULTSUBCLASS
LOAD	- OFFLINE	LOAD	SAMPLE	EXECUTE	USER	2	BATCH	SYSDEFAULTSUBCLASS
LOAD	- OFFLINE	LOAD	SAMPLE	EXECUTE	USER	2	BATCH	SYSDEFAULTSUBCLASS
LOAD	- OFFLINE	LOAD	SAMPLE	EXECUTE	USER	2	BATCH	SYSDEFAULTSUBCLASS
LOAD	- OFFLINE	LOAD	SAMPLE	EXECUTE	USER	3	BATCH	SYSDEFAULTSUBCLASS
LOAD	- OFFLINE	LOAD	SAMPLE	EXECUTE	USER	3	BATCH	SYSDEFAULTSUBCLASS
LOAD	- OFFLINE	LOAD	SAMPLE	EXECUTE	USER	3	BATCH	SYSDEFAULTSUBCLASS

Monitoring threshold violations

When a DB2 workload manager threshold is violated, a threshold violation record is written to the active THRESHOLD VIOLATIONS event monitor, if one exists.

About this task

The threshold violation record includes the following information:

- A description of the threshold that was violated (the identifier, maximum value, and so on).
- An identification of the activity that violated the threshold, including the identifier of the application that submitted the activity, the unique activity identifier, and the unit of work identifier.
- The time that the threshold was violated.
- The action that was taken. The action indicates whether the activity that violated the threshold was permitted to continue or was stopped. If the activity was stopped, the application that submitted the activity will have received an SQL4712N error.

When a threshold violation occurs for a threshold that has a REMAP ACTIVITY action defined for it, a threshold violation record is optional. Whether or not a threshold violation record is recorded is determined by the NO EVENT MONITOR RECORD or LOG EVENT MONITOR RECORD clause of your CREATE THRESHOLD statement.

You can optionally have detailed activity information (including statement text) written to an active activities event monitor if the threshold violation is caused by an activity. The activity information is written when the activity completes, not when the threshold is violated. Specify that activity information should be collected when a threshold is violated by using the COLLECT ACTIVITY DATA keyword on either the CREATE or ALTER threshold or work action set statements.

Procedure

To monitor threshold violations:

1. Use the CREATE EVENT MONITOR statement to create an event monitor of type THRESHOLD VIOLATIONS. For example:

```
CREATE EVENT MONITOR VIOLATIONS FOR THRESHOLD VIOLATIONS WRITE TO TABLE
```
2. Use the COMMIT statement to commit your changes.
3. Use the SET EVENT MONITOR STATE statement to activate the event monitor. Instead of using the SET EVENT MONITOR STATE statement, you can use the AUTOSTART default for the THRESHOLD VIOLATIONS event monitor to have it activated the next time that the database is activated. However, only one event monitor of the THRESHOLD VIOLATIONS type can be active on a database partition at one time. If you want to define multiple THRESHOLD VIOLATIONS event monitors, you should not use the AUTOSTART option.
4. Use the COMMIT statement to commit your changes.

Note: If you create any thresholds, you should create and activate a threshold violations event monitor so you can monitor any threshold violations that occur. A threshold violations event monitor does not have any impact unless thresholds are violated.

Example

This example shows how you can determine what remappings of a particular activity occurred as the result of a threshold violation that included a REMAP ACTIVITY action. To find the activities that were remapped, use a statement like the following:

```
SELECT VARCHAR(APPL_ID, 30) AS APPLID,
       UOW_ID,
       ACTIVITY_ID,
       VARCHAR(T.PARENTSERVICECLASSNAME,20) AS SERVICE_SUPERCLASS,
       VARCHAR(T.SERVICECLASSNAME,20) AS FROM_SERVICE_SUBCLASS,
       VARCHAR(S.SERVICECLASSNAME,20) AS TO_SERVICE_SUBCLASS
FROM THRESHOLDVIOLATIONS_TH1,
     SYSCAT.SERVICECLASSES AS T,
     SYSCAT.SERVICECLASSES AS S
WHERE SOURCE_SERVICE_CLASS_ID = T.SERVICECLASSID AND
       DESTINATION_SERVICE_CLASS_ID = S.SERVICECLASSID AND
       THRESHOLD_ACTION = 'REMAP'
ORDER BY APPLID, ACTIVITY_ID, UOW_ID, TIME_OF_VIOLATION ASC;
```

In this example, two remappings occurred for the activity submitted by the application with the ID *N0.swalkty.080613140844 which is identified by activity ID 1 and unit of work (UOW) ID 1:

APPLID	UOW_ID	ACTIVITY_ID	SERVICE_SUPERCLASS	FROM_SERVICE_SUBCLASS	TO_SERVICE_SUBCLASS
*N0.swalkty.080613140844	1	1	1 WORK	HIGH	MED
*N0.swalkty.080613140844	1	1	1 WORK	MED	LOW

2 record(s) selected.

The output is ordered by the time of threshold violation and shows that the activity was remapped twice after it started executing. Although not shown in the output, the initial service subclass the activity was mapped to is likely a high priority service subclass, typical of a three-tiered configuration that permits shorter running queries to complete more quickly. Because the activity did not complete quickly enough in the high priority service subclass, it violated a threshold and was remapped to a medium priority service subclass, and then remapped again to a low priority service subclass after a second threshold violation later on.

Collecting data for individual activities

You can use an ACTIVITIES event monitor to collect data for individual activities that run in your system. The data collected includes items such as statement text and compilation environment, and can be used to investigate and diagnose problems, and as input to other tools (for example, the Design Advisor).

You can collect information about individual activities for service subclasses, workloads, work classes (through work actions), and threshold violations. You enable activity collection using the COLLECT ACTIVITY DATA keyword of the CREATE and ALTER statements for these DB2 workload manager objects. When an activity completes, information about the activity is sent to the active ACTIVITIES event monitor if:

- The activity was submitted by an application that is mapped to a workload for which COLLECT ACTIVITY DATA is specified, or
- The activity runs in a service subclass for which COLLECT ACTIVITY DATA is specified, or
- The activity has a COLLECT ACTIVITY DATA work action applied to it, or
- The activity violates a threshold that was defined with the COLLECT ACTIVITY DATA action

The COLLECT ACTIVITY DATA keyword also controls the amount of information that is sent to the ACTIVITIES event monitor. If the keyword specifies WITH DETAILS, statement information (such as statement text) is collected. If the keyword specifies WITH DETAILS AND VALUES, data values are collected as well.

An activity might have multiple COLLECT ACTIVITY DATA keywords applied to it. For example, the activity might run in a service subclass for which COLLECT ACTIVITY DATA is specified, and while executing it might violate a threshold that has the COLLECT ACTIVITY DATA action. In this situation, the activity is only collected once. The COLLECT keyword that specifies the largest amount of information to be collected is applied to the activity. For example, if both COLLECT ACTIVITY DATA WITHOUT DETAILS and COLLECT ACTIVITY DATA WITH DETAILS are applied to an activity, the activity is collected with detailed information.

To enable collection of activities for a given DB2 workload manager object:

1. Use the CREATE EVENT MONITOR statement to create an ACTIVITIES event monitor.
2. Use the COMMIT statement to commit your changes.
3. Use the SET EVENT MONITOR STATE statement to activate the event monitor. Instead of using the SET EVENT MONITOR STATE statement, you can use the AUTOSTART default for the ACTIVITIES event monitor to have it activated the next time that the database is activated. If you want to define multiple ACTIVITIES event monitors, you should not use the AUTOSTART option.
4. Use the COMMIT statement to commit your changes.
5. Identify the objects for which you want to collect activities by using the ALTER SERVICE CLASS, ALTER WORK ACTION SET, ALTER THRESHOLD, or ALTER WORKLOAD statement and specify the COLLECT ACTIVITY DATA keywords.
6. Use the COMMIT statement to commit your changes.

Note: Individual activity collection is more expensive than workload management statistics collection. You should try to set up activity collection to collect as few activities as possible. For example, if you need to investigate activities submitted by a specific application, you could isolate that application by creating a workload or service class specifically for that application, and only enable activity collection for that workload or service class.

You might not always know in advance that you will want to capture an activity. For example, you might have a query that is taking a long time to run and you want to collect information about it for later analysis. In this situation, it is too late to specify the COLLECT ACTIVITY DATA keyword on the DB2 workload manager objects, because the activity has already entered the system. In this situation, you can use the WLM_CAPTURE_ACTIVITY_IN_PROGRESS stored procedure. The WLM_CAPTURE_ACTIVITY_IN_PROGRESS stored procedure sends information about an executing activity to the active ACTIVITIES event monitor. You identify the activity to be collected using the application handle, unit of work identifier, and activity identifier. Information about the activity is immediately be sent to the ACTIVITIES event monitor when the procedure is invoked: you do not need to wait for the activity to complete.

Importing activity information into the Design Advisor

You can import activities collected by an activities event monitor into the Design Advisor to help you make decisions about the database objects accessed by these activities.

Activities imported into the design advisor must have been collected using the COLLECT ACTIVITY DATA WITH DETAILS or COLLECT ACTIVITY DATA WITH DETAILS AND VALUES options. The COLLECT ACTIVITY DATA WITHOUT DETAILS option is not sufficient, it will not capture the statement text which is required by the Design Advisor.

To import activity information from the activity event monitor tables into the Design Advisor, run the db2advis command with the **-wlm** parameter, followed by additional parameters:

1. The activities event monitor name
2. Optional: the workload or service class name
3. Optional: the start time and end time

For example, to import information about all the activities collected by the DB2ACTIVITIES event monitor in the SAMPLE database, use the following command:

```
db2advis -d SAMPLE -wlm DB2ACTIVITIES
```

Note: You can only import information from activities event monitor tables through the Design Advisor command line interface.

Cancelling activities

If an activity is consuming too many resources, or is running too long, you can cancel it. Cancelling an activity is gentler than forcing the application that submitted the activity. A cancelled activity returns SQL4725N to the user, but does not end the connection or affect any other user activity. Forcing the application ends both the connection and user activities.

You can only explicitly cancel an activity if a coordinator activity is currently working on a request for the activity. If you cancel an activity in the IDLE state (that is, no requests are being processed), the activity is placed in the CANCEL_PENDING state and is cancelled on the next request that is received. For example, if you attempt to cancel a CURSOR activity between fetches, the SQL4725N error is not returned to the user until the next fetch after the cancel.

All user activities are cancellable, including the load utility and stored procedures.

To cancel an activity:

1. Identify the activity that you want to cancel. You can use the WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97 table function to identify the activities running in an application. You can also use the MON_GET_ACTIVITY_DETAILS_COMPLETE table function to view additional details about a particular activity if the information in WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97 is not sufficient to identify the work that the activities are performing.
2. Cancel the activity using the WLM_CANCEL_ACTIVITY stored procedure. The stored procedure takes the following arguments: *application_handle*, *uow_id*, and

activity_id. For an example of how to use this stored procedure, see “Scenario: Identifying activities that are taking too long to complete” on page 267.

Guidelines for capturing information about and investigating a rogue activity

This topic provides guidelines for capturing information about, and investigating, a rogue activity.

First establish a set of criteria for what you would consider a rogue activity. For example:

- An activity in that runs in a service class for activities with a low estimated cost, and runs for more than 1 hour
- An activity that returns an unusually large number of rows
- An activity that consumes an unusually high amount of temporary table space

Then create thresholds that describe these criteria and contain a COLLECT ACTIVITY DATA WITH DETAILS action. When the threshold is violated, information about the activity that violated the threshold is sent to the active ACTIVITIES event monitor when the activity completes.

For example, to collect information about any database activity that runs for more than 3 hours, create a threshold such as the following threshold:

```
CREATE THRESHOLD LONGRUNNINGACTIVITIES
  FOR DATABASE ACTIVITIES ENFORCEMENT DATABASE
  WHEN ACTIVITYTOTALTIME > 3 HOURS COLLECT ACTIVITY DATA WITH DETAILS
  CONTINUE
```

Monitoring with DB2 workload manager is lightweight, if you are careful to apply it only to a small subset of your queries, as shown in the example, where only queries running for at least three hours are monitored. You can refine this example further by creating a threshold not at the global database level, but at the level of a user-defined superclass. If this more narrowly scoped monitoring suits your purpose, it can further reduce the cost of monitoring and it will provide information only at the level you need:

```
CREATE SERVICE CLASS LONGQUERIES
  AGENT PRIORITY 20
  PREFETCH PRIORITY LOW

CREATE THRESHOLD LONGRUNNINGACTIVITIES2
  FOR SERVICE CLASS LONGQUERIES ACTIVITIES ENFORCEMENT DATABASE
  WHEN ACTIVITYTOTALTIME > 3 HOURS COLLECT ACTIVITY DATA WITH DETAILS
  CONTINUE
```

. The service class created for the threshold is assigned low agent and prefetch priority, because it is intended to be used for long running queries (this SQL statement works on UNIX operating systems and Linux; on Windows operating systems, substitute an agent priority of -6).

After your data server has performed some work, you can analyze the information that is written to the threshold violations and activities event monitors. DML activities also have their statement text and compilation environment information written to the activities event monitor, so you can run DB2 explain on them to further investigate the performance of the activity.

Workload management performance modelling

The workload on your system can be modelled as a set of activities that arrive at the system at a rate governed by an arrival rate distribution for activities (often measured as its inverse, the *inter-arrival* time distribution) and the amount of time activities spend executing in the system following a service time distribution.

Inter-arrival time is the time between the arrival of one activity and the arrival of the next activity. Service time is the time that an activity spends executing on the system. For example, if you submit a query at time 0 seconds, it spends 2 seconds in a queue, and it finishes at time 5 seconds, the service time is $5 - 2 = 3$ seconds. Service time assumes no other work executing on the system (that is, it is not the observed execution time, but rather the time it would take to execute the activity in isolation). The service time distribution can be approximated for DML activities using the estimated cost in timerons, which considers both processor and I/O time for an activity.

You can build a workload model for your system by measuring the inter-arrival time distribution and the service time distribution of the activities on the system. Inter-arrival time distributions and approximate service time distributions (using estimated cost) can be obtained by using extended aggregate activity statistics for service subclasses or work classes (using work actions) and a statistics event monitor. These statistics are not collected by default. See “Statistics for DB2 workload manager objects” on page 169 for more information.

Example: Capturing information about an activity for later analysis

You can use workload management features to capture information about an activity for later analysis.

Assume that you have a stored procedure called `MYSHEMA.MYSLOWSTP` and that it is running more slowly than usual. You begin to receive complaints about this situation and decide to investigate the cause of the slowdown. If investigating while the stored procedure is running is impractical, you can capture information about the stored procedure activity and any activities nested in it.

Assuming that you have an active activities event monitor called `DB2ACTIVITIES`, you can create a work class for `CALL` statements that apply to the schema of the `MYSHEMA.MYSLOWSTP` stored procedure. Then you can create a work action to map the `CALL` activity and all nested activities to a service class that has activity collection enabled. The `CALL` activity, and any activities nested in it, are sent to the event monitor. Following are examples of the DDL required to create the DB2 workload manager objects:

```
CREATE SERVICE CLASS SC1;
CREATE WORKLOAD WL1 APPLNAME ('DB2BP') SERVICE CLASS SC1;
CREATE SERVICE CLASS PROBLEMQUERIESSC UNDER SC1 COLLECT ACTIVITY DATA ON COORDINATOR WITH DETAILS;

CREATE WORK CLASS SET PROBLEMQUERIES
(WORK CLASS CALLSTATEMENTS WORK TYPE CALL ROUTINES IN SCHEMA MYSCHEMA);

CREATE WORK ACTION SET DATABASEACTIONS FOR SERVICE CLASS SC1 USING WORK CLASS SET PROBLEMQUERIES
(WORK ACTION CAPTURECALL ON WORK CLASS CALLSTATEMENTS MAP ACTIVITY WITH NESTED TO PROBLEMQUERIESSC);
```

After the `MYSHEMA.MYSLOWSTP` stored procedure runs, you can issue the following query to obtain the application handle, the unit of work identifier, and the activity identifier for the activity:

```
SELECT AGENT_ID,
       UOW_ID,
       ACTIVITY_ID
```

```

FROM ACTIVITY_DB2ACTIVITIES
WHERE SC_WORK_ACTION_SET_ID = (SELECT ACTIONSETID
                               FROM SYSCAT.WORKACTIONSETS
                               WHERE ACTIONSETNAME = 'DATABASEACTIONS')
AND SC_WORK_CLASS_ID = (SELECT WORKCLASSID
                        FROM SYSCAT.WORKCLASSES
                        WHERE WORKCLASSNAME = 'CALLSTATEMENTS'
                        AND WORKCLASSETID =
                        (SELECT WORKCLASSETID FROM SYSCAT.WORKACTIONSETS WHERE ACTIONSETNAME
                        = 'DATABASEACTIONS'));

```

Assuming that the captured activity has an application handle of 1, a unit of work identifier of 2, and an activity identifier of 3, the following results are generated:

```

AGENT_ID          UOW_ID          ACTIVITY_ID
=====          =====          =====
                  1                2                3

```

Using this information, you can issue the following query against the ACTIVITY_DB2ACTIVITIES and the ACTIVITYSTMT_DB2ACTIVITIES tables to determine where the activity spent its time:

```

WITH RAH (LEVEL, APPL_ID, PARENT_UOW_ID, PARENT_ACTIVITY_ID,
          UOW_ID, ACTIVITY_ID, STMT_TEXT, TIME_CREATED, TIME_COMPLETED) AS
  (SELECT 1, ROOT.APPL_ID, ROOT.PARENT_UOW_ID,
         ROOT.PARENT_ACTIVITY_ID, ROOT.UOW_ID, ROOT.ACTIVITY_ID,
         ROOTSTMT.STMT_TEXT, ROOT.TIME_CREATED, ROOT.TIME_COMPLETED
   FROM ACTIVITY_DB2ACTIVITIES ROOT, ACTIVITYSTMT_DB2ACTIVITIES ROOTSTMT
   WHERE ROOT.APPL_ID = ROOTSTMT.APPL_ID AND ROOT.AGENT_ID = 1
        AND ROOT.UOW_ID = ROOTSTMT.UOW_ID AND ROOT.UOW_ID = 2
        AND ROOT.ACTIVITY_ID = ROOTSTMT.ACTIVITY_ID AND ROOT.ACTIVITY_ID = 3
   UNION ALL
   SELECT PARENT.LEVEL +1, CHILD.APPL_ID, CHILD.PARENT_UOW_ID,
          CHILD.PARENT_ACTIVITY_ID, CHILD.UOW_ID,
          CHILD.ACTIVITY_ID, CHILDSTMT.STMT_TEXT, CHILD.TIME_CREATED,
          CHILD.TIME_COMPLETED
   FROM RAH PARENT, ACTIVITY_DB2ACTIVITIES CHILD,
        ACTIVITYSTMT_DB2ACTIVITIES CHILDSTMT
   WHERE PARENT.APPL_ID = CHILD.APPL_ID AND
        CHILD.APPL_ID = CHILDSTMT.APPL_ID AND
        PARENT.UOW_ID = CHILD.PARENT_UOW_ID AND
        CHILD.UOW_ID = CHILDSTMT.UOW_ID AND
        PARENT.ACTIVITY_ID = CHILD.PARENT_ACTIVITY_ID AND
        CHILD.ACTIVITY_ID = CHILDSTMT.ACTIVITY_ID AND
        PARENT.LEVEL < 64
  )
SELECT UOW_ID, ACTIVITY_ID, SUBSTR(STMT_TEXT,1,40),
       TIMESTAMPDIFF(2, CHAR(TIME_COMPLETED - TIME_CREATED)) AS
       LIFE_TIME
FROM RAH
ORDER BY UOW_ID, ACTIVITY_ID;

```

The results would resemble the following ones:

```

UOW_ID ACTIVITY_ID STMT_TEXT                                     LIFE_TIME
=====
2       3          CALL SLOWPROC                                             1000
2       4          SELECT COUNT(*) FROM ORG                                1
2       5          SELECT * FROM MYHUGETABLE                               999

```

The results indicate that the stored procedure is spending most of its time querying the MYHUGETABLE table. Your next step is to investigate what changes to the MYHUGETABLE table might cause queries running against it to slow down.

When many stored procedures run simultaneously, greater overhead is incurred when performing the analysis. To solve this problem, you can create a workload and service class for running a stored procedure that is issued by a specific

authorization identifier, a specific application, or both. You can then use the preceding method to analyze the behavior of the stored procedure.

Chapter 5. Integration with operating system workload managers

If available, use DB2 workload manager in conjunction with an operating system workload manager, which provides you with additional capabilities.

The point of integration between DB2 workload manager and operating system workload managers is the DB2 service class. You create a mapping between a DB2 service class and an operating system workload manager class when you define a DB2 service class by using the `OUTBOUND CORRELATOR` option of the `CREATE SERVICE CLASS` or the `ALTER SERVICE CLASS` statement.

If the outbound correlator is set, all threads in the DB2 service class are associated with the operating system workload manager using the outbound correlator when the next activity begins.

Integration of AIX Workload Manager with DB2 workload manager

On the AIX operating system, the optional integration between DB2 service classes and AIX WLM classes permits you to control the amount of processor resource allocated to each service class.

Implementing AIX WLM controls may not be needed to meet your performance objectives, but even if you do not need to exercise AIX WLM, the operating system statistics provided by AIX WLM per AIX class are often useful for monitoring and tuning efforts.

AIX WLM assigns relative or absolute amounts of processor resource as shares to classes which benefit from controls that you can change dynamically and that become effective immediately. If relative AIX CPU shares do not provide the level of control you require, you also have the choice of assigning hard maximum percentage of CPU resource. By doing so, you surrender some of the flexibility of relative CPU allocation, which is useful during off-peak times, but you also gain excellent and guaranteed control with a hard maximum limit on CPU time resource allocation.

Recommended mappings between DB2 service classes and AIX classes

Use a 1:1 mapping of DB2 service classes to AIX Workload Manager service classes to take advantage of AIX WLM processor controls. By having a 1:1 mapping between DB2 service classes and AIX Workload Manager service classes, you can adjust the AIX processor resource for each DB2 service class individually to meet your business priority goals.

The following figure shows the integration of the DB2 workload manager with the AIX Workload Manager. Note the 1:1 mapping between each DB2 service class and AIX Workload Manager service class at the service superclass and service subclass levels.

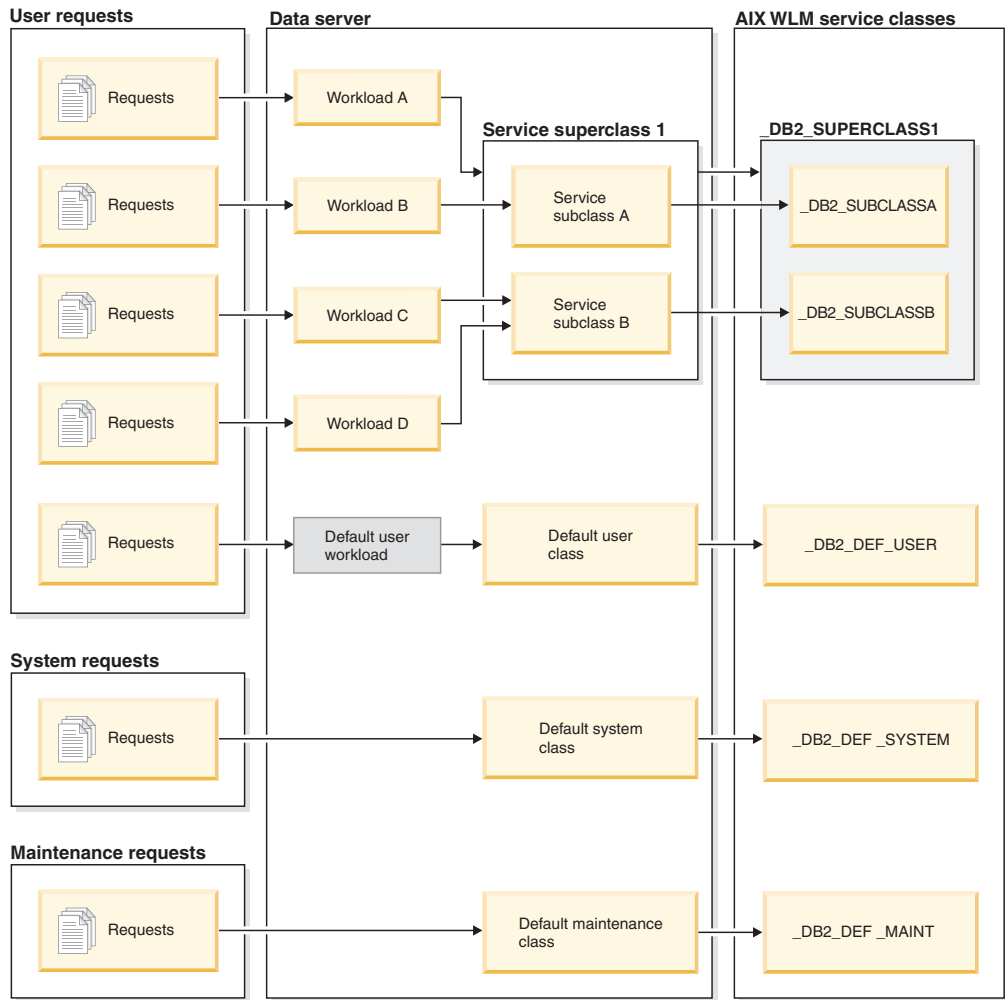


Figure 28. Integration of the DB2 workload manager with the AIX Workload Manager

When a DB2 environment consists of a single database in a single DB2 instance, such as the example portrayed in the previous figure, it is possible to map directly between DB2 service classes and AIX Workload Manager classes. Each DB2 service superclass can have a corresponding AIX Workload Manager service superclass and each DB2 service subclass can map to a corresponding AIX service subclass.

In situations where the DB2 environment consists of multiple databases and DB2 instances, several levels might be candidates for resource control. Because the AIX Workload Manager supports a two-level hierarchy, that is, superclass and subclass, only two levels of a DB2 environment can be mapped to AIX Workload Manager classes at any time. The following figure shows one way to achieve a 1:1 mapping with multiple databases, each with multiple superclasses. Here, each database has its own AIX Workload Manager superclass and each DB2 service superclass is mapped to an AIX Workload Manager subclass.




Figure 29. DB2 service classes mapped to AIX classes (with DB2 service superclasses only)

An alternative configuration is to map each DB2 service superclass to its own AIX Workload Manager superclass, which results in four superclasses in this example. In this situation, the database level of resource control is represented explicitly in the AIX Workload Manager service class definitions.

The following figure shows one way to achieve the 1:1 mapping in the situation where you have multiple databases, each with service superclasses and service subclasses. Here, each database corresponds to an AIX superclass and each DB2 service subclass is mapped to an AIX Workload Manager subclass. The DB2 service superclass is not shown explicitly in the AIX Workload Manager service class

definitions.

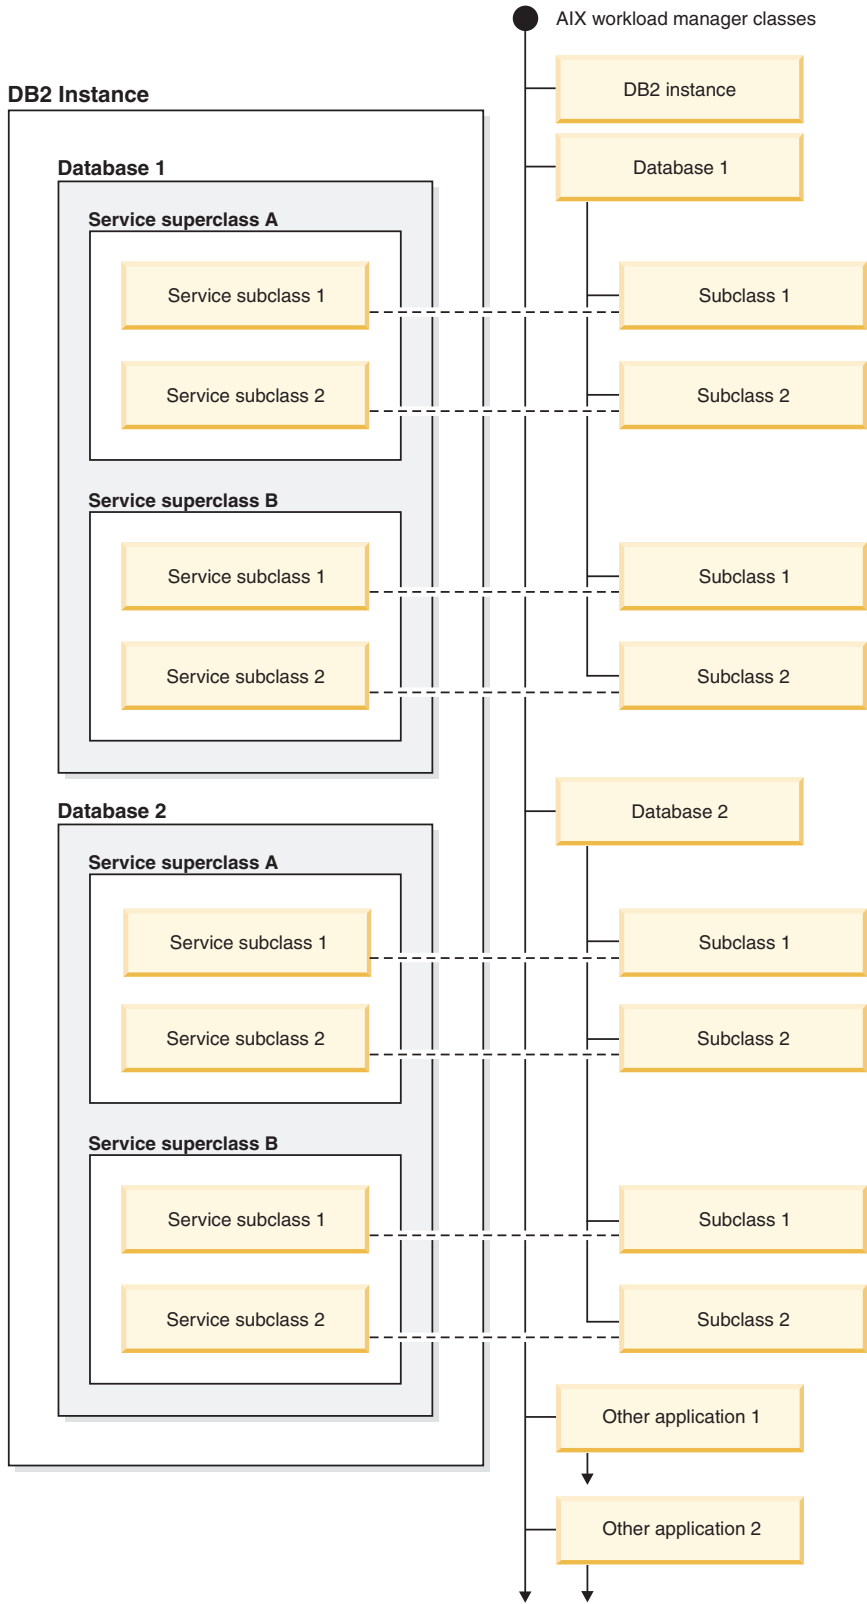


Figure 30. DB2 service classes mapped to AIX Workload Manager classes (with DB2 service subclasses)

Defining mappings between DB2 service classes and AIX classes

Mapping between DB2 service classes and AIX Workload Manager classes is specified for the DB2 service class using the `OUTBOUND CORRELATOR` keyword of the `CREATE SERVICE CLASS` or the `ALTER SERVICE CLASS` statements.

The steps for setting up the AIX Workload Manager classes with the DB2 data server are:

1. Create the DB2 service superclasses and service subclasses, and specify the `OUTBOUND CORRELATOR` tags.
2. Create the corresponding AIX classes.
3. Create the associated AIX Workload Manager rules files to contain the DB2 workload manager to AIX Workload Manager mappings using the `OUTBOUND CORRELATOR` tags under the tag columns.
4. Start the AIX Workload Manager.
5. If required, set this AIX Workload Manager configuration as active.

When a thread joins a DB2 service class, the DB2 data server calls the appropriate AIX Workload Manager API to associate the thread to the corresponding AIX service class. The DB2 data server sends the thread's target AIX service class to the AIX Workload Manager by passing it the application tag set in the `OUTBOUND CORRELATOR` parameter.

You must ensure that the AIX Workload Manager is properly installed, configured, and active. If the DB2 data server cannot communicate with the AIX Workload Manager, a message is logged to the `db2diag` log files and DB2 administrator log. The database activity continues.

The DB2 data server cannot detect whether the `OUTBOUND CORRELATOR` value that it passes to the AIX Workload Manager is recognized by the AIX Workload Manager. You must verify that the value specified for the DB2 service class matches the application tags that map DB2 threads to the AIX service classes. If the `OUTBOUND CORRELATOR` value is not recognized by the AIX Workload Manager, the database activity continues to execute.

Other points to note are:

- DB2 service classes cannot work with the AIX Workload Manager inheritance feature. Inheritance is the default setting for an AIX service class; inheritance must be explicitly disabled by setting the inheritance attribute to `N0`. AIX Workload Manager inheritance forces all child threads and processes to map to the same class as the parent thread or process. If inheritance is enabled, the DB2 workload manager cannot change the AIX Workload Manager class of a thread using tagging. This restriction makes any integration of the DB2 workload manager and the AIX Workload Manager unusable. The DB2 data server cannot detect whether AIX Workload Manager inheritance is enabled and does not issue an error message if inheritance is enabled.
- DB2 service classes are not compatible with the AIX Workload Manager manual assignment feature. With the manual assignment feature, users can manually assign a process to a specific AIX Workload Manager class. By manually assigning the DB2 process, all threads in the process are assigned to a target AIX Workload Manager class, the DB2 service class mapping logic is defeated and results are not predictable.

For more information on the AIX Workload Manager, see the AIX Information Center at <http://publib.boulder.ibm.com/infocenter/pseries/v5r3/index.jsp>

Setting processor controls on AIX classes

The AIX Workload Manager can be used to control the amount of processor resource allocated to each service class. Options include setting a minimum, maximum, or relative proportion share of processor resource for each service class.

When integrating the AIX Workload Manager with DB2 Workload Management, only processor resource allocation is supported. You should not set memory and I/O settings for the AIX classes. DB2 database-level memory is shared among all agents from different DB2 service classes, so you cannot divide memory allocation between different service classes. AIX-level I/O control does not support the DB2 engine threaded model. To control I/O, you can use the prefetcher priority attribute of a DB2 service class to differentiate I/O priorities between different DB2 service classes.

If you use AIX to control the amount of processor resource allocated to a service class, do not also change the agent priority setting for that DB2 service class. Use only one of these mechanisms to govern the access to processor resource. You cannot set both the AGENT PRIORITY and the OUTBOUND CORRELATOR value for a service class. See “Agent priority of service classes” on page 72 for more information.

AIX Workload Manager settings should be consistent on all physical computers that participate in an instance. For example, if the resource setting for an AIX service class is set high on one computer, the same setting should be used for that AIX service class on all other computers. If the resource usage settings are inconsistent across computers, requests running in the same AIX service class will exhibit different performance levels on different database partitions. This situation can lead to poor overall throughput for connections in an AIX service class.

Integration of Linux workload management with DB2 workload manager

On the Linux operating system, the optional integration between DB2 service classes and Linux classes (control groups) permits you to control the amount of processor resource allocated to each service class. If enabled, all threads running in a DB2 service class are mapped to a Linux class where they are subject to the processor resource controls you define.

To make use of Linux workload management support, you require a Linux kernel version 2.6.26 or later and the libcgrou library package.

Linux workload management supports a hierarchy of classes with superclasses and subclasses, with processor shares for subclasses divided in proportion to the shares of the parent class. These shares provide a method of control over processor resource such that all threads in the system will always run, but the amount of processor time each thread receives is dependent on the number of shares assigned to the Linux class.

Processor resource on the Linux operating system is assigned in shares relative to the Linux workload management default class, which by default has a processor share at a value of 1024. If you define no other Linux classes, all threads run in

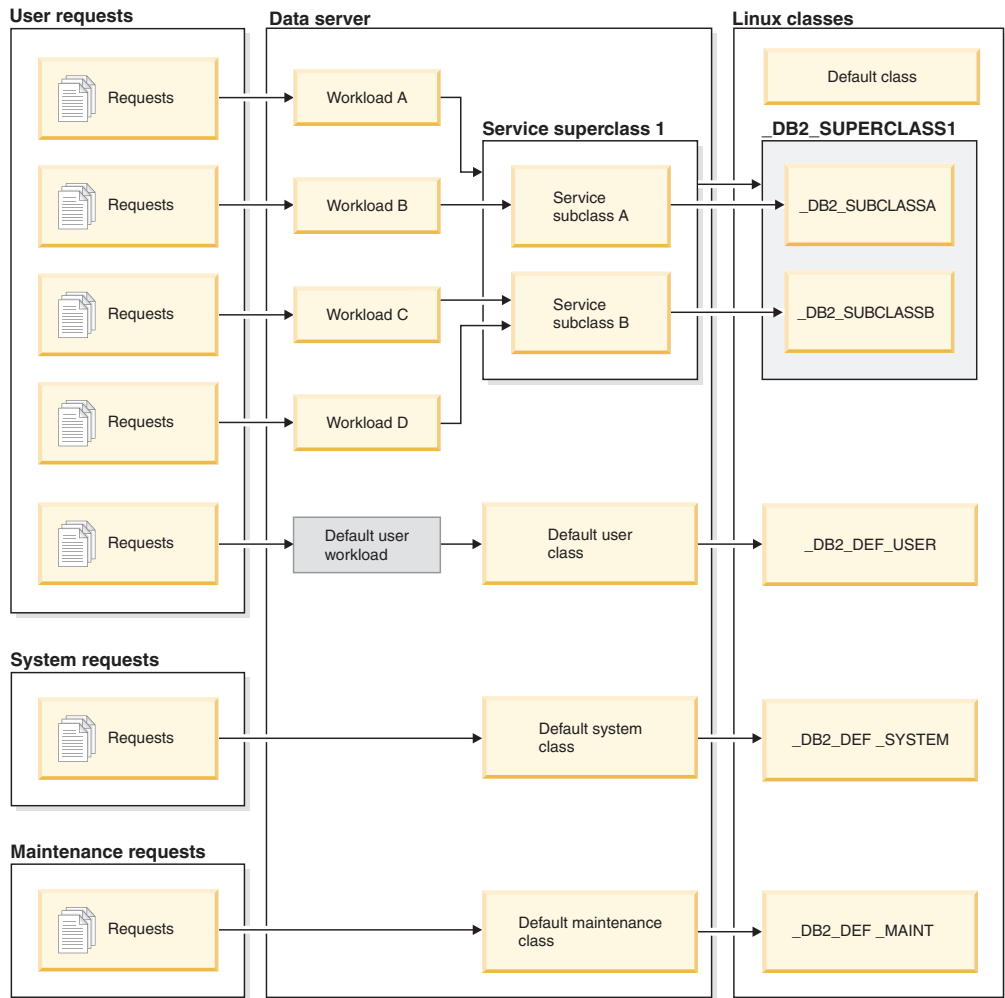
this default class. If you define a class that has a share value equal to 1024, then this class receives the same amount of processor resource as the Linux default class with the default processor share. Similarly, a class with a share of 2048 receives a target processor usage quota twice that of the default class. On more complex systems, you should consider raising the processor share of the Linux default class, which improves the granularity for shares across the system so that you can assign processor resources more accurately.

Recommended mappings between DB2 service classes and Linux classes

You should use a 1:1 mapping between DB2 service classes and Linux classes which permits you to adjust the Linux processor shares assigned to activities in each DB2 service class individually according to business priority. It is important that you associate every DB2 service class with a Linux WLM class, either by setting an outbound correlator for each service superclass and subclass, or through inheritance from the parent service class for subclasses. This includes the default SYSDEFAULTSYSTEMCLASS, SYSDEFAULTMAINTENANCECLASS and SYSDEFAULTUSERCLASS service classes.

The following figure shows how two DB2 service subclasses under the same user defined service superclass can get mapped 1:1 to Linux subclasses under a common superclass. In this example, the work identified and assigned by two workloads for each DB2 service subclass is subject to the processor resource controls imposed by the corresponding Linux subclasses (`_DB2_SUBCLASSA`, `_DB2_SUBCLASSB`). Also shown are three Linux classes that correspond to the default DB2 workload manager service classes (`_DB2_DEF_USER`, `_DB2_DEF_SYSTEM`, `_DB2_DEF_MAINT`). If you integrate DB2 workload manager with Linux workload management, you should always create these additional Linux classes to match the default DB2 service classes. To avoid any bottleneck, the Linux class corresponding to the DB2 default system class should receive more processor shares than any other Linux class that DB2 activities map to, whilst the Linux class corresponding to the default maintenance class should receive less processor shares.

Figure 31. Integration of the DB2 workload manager with Linux workload management



Defining mappings between DB2 service classes and Linux workload manager classes

The steps for integrating DB2 workload manager with Linux workload management, which runs as an operating system service, are as follows:

1. Define the Linux classes, class permissions, and processor shares by editing the `/etc/cgconfig.conf` control groups configuration file. What Linux classes you create depends on the conditions dictated by your business priorities for the work your data server performs. If you want to apply processor resource based on the source of certain work, for example, create a Linux class to match the DB2 service class that work is going to be assigned to by the workload identifying the work. Define an entry for each Linux class corresponding to the DB2 service class to be created that you want to use for the mapping. The following sections must be provided in the `/etc/cgconfig.conf` configuration file:
 - **group:** The Linux class name. For example, if you specify `group _class1`, you create a superclass `_class1`. If you specify `group _class1/_subclass1`, you create the subclass `_subclass1` under the superclass `_class1`.
 - **perm:** The permissions section that determines who can control what threads are assigned to a Linux class and who can change the processor shares of classes in the `/etc/cgconfig.conf` configuration file.

- task: The user ID (**uid**) and group ID (**gid**) whose threads can run in the Linux workload management class. To enable Linux workload management to work with DB2 workload manager, you should set **uid** to the DB2 instance owner user ID.
- admin: The user ID (**uid**) and group ID (**gid**) that can change processor shares for a Linux workload management class.
- cpu: The processor shares definition section
 - cpu.shares: The share assigned to this Linux class relative to the default class

The `/etc/cgconfig.conf` configuration file must contain these sections in the following format:

```
# Superclass name
group _name
{
    perm
    {
        task
        {
            uid = db2inst1;
            gid = db2iadml;
        }
        admin
        {
            uid = db2inst1;
            gid = db2iadml;
        }
    }

    cpu
    {
        cpu.shares = 1024;
    }
}
```

2. Start the Linux workload management service daemon with the service `cgconfig start` command, then start your DB2 data server with the `db2start` command.
3. To map a DB2 service class to one of the Linux classes, include the Linux class name in the `OUTBOUND CORRELATOR` clause when you create or alter the service class, which associates threads from the DB2 service class with the external Linux class.
4. If you want to find out what threads are assigned to a particular Linux class, you can use the `cat` command on the `/cgroup/class_name/tasks` file, where `class_name` represents the name of the Linux class you are interested in. All threads that are not mapped to a user-defined Linux class are assigned to the Linux default class, which you can find at `MOUNTPOINT/sysdefault`, where `MOUNTPOINT` is defined in the `cgconfig.conf` configuration file.
5. To add or remove Linux classes, you must stop with the Linux workload management service with the service `cgconfig stop` command, make your changes, and then restart the service. Note that stopping the service affects the entire system, because all tasks are moved to the default class. If you used the `/etc/init.d/cgred` script to start the service daemon, issue `/etc/init.d/cgred stop` to stop it.

For the integration with DB2 workload manager to work, you must ensure that the Linux workload management service is properly installed, configured, and active. If the DB2 data server cannot communicate with the Linux workload management

service, a message is logged to the db2diag log files and DB2 administrator log. Database activities will continue to execute.

The DB2 data server cannot detect whether the outbound correlator that it passes to external workload managers is recognized by Linux workload management. You must verify that the OUTBOUND CORRELATOR value specified for a DB2 service class matches the Linux class name so that DB2 threads are mapped to the Linux class. If an outbound correlator is not recognized, database activities will continue to execute.

Example

The following example illustrates how you can make use of Linux workload management processor controls by integrating with DB2 workload manager. In this example, we create two user-defined DB2 service classes, one for batch applications (BATCHAPPS) and one for online applications (ONLINEAPPS). For simplicity, this example does not show the default service classes, which should be included in an implementation that creates the recommended 1:1 mapping between DB2 service classes and Linux classes. Because response time is critical for the online applications, we want the ONLINEAPPS service class to receive three times the amount of processor shares relative to work that runs in the Linux default class ($3 \times 1024 = 3072$ shares). Batch applications have a lower business priority, and the BATCHAPPS class should be assigned half the processor resource of work that runs in the Linux default class ($1024 / 2 = 512$ shares). All other work on the system will run in the Linux default class. Note that this example does not create Linux classes corresponding to the three default DB2 workload manager service classes.

To create this setup, first create the two corresponding Linux classes `_BATCHAPPS` and `_ONLINEAPPS` and set their relative processor shares by editing the `/etc/cgconfig.conf` tasks file. After editing, the tasks file contains the following two entries, one for each Linux class:

```
# Superclass ONLINEAPPS
group _ONLINEAPPS
{
    perm
    {
        task
        {
            uid = db2inst1;
            gid = db2iadm1;
        }
        admin
        {
            uid = db2inst1;
            gid = db2iadm1;
        }
    }

    cpu
    {
        # 3 x 1024 = 3072 shares
        cpu.shares = 3072;
    }
}

# Superclass BATCHAPPS
group _BATCHAPPS
{
    perm
```

```

{
  task
  {
    uid = db2inst1;
    gid = db2iadml;
  }
  admin
  {
    uid = db2inst1;
    gid = db2iadml;
  }
}

cpu
{
  # 1024 / 2 = 512 shares
  cpu.shares = 512;
}
}

```

The absolute processor time in percent assigned to each Linux class as processor shares is as follows:

Table 56. Processor shares and absolute processor time assigned to Linux classes

Linux class	Shares	Absolute processor time in percent
Default class	1024 (default)	1024 / 4608 = 22%
_ONLINEAPPS	1024 x 3 = 3072	3072 / 4608 = 67%
_BATCHAPPS	1024 x ½ = 512	512 / 4608 = 11%
	Total = 1024 + 3072 + 512 = 4608 shares	

Once the Linux WLM classes are created, you can start the Linux workload management service:

```
service cgconfig start
```

Next, create the associated DB2 service classes with the following statements:

```
DB2 CREATE SERVICE CLASS BATCHAPPS OUTBOUND CORRELATOR '_BATCHAPPS'
DB2 CREATE SERVICE CLASS ONLINEAPPS OUTBOUND CORRELATOR '_ONLINEAPPS'
```

To find out which threads are running in a Linux class, issue the cat command. For the business critical _ONLINEAPPS Linux class, the command and output look as follows. You can see that there are six thread running in this Linux class:

```
cat /cgroup/_ONLINEAPPS/tasks
```

```
1056
1087
1107
985
1036
1205
```

Chapter 6. Tutorial for DB2 workload manager

The exercises in this tutorial were designed to provide you with a hands-on introduction to DB2 workload manager (WLM). Each exercise highlights one or more of the workload management features available with DB2 workload manager.

These exercises provide some guidance for using DB2 workload manager features which you can adapt for your own purposes, but you should note that the initial configuration you chose for your own data server may differ and should be based on your specific workload management objectives.

Before you begin

This tutorial is designed to be run against the SAMPLE database and, unless noted otherwise, requires DBADM or WLMADM authority (or SQLADM authority if only the COLLECT ACTIVITY DATA clause is specified). You should also start the instance and activate the SAMPLE database before continuing:

```
db2start
db2 activate db sample
```

Some of the command and query statements shown in these exercises are quite long. You can find most of these statements in the text file wlm-tutorial-steps.txt, which you can copy from when working through the exercises. The scripts representing the workloads that are required for the different exercises are also included.

Both wlm-tutorial-steps.txt and the workload scripts can be found here.

Exercise 1: Getting started with basic monitoring using default DB2 workload manager objects

This exercise demonstrates the basic types of monitoring information that can be obtained from the default workload and service class objects.

Estimated time: 20-25 minutes

By default, the user workload (SYSDEFAULTUSERWORKLOAD) and a default user service class (SYSDEFAULTUSERCLASS) are always created for each database. These default objects can be used to take advantage of the new DB2 workload manager monitoring features without having to create any user defined workloads or service classes. If no user defined workloads and service classes are created, all user activities will be associated with these default objects.

There are two separate features of monitoring that are demonstrated by this exercise:

1. The ability to collect aggregate statistics for all activities that run in a service class. Aggregate activity statistics provide an inexpensive way of looking at work in a service class as a whole. They show information like the number of activities that ran in the service class, and the average lifetime of those activities.
2. The ability to capture information about individual activities. Activity information can be useful when investigating the performance or behavior of a

particular activity. Activity information includes things such as statement text, compilation environment, etc. Activity information is more expensive to collect than aggregate activity statistics and is usually targeted towards a specific subset of activities.

Step 1: Create and enable event monitors

Connect to the database and create and enable event monitors for activities and statistics.

```
CONNECT TO SAMPLE
```

```
CREATE EVENT MONITOR DB2ACTIVITIES FOR ACTIVITIES WRITE TO TABLE  
CREATE EVENT MONITOR DB2STATISTICS FOR STATISTICS WRITE TO TABLE
```

```
SET EVENT MONITOR DB2ACTIVITIES STATE 1  
SET EVENT MONITOR DB2STATISTICS STATE 1
```

Step 2: Collect individual activities

Enable collection of individual activities using the `COLLECT ACTIVITY DATA` clause on the `CREATE` or `ALTER WORKLOAD STATEMENT`. When the `COLLECT ACTIVITY DATA` clause is specified for a workload, information about any activity submitted by an occurrence of that workload will be sent to the active `ACTIVITIES` event monitor when the activity completes. The `COLLECT ACTIVITY DATA` clause permits you to specify how much information should be collected by applying one of the following options:

- `WITHOUT DETAILS`: Collect activity information without statement and compilation environment.
- `WITH DETAILS`: Collect activity information including statement and compilation environment.
- `WITH DETAILS AND VALUES`: Collect activity information including statement and compilation environment, and input data values.

For this exercise, you will specify the `WITH DETAILS` clause so that the statement text information is captured.

```
ALTER WORKLOAD SYSDEFAULTUSERWORKLOAD  
  COLLECT ACTIVITY DATA ON COORDINATOR WITH DETAILS
```

In this example activity data is collected for the default user workload. This results in information about all user activities being collected since no other user defined workloads are currently active. This would be too expensive in a production environment. A better approach would be to isolate the activities of interest using a specific user defined workload or service class and apply the `COLLECT ACTIVITY DATA` clause to that workload or service class only.

Additional Information: The `COLLECT ACTIVITY DATA` clause can also be specified on a service class, work class (using a work action) or a threshold. If the clause is specified for a service class, information will be collected for any activity that runs in the service class. If it is specified for a work class (using a work action), any activity that has the work action applied to it will be collected. If the clause is specified for a threshold, activity information will be collected if the threshold is violated.

Step 3: Collect aggregate activity statistics

Enable collection of aggregate activity statistics for the default subclass under the default user service class using the COLLECT AGGREGATE ACTIVITY DATA clause. When this clause is specified, aggregate statistics will be maintained in memory for the corresponding service class (for example, statistics such as average activity lifetime). The in-memory statistics can be viewed using the service subclass statistics table function, or can be collected and sent to the active statistics event monitor for later analysis.

```
ALTER SERVICE CLASS SYSDEFAULTSUBCLASS UNDER SYSDEFAULTUSERCLASS
  COLLECT AGGREGATE ACTIVITY DATA BASE
```

Additional Information: There is a set of statistics collected by default for all DB2 workload manager objects. The COLLECT AGGREGATE ACTIVITY DATA clause enables collection of a number of additional optional statistics, such as the activity lifetime histogram.

In this example all user activities will be run in the SYSDEFAULTSUBCLASS service subclass under the SYSDEFAULTUSERCLASS service super class because no user defined service classes have been created. Therefore, information will be collected for all user activities.

Step 4: Run some activities

Run some activities, which will result in statistics being updated and the activities being collected.

```
db2 -o -tvf work1.db2
db2 -o -tvf work2.db2
```

The scripts representing applications (such as work1.db2 and work2.db2) disconnect you from the database, so that after running them you will need to reconnect.

Step 5: View in-memory statistics

You can view the in-memory service class statistics using the WLM_GET_SERVICE_SUBCLASS_STATS_V97 table function. For example:

```
CONNECT TO SAMPLE

SELECT VARCHAR(SERVICE_SUPERCLASS_NAME, 30) AS SUPERCLASS,
       VARCHAR(SERVICE_SUBCLASS_NAME, 30) AS SUBCLASS,
       LAST_RESET,
       COORD_ACT_COMPLETED_TOTAL,
       COORD_ACT_REJECTED_TOTAL,
       COORD_ACT_ABORTED_TOTAL,
       COORD_ACT_LIFETIME_AVG
FROM TABLE(SYSPROC.WLM_GET_SERVICE_SUBCLASS_STATS_V97 ( 'SYSDEFAULTUSERCLASS',
                                                         'SYSDEFAULTSUBCLASS',
                                                         -1 )) AS T
```

The output from this query will look something such as the following:

```
SUPERCLASS          SUBCLASS          LAST_RESET
COORD_ACT_COMPLETED_TOTAL COORD_ACT_REJECTED_TOTAL COORD_ACT_ABORTED_TOTAL
COORD_ACT_LIFETIME_AVG
-----
-----
-----
```

```

SYSDEFAULTUSERCLASS      SYSDEFAULTSUBCLASS      2007-07-18-16.03.51.752190
74                        0                        0      +1.40288000000000E+002

```

1 record(s) selected.

The COORD_ACT_COMPLETED_TOTAL column indicates how many activities have completed successfully in this service class. The last reset time indicates the last time that statistics were reset for this service class.

Additional Information: If you do not enable aggregate activity statistics for a service class using the COLLECT AGGREGATE ACTIVITY DATA clause, some statistics reported by the WLM_GET_SERVICE_SUBCLASS_STATS_V97 table function will be NULL.

Step 6: Send in-memory statistics to event monitor

Use the WLM_COLLECT_STATS stored procedure to send the in-memory statistics for all DB2 workload manager objects to the active statistics event monitor. When statistics are collected and sent to the statistics event monitor, the in-memory values are reset.

```
CALL SYSPROC.WLM_COLLECT_STATS()
```

Additional Information: If there is no active statistics event monitor, you can still use the WLM_COLLECT_STATS procedure to reset the in-memory statistics, but the current in-memory values will be lost. It is possible to automate workload management statistics collection using the WLM_COLLECT_INT database configuration parameter. If you set this parameter to a nonzero value, workload management statistics will be collected automatically every **wlm_collect_int** minutes (as if you manually invoked the WLM_COLLECT_STATS procedure every **wlm_collect_int** minutes).

Step 7: View in-memory statistics again

Invoke the WLM_GET_SERVICE_SUBCLASS_STATS_V97 table function again. Note that the LAST_RESET timestamp has been updated and the statistics have been reset.

```

SELECT VARCHAR(SERVICE_SUPERCLASS_NAME, 30) AS SUPERCLASS,
       VARCHAR(SERVICE_SUBCLASS_NAME, 30) AS SUBCLASS,
       LAST_RESET,
       COORD_ACT_COMPLETED_TOTAL,
       COORD_ACT_REJECTED_TOTAL,
       COORD_ACT_ABORTED_TOTAL,
       COORD_ACT_LIFETIME_AVG
FROM TABLE(SYSPROC.WLM_GET_SERVICE_SUBCLASS_STATS_V97 ( 'SYSDEFAULTUSERCLASS',
                                                         'SYSDEFAULTSUBCLASS',
                                                         -1 )) AS T

```

The output will look something like:

```

SUPERCLASS      SUBCLASS      LAST_RESET
COORD_ACT_COMPLETED_TOTAL COORD_ACT_REJECTED_TOTAL
COORD_ACT_ABORTED_TOTAL COORD_ACT_LIFETIME_AVG
-----
-----
SYSDEFAULTUSERCLASS      SYSDEFAULTSUBCLASS      2007-07-18-
16.04.03.505818      0      0
0      +0.00000000000000E+000

```

1 record(s) selected.

Step 8: View service class statistics collected by the statistics event monitor

The WLM_COLLECT_STATS procedure sent the in-memory service class statistics to the statistics event monitor. You can look at the statistics that were collected by the event monitor using statement such as the following:

```
SELECT VARCHAR(SERVICE_SUPERCLASS_NAME, 30) AS SUPERCLASS,
       VARCHAR(SERVICE_SUBCLASS_NAME, 30) AS SUBCLASS,
       LAST_WLM_RESET,
       STATISTICS_TIMESTAMP,
       COORD_ACT_COMPLETED_TOTAL,
       COORD_ACT_REJECTED_TOTAL,
       COORD_ACT_ABORTED_TOTAL,
       COORD_ACT_LIFETIME_AVG
FROM SCSTATS_DB2STATISTICS
```

The output will look something like:

```
SUPERCLASS          SUBCLASS
LAST_WLM_RESET      STATISTICS_TIMESTAMP
COORD_ACT_COMPLETED_TOTAL COORD_ACT_REJECTED_TOTAL
COORD_ACT_ABORTED_TOTAL COORD_ACT_LIFETIME_AVG
-----
SYSDEFAULTSYSTEMCLASS      SYSDEFAULTSUBCLASS      2007-07-18-
16.03.46.333724 2007-07-18-16.04.03.505818      0
0              0              -1
SYSDEFAULTMAINTENANCECLASS      SYSDEFAULTSUBCLASS      2007-07-18-
16.03.46.334301 2007-07-18-16.04.03.505818      0
0              0              -1
SYSDEFAULTUSERCLASS      SYSDEFAULTSUBCLASS      2007-07-18-
16.03.51.752190 2007-07-18-16.04.03.505818      75
0              0              136
```

3 record(s) selected.

Every time statistics are sent to the event monitor, a statistics record will be created for each DB2 workload manager object. Note the two timestamps LAST_WLM_RESET and STATISTICS_TIMESTAMP. The interval of time from LAST_WLM_RESET to STATISTICS_TIMESTAMP indicates the period of time over which the statistics in that record were collected. The STATISTICS_TIMESTAMP indicates when the statistics were collected. Note that the average lifetime for activities on the coordinator is -1 for the default system and maintenance service classes. The average activity lifetime statistic is only maintained for a service class if aggregate activity statistics are enabled using the COLLECT AGGREGATE ACTIVITY DATA clause.

Step 9: View activity information

Information about every individual activity associated with the default user workload was also collected by the activities event monitor, due to the specification of the COLLECT ACTIVITY DATA clause on the default workload in step 2. You can look at this activity information using a query such as the following:

```
SELECT VARCHAR(A.APPL_NAME, 15) as APPL_NAME,
       VARCHAR(A.TPMON_CLIENT_APP, 20) AS CLIENT_APP_NAME,
       VARCHAR(A.APPL_ID, 30) as APPL_ID,
       A.ACTIVITY_ID,
       A.UOW_ID,
       VARCHAR(S.STMT_TEXT, 300) AS STMT_TEXT
FROM ACTIVITY_DB2ACTIVITIES AS A,
```

```

ACTIVITYSTMT_DB2ACTIVITIES AS S
WHERE A.APPL_ID = S.APPL_ID AND
      A.ACTIVITY_ID = S.ACTIVITY_ID AND
      A.UOW_ID = S.UOW_ID

```

The output will look something like:

```

APPL_NAME      CLIENT_APP_NAME  APPL_ID
ACTIVITY_ID    UOW_ID          STMT_TEXT
-----
-----
-----
-----
-----
-----
-----
db2bp          CLP wlmmonbasic.db2 *LOCAL.db2inst1.070718200344
1              8 ALTER SERVICE CLASS SYSDEFAULTSUBCLASS UNDER
SYSDEFAULTUSERCLASS COLLECT AGGREGATE ACTIVITY DATA BASE
db2bp          CLP work1.db2      *LOCAL.db2inst1.070718200352
1              1 values(current client_applname)
db2bp          CLP work1.db2      *LOCAL.db2inst1.070718200352
2              1 select * from org
db2bp          CLP work1.db2      *LOCAL.db2inst1.070718200352
3              1 select * from employee
db2bp          CLP work1.db2      *LOCAL.db2inst1.070718200352
4              1 select * from sales
...

```

Note that you may see some truncation warnings (SQL0445).

When CLP executes a script, it will set the CURRENT CLIENT_APPLNAME special register to "CLP **script name**". So you can tell from the query above which script submitted each activity.

Step 10: Reset for next exercise

Update the SYSDEFAULTUSERWORKLOAD workload and the SYSDEFAULTSUBCLASS service subclass so that no activity data or aggregate activity statistics is collected, disable event monitors and clear out the activity and statistics tables, and call WLM_COLLECT_STATS() to reset the statistics.

```

ALTER WORKLOAD SYSDEFAULTUSERWORKLOAD COLLECT ACTIVITY DATA NONE

ALTER SERVICE CLASS SYSDEFAULTSUBCLASS UNDER SYSDEFAULTUSERCLASS
COLLECT AGGREGATE ACTIVITY DATA NONE

SET EVENT MONITOR DB2ACTIVITIES STATE 0
SET EVENT MONITOR DB2STATISTICS STATE 0

DELETE FROM ACTIVITY_DB2ACTIVITIES
DELETE FROM ACTIVITYSTMT_DB2ACTIVITIES
DELETE FROM SCSTATS_DB2STATISTICS
DELETE FROM WLSTATS_DB2STATISTICS

CALL WLM_COLLECT_STATS()

```

Exercise 2: Isolating activities using service classes and workloads

This exercise demonstrates how to create service classes and how to send activities to a service class using a workload. It also demonstrates how to use some of the WLM monitoring features to determine the workload that activities are being mapped to and to get information about activities being run in a service class and under a workload.

Estimated time: 20-25 minutes

Service classes are the primary point of resource control for database activities. They are also useful for monitoring. For example, you can collect statistics for activities in a particular service class to determine whether the performance goals for that service class are being met. By default, three default service classes (SYSDEFAULTSYSTEMCLASS, SYSDEFAULTMAINTENANCECLASS, and SYSDEFAULTUSERCLASS) are created for each database. If no user defined service classes are created, user activities are run under the default user service class (SYSDEFAULTUSERCLASS).

A workload is an entity that groups one or more units of work based on criteria such as system user ID, session user ID, etc. Workloads provide a means of assigning work to a service class so that the work can later be managed. A default user workload (SYSDEFAULTUSERWORKLOAD) and a default administration workload (SYSDEFAULTADMWORKLOAD) are created for each database. If no user defined workloads are created, all user activities are associated with the default user workload.

There are four separate features that are demonstrated in this exercise:

- How to create a service class.
- How to create a workload.
- How to examine basic workload statistics.
- How to collect activity information for activities run under an individual workload.

Step 1: Examine where activities are run with no user-defined service classes and workloads

First examine where activities are executed if there is no user defined service class or workload. All DB2 activities are assigned to a workload and run in a service class. If no user defined service classes are created, activities run in the default subclass (SYSDEFAULTSUBCLASS) under the default user service class (SYSDEFAULTUSERCLASS) and if no user defined workloads are created, activities run under the default user workload (SYSDEFAULTUSERWORKLOAD).

Run the work1.db2 and work2.db2 scripts and then examine the in-memory statistics for the SYSDEFAULTSUBCLASS of SYSDEFAULTUSERCLASS using the WLM_GET_SERVICE_SUBCLASS_STATS_V97 .

```
db2 -o -tvf work1.db2
db2 -o -tvf work2.db2
```

```
CONNECT TO SAMPLE
```

```
SELECT VARCHAR( SERVICE_SUPERCLASS_NAME, 30) SUPERCLASS,
       VARCHAR( SERVICE_SUBCLASS_NAME, 30) SUBCLASS,
       COORD_ACT_COMPLETED_TOTAL
FROM TABLE(WLM_GET_SERVICE_SUBCLASS_STATS_V97('',' ',-1)) AS T
```

You will see output such as the following:

SUPERCLASS D_TOTAL	SUBCLASS	COORD_ACT_COMPLETE
-----	-----	-----

SYSDEFAULTSYSTEMCLASS	SYSDEFAULTSUBCLASS	
0		
SYSDEFAULTMAINTENANCECLASS	SYSDEFAULTSUBCLASS	
0		

```
SYSDEFAULTUSERCLASS      SYSDEFAULTSUBCLASS
      75
```

3 record(s) selected.

Note all the activities are run in the SYSDEFAULTUSERCLASS service super class.

Additional Information: There are 2 other service classes as well, SYSDEFAULTSYSTEMCLASS and SYSDEFAULTMAINTENANCECLASS. These service classes are used for internal maintenance and system level tasks. User activities will not run in these service classes. You may notice nonzero activity counts in these service classes as well if the DB2 data server has issued any internal activities.

Use the WLM_GET_WORKLOAD_STATS_V97 table function to view workload statistics to determine which workload the applications are being associated with.

```
SELECT SUBSTR(WORKLOAD_NAME, 1, 22) AS WL_DEF_NAME,
       WLO_COMPLETED_TOTAL,
       CONCURRENT_WLO_ACT_TOP FROM
       TABLE(WLM_GET_WORKLOAD_STATS_V97(CAST(NULL AS VARCHAR(128))), -2))
AS WLSTATS
```

The output will look something such as the following:

WL_DEF_NAME	WLO_COMPLETED_TOTAL	CONCURRENT_WLO_ACT_TOP
SYSDEFAULTUSERWORKLOAD	3	5
SYSDEFAULTADMWORKLOAD	0	0

2 record(s) selected.

Note there is one workload occurrence completed for both of the scripts (work1.db2 and work2.db2) as well as a workload occurrence for the connection used to execute the previous command.

Step 2: Create a service class and workload

Create a service class and then create a workload such that all activities run from the work1.db2 script get mapped to the newly created service class. When CLP executes a script, the CURRENT_CLIENT_APPLNAME special register value is set to "CLP script name".

```
CREATE SERVICE CLASS work1_sc

CREATE WORKLOAD work1_w1 CURRENT CLIENT_APPLNAME('CLP work1.db2')
      SERVICE CLASS work1_sc
```

Additional Information: There are a number of attributes that can be specified when creating a workload or a service class. For example, when creating a workload, you can identify the connection based on application name, session user, etc. For more information, refer the CREATE WORKLOAD and the CREATE SERVICE CLASS documentation.

Step 3: Grant usage on workload

Grant usage on the workload (requires ACCESSCTRL or SECADM authority).

```
GRANT USAGE ON WORKLOAD work1_w1 TO PUBLIC
```

Additional information: A connection can be associated with a workload only if the session user has USAGE privilege on the workload. This is necessary to prevent

users from changing connection attributes of their application in an attempt to run their work in a higher priority service class. Some connection attributes can be changed programmatically (using the `sqlseti` API, for example). In this exercise, we just grant `USAGE` privilege to `PUBLIC`. You would want to be more discriminating on a real system. Since the sample will be run as `DBADM`, this step could be skipped altogether.

Step 4: Reset in-memory statistics

Reset the in-memory statistics using the `WLM_COLLECT_STATS` function, to clear the statistics collected.

```
CALL SYSPROC.WLM_COLLECT_STATS()
```

Step 5: Run some activities

Run both the `work1.db2` and the `work2.db2` scripts.

```
db2 -o -tvf work1.db2
db2 -o -tvf work2.db2
```

Step 6: View workload and service class statistics

Use the `WLM_GET_WORKLOAD_STATS_V97` table function to view workload statistics to determine which workload the applications are being associated with.

```
CONNECT TO SAMPLE
```

```
SELECT SUBSTR(WORKLOAD_NAME, 1, 22) AS WL_DEF_NAME,
       WLO_COMPLETED_TOTAL,
       CONCURRENT_WLO_ACT_TOP
FROM TABLE(WLM_GET_WORKLOAD_STATS_V97(CAST(NULL AS VARCHAR(128)), -2))
AS WLSTATS
```

The output will look something such as the following:

WL_DEF_NAME	WLO_COMPLETED_TOTAL	CONCURRENT_WLO_ACT_TOP
WORK1_WL	1	5
SYSDEFAULTUSERWORKLOAD	1	5
SYSDEFAULTADMWORKLOAD	0	0

Note that one workload occurrence completed under `WORK1_WL` which is the `work1.db2` script. One workload occurrence completed under `SYSDEFAULTUSERWORKLOAD` which is the `work2.db2` script.

You may see a 2nd workload occurrence completed for the `SYSDEFAULTUSERWORKLOAD` which is the connection that was used to call the `WLM_COLLECT_STATS` procedure. `WLM_COLLECT_STATS` is an asynchronous procedure which might be completed before the statistics are actually collected and therefore might be included.

You can also use the `WLM_GET_SERVICE_SUBCLASS_STATS_V97` table function to show which service class the activities are being run under as a result of creating the new workload.

```
SELECT VARCHAR( SERVICE_SUPERCLASS_NAME, 30) SUPERCLASS,
       VARCHAR( SERVICE_SUBCLASS_NAME, 23) SUBCLASS,
       COORD_ACT_COMPLETED_TOTAL COORDACTCOMP
FROM TABLE(WLM_GET_SERVICE_SUBCLASS_STATS_V97('','",-1)) AS T
```

The results looking something such as the following:

SUPERCLASS	SUBCLASS	COORDACTCOMP
SYSDEFAULTSYSTEMCLASS	SYSDEFAULTSUBCLASS	0
SYSDEFAULTMAINTENANCECLASS	SYSDEFAULTSUBCLASS	0
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	37
WORK1_SC	SYSDEFAULTSUBCLASS	37

Note the activities that completed under the WORK1_SC due to the WORK1_WL workload mapping.

Step 7: Create another service class and workload

Create a second service class and then create a workload such that all activities run from the work2.db2 application get mapped to the newly created service class. In addition, set up the workload so that it will collect some activity data. For this example, we just collect activity data without any additional details or values.

```
CREATE SERVICE CLASS work2_sc

CREATE WORKLOAD work2_w1
    CURRENT_CLIENT_APPLNAME('CLP work2.db2')
    SERVICE CLASS work2_sc
    COLLECT ACTIVITY DATA
```

Additional information: When the COLLECT ACTIVITY DATA clause is specified for a workload, information about any activity submitted by an occurrence of that workload will be sent to the active ACTIVITIES event monitor when the activity completes. The COLLECT ACTIVITY DATA clause permits you to specify how much information should be collected by applying one of the following options:

- WITHOUT DETAILS: Collect activity information without statement and compilation environment (the default)
- WITH DETAILS: Collect activity information including statement and compilation environment.
- WITH DETAILS AND VALUES: Collect activity information including statement and compilation environment, and input data values.

Step 8: Enable the activities event monitor

Enable the event monitors for activities.

The activity monitor was created in Exercise 1.

```
SET EVENT MONITOR DB2ACTIVITIES STATE 1
```

Step 9: Reset in-memory statistics and run some activities

Use the WLM_COLLECT_STATS stored procedure to reset the statistics again and run the work1.db2 and work2.db2 scripts again.

```
CALL SYSPROC.WLM_COLLECT_STATS()
```

```
db2 -o -tvf work1.db2
db2 -o -tvf work2.db2
```

Step 10: View workload and service class statistics

Use the WLM_GET_WORKLOAD_STATS_V97 table function again to determine which workload the applications are being associated with

```

CONNECT TO SAMPLE

SELECT SUBSTR(WORKLOAD_NAME, 1, 22) AS WL_DEF_NAME,
       WLO_COMPLETED_TOTAL,
       CONCURRENT_WLO_ACT_TOP
FROM TABLE(WLM_GET_WORKLOAD_STATS_V97(CAST(NULL AS VARCHAR(128)), -2))
AS WLSTATS

```

The output will look something such as the following:

WL_DEF_NAME	WLO_COMPLETED_TOTAL	CONCURRENT_WLO_ACT_TOP
WORK1_WL	1	5
WORK2_WL	1	5
SYSDEFAULTUSERWORKLOAD	0	0
SYSDEFAULTADMWORKLOAD	0	0

Note this time both workload definitions have a workload occurrence run, once for each script.

You may or may not see a workload occurrence completed for the SYSDEFAULTUSERWORKLOAD depending on whether workload occurrence over which the call to the WLM_COLLECT_STATS procedure was submitted is closed before the statistics are collected.

Use WLM_GET_SERVICE_SUBCLASS_STATS_V97 again to show which service class the activities are being run under as a result of creating the new workload.

```

SELECT VARCHAR( SERVICE_SUPERCLASS_NAME, 30) SUPERCLASS,
       VARCHAR( SERVICE_SUBCLASS_NAME, 23) SUBCLASS,
       COORD_ACT_COMPLETED_TOTAL COORDACTCOMP
FROM TABLE(WLM_GET_SERVICE_SUBCLASS_STATS_V97('','",-1)) AS T

```

With the results looking something like:

SUPERCLASS	SUBCLASS	COORDACTCOMP
SYSDEFAULTSYSTEMCLASS	SYSDEFAULTSUBCLASS	0
SYSDEFAULTMAINTENANCECLASS	SYSDEFAULTSUBCLASS	0
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	1
WORK1_SC	SYSDEFAULTSUBCLASS	37
WORK2_SC	SYSDEFAULTSUBCLASS	37

Note this time service super class work2_sc has some activities run under it due to the WORK2_WL mapping. The one activity under SYSDEFAULTUSERCLASS is the query previously run on WLM_GET_WORKLOAD_STATS_V97.

Step 11: View the activity data collected

Query the activity table for information on the activities that have been run. Note that only the activities from the work2.db2 script have been collected because only the work2_wl workload definition has the COLLECT ACTIVITY DATA attribute specified.

```

SELECT SUBSTR(WORKLOADNAME, 1, 20) WL_DEF_NAME,
       SUBSTR(APPL_NAME, 1, 20) APPL_NAME,
       SUBSTR(ACTIVITY_TYPE, 1, 10) ACT_TYPE
FROM SYSIBM.SYSWORKLOADS, ACTIVITY_DB2ACTIVITIES
WHERE WORKLOADID = WORKLOAD_ID

```

The results look something like:

WL_DEF_NAME	APPL_NAME	ACT_TYPE
WORK2_WL	db2bp	READ_DML

WORK2_WL	db2bp	READ_DML
WORK2_WL	db2bp	READ_DML
WORK2_WL	db2bp	READ_DML
WORK2_WL	db2bp	WRITE_DML
WORK2_WL	db2bp	WRITE_DML
WORK2_WL	db2bp	WRITE_DML
WORK2_WL	db2bp	WRITE_DML
WORK2_WL	db2bp	WRITE_DML
WORK2_WL	db2bp	DDL
WORK2_WL	db2bp	DDL
WORK2_WL	db2bp	DDL
WORK2_WL	db2bp	OTHER
WORK2_WL	db2bp	READ_DML
WORK2_WL	db2bp	READ_DML
WORK2_WL	db2bp	READ_DML
WORK2_WL	db2bp	READ_DML
WORK2_WL	db2bp	READ_DML
WORK2_WL	db2bp	READ_DML
WORK2_WL	db2bp	OTHER
WORK2_WL	db2bp	DDL
WORK2_WL	db2bp	LOAD
WORK2_WL	db2bp	DDL
WORK2_WL	db2bp	DDL
WORK2_WL	db2bp	DDL
WORK2_WL	db2bp	DDL
WORK2_WL	db2bp	DDL
WORK2_WL	db2bp	READ_DML
WORK2_WL	db2bp	CALL
WORK2_WL	db2bp	READ_DML
WORK2_WL	db2bp	CALL
WORK2_WL	db2bp	CALL
WORK2_WL	db2bp	CALL
WORK2_WL	db2bp	DDL
WORK2_WL	db2bp	DDL
WORK2_WL	db2bp	DDL
WORK2_WL	db2bp	DDL
:		
:		

Step 12: Assign resources to service classes

Now that you have isolated the activities issued by these two scripts into separate service classes, you can assign resources to the service classes or monitor the activities that run in those service classes. A few examples: If the work performed by the script work2.db2 is more important than the work performed by the script work1.db2, you could increase the priority of agents running in the WORK2_SC service class using a statement such as the following.

On UNIX operating environments (a negative value specifies a higher priority):

```
ALTER SERVICE CLASS WORK2_SC AGENT PRIORITY -6
```

On Windows operating environments (a positive value specifies a higher priority):

```
ALTER SERVICE CLASS WORK2_SC AGENT PRIORITY 6
```

If you wanted to capture details about every individual activity that executes in the WORK2_SC service class, you could enable activity collection for that service class using the following:

```
ALTER SERVICE CLASS SYSDEFAULTSUBCLASS UNDER WORK2_SC
  COLLECT ACTIVITY DATA ON COORDINATOR WITH DETAILS
```

Step 13: Reset for next exercise

Update workload work2_w1 so that no activity data is collected, disable the event monitor and clean up the event monitor table, and call WLM_COLLECT_STATS() to reset the statistics.

```
ALTER WORKLOAD work2_w1
    COLLECT ACTIVITY DATA NONE

SET EVENT MONITOR DB2ACTIVITIES STATE 0

DELETE FROM ACTIVITY_DB2ACTIVITIES

CALL WLM_COLLECT_STATS()
```

Exercise 3: Using thresholds to control rogue activities and using the threshold violation monitor

This exercise demonstrates how you can use thresholds to detect resource misuse or the beginning of system overload by establishing limits over the consumption of a specific resource.

Estimated time: 15-20 minutes

If a threshold is violated, a specified action can be triggered. The supported actions are:

- **STOP EXECUTION:** Stop processing the activity that caused the threshold to be violated.
- **CONTINUE:** Continue processing
- **Collect information about the activity that violated the threshold.** This action can be specified in conjunction with the **CONTINUE** or **STOP EXECUTION** action.

Regardless of whether an activity that violates a threshold is stopped or permitted to continue running, a record of the violation is written to an active **THRESHOLD VIOLATIONS** event monitor (assuming one is defined in advance) each time a threshold is violated. The record contains information such as which threshold was violated, the time of the violation, and the threshold action.

This exercise demonstrates how thresholds can be used to detect or prevent rogue activities from running on your system and using up system resources. A rogue activity is any activity that uses an unexpectedly high amount of resources. For example, a query that runs for an abnormally long time, or returns an unexpectedly large result set.

Step 1: Create a threshold violation event monitor

Create and enable a write-to-table event monitor that will be used to capture the threshold violation information and enable the activity event monitor that was created in Exercise 1.

```
CREATE EVENT MONITOR threvio FOR THRESHOLD VIOLATIONS WRITE TO TABLE
    THRESHOLDVIOLATIONS(IN userspace1),
    CONTROL(IN userspace1)

SET EVENT MONITOR threvio STATE 1

SET EVENT MONITOR db2activities STATE 1
```

Step 2: Create a workload

Create a workload such that all activities run from the workth.db2 script will get mapped to the work1_sc service class.

The work1_sc service class already exists since it was created in Exercise 2.

```
CREATE WORKLOAD workth_wl
    CURRENT CLIENT_APPLNAME('CLP workth.db2')
    SERVICE CLASS work1_sc
```

Step 3: Create thresholds

Create two thresholds, one of which (th_estcost) is an ESTIMATEDSQLCOST threshold and another (th_sqlrows) is a SQLROWSRETURNED threshold and apply them to the service class you wish to control the activities for (in this case, work1_sc service class).

The th_estcost threshold specifies an upper bound (10000 timerons) for the optimizer-estimated cost (in timerons) for an activity running in the work1_sc service class. If any query with an estimated cost greater than 10000 timerons, tries to execute in the work1_sc service class, this threshold is violated and the query is not permitted to run.

The th_sqlrows threshold specifies that any activity running in the work1_sc service class can return at most 30 rows from the data server. If any query tries to return more than 30 rows, this threshold is violated, only 30 rows will be returned to the client and the query will be stopped. In addition, data about the activity that caused the threshold violation will be collected.

In either case, when an activity violates the threshold, a threshold violation record is written to the THRESHOLD VIOLATIONS event monitor as defined in step 1 and the execution of the activity is stopped (because of the STOP EXECUTION action). The application that submitted the activity will receive an SQL4712N error.

```
CREATE THRESHOLD th_estcost
    FOR SERVICE CLASS work1_sc ACTIVITIES
    ENFORCEMENT DATABASE
    WHEN ESTIMATEDSQLCOST > 10000
    STOP EXECUTION

CREATE THRESHOLD th_sqlrows
    FOR SERVICE CLASS work1_sc ACTIVITIES
    ENFORCEMENT DATABASE
    WHEN SQLROWSRETURNED > 30
    COLLECT ACTIVITY DATA WITH DETAILS AND VALUES
    STOP EXECUTION
```

Additional information: A threshold can be either predictive or reactive:

- Predictive threshold: The boundaries of a predictive threshold are checked before the activity starts running. To check whether a predictive threshold would be violated, the data server obtains usage estimates from the query compiler. For this example, the th_estcost threshold is a predictive threshold.
- Reactive threshold: The boundaries of a reactive threshold are checked while an activity is executing. Approximate runtime usage estimates of the controlled resource are used to evaluate the boundaries of reactive thresholds. The runtime usage estimates are not obtained continuously but rather at selected predefined checkpoints during the lifetime of the tracked work. For this example, the th_sqlrows is a reactive threshold.

Step 4: Run some activities

Run some activities, some of which violate the threshold upper bounds defined in the previous step.

```
db2 -o -tvf workth.db2
```

Note that the statements which violate the thresholds defined above fail with an error of SQL4712N/SQLSTATE 5U026.

Step 5: View the threshold violation event monitor

Information about every threshold violation is collected by the THRESHOLD VIOLATIONS event monitor. You can query the threshold violation information by issuing regular SQL statements against the threshold violation monitor table as shown in the following example.

```
CONNECT TO SAMPLE
SELECT APPL_ID,
       UOW_ID,
       ACTIVITY_ID,
       COORD_PARTITION_NUM AS COORDPART,
       THRESHOLD_PREDICATE,
       THRESHOLD_ACTION,
       TIME_OF_VIOLATION
FROM THRESHOLDVIOLATIONS_THREVI0
ORDER BY THRESHOLD_ACTION, THRESHOLD_PREDICATE, TIME_OF_VIOLATION
```

The output will look something such as the following:

```
APPL_ID          UOW_ID
ACTIVITY_ID      COORDPART  THRESHOLD_PREDICATE  THRESHOLD_ACTION  TIME_OF_VIOLATION
-----
--
*LOCAL.DB2.070821150008          11
1          0 EstimatedSQLCost
          Stop          2007-08-21-
11.00.11.000000
*LOCAL.DB2.070821150008          10
1          0 SQLRowsReturned
          Stop          2007-08-21-
11.00.10.000000

2 record(s) selected.
```

Step 6: View information for the activity that violated the threshold

Activity information is collected for any activity that violates a threshold that is defined with a COLLECT clause. Show the detailed information about the activities that violated a threshold using the following query:

```
SELECT VARCHAR(A.APPL_NAME, 15) AS APPL_NAME,
       VARCHAR(A.TPMON_CLIENT_APP, 20) AS CLIENT_APP_NAME,
       A.ACTIVITY_ID,
       A.ACTIVITY_TYPE,
       A.WORKLOAD_ID,
       T.THRESHOLD_PREDICATE,
       A.QUERY_CARD_ESTIMATE,
       T.THRESHOLD_MAXVALUE,
       T.TIME_OF_VIOLATION,
       VARCHAR(A5.STMT_TEXT, 100) AS STMT_TEXT
```



```

FROM THRESHOLDVIOLATIONS_THREVIO AS T,
     ACTIVITY_DB2ACTIVITIES AS A,
     ACTIVITYSTMT_DB2ACTIVITIES AS AS
WHERE T.APPL_ID = A.APPL_ID AND
      T.UOW_ID = A.UOW_ID AND
      T.ACTIVITY_ID = A.ACTIVITY_ID AND
      A.APPL_ID = AS.APPL_ID AND
      A.ACTIVITY_ID = AS.ACTIVITY_ID AND
      A.UOW_ID = AS.UOW_ID

```

The output will look something such as the following:

```

APPL_NAME          CLIENT_APP_NAME      ACTIVITY_ID          ACTIVITY_TYPE
                   WORKLOAD_ID THRESHOLD_PREDICATE
                   QUERY_CARD_ESTIMATE THRESHOLD_MAXVALUE
TIME_OF_VIOLATION  STMT_TEXT
-----
db2bp              CLP workth.db2      3 READ_DML
                   3 SQLRowsReturned
                   41                               3
0 2007-08-31-09.01.16.000000 SELECT * FROM SALES

```

Note that the activity that violated the `th_estcost` (EstimatedSqlCost) threshold is not shown. The reason is that the threshold did not specify the `COLLECT ACTIVITY DATA` clause, so that no activity data was collected for that activity.

Step 7: Reset for next exercise

Disable the event monitors that were enabled. Also disable and drop the `th_estcost` and `th_sqlrows` thresholds that were created.

```

SET EVENT MONITOR threvio STATE 0
SET EVENT MONITOR db2activities STATE 0

```

```

ALTER THRESHOLD th_estcost DISABLE
DROP THRESHOLD th_estcost

```

```

ALTER THRESHOLD th_sqlrows DISABLE
DROP THRESHOLD th_sqlrows

```

Also clean up the activities event monitor tables and the threshold violation table

```

DELETE from ACTIVITY_DB2ACTIVITIES
DELETE from ACTIVITYSTMT_DB2ACTIVITIES
DELETE from THRESHOLDVIOLATIONS_THREVIO

CALL WLM_COLLECT_STATS()

```

Exercise 4: Differentiating activities by activity type

This exercise demonstrates how a work action set can be used to: Collect information about all activities of a certain type; apply a threshold to all activities of a certain type; isolate activities of a certain type by mapping them to a specific service subclass

Estimated time: 25-30 minutes

Work action sets are used to apply an action to an activity based on what the activity is doing rather than who submitted it (as is done with workloads).

Actions can be applied to either:

- All database activities of a certain type (using a database work action set).
- Only to activities of a certain type in a particular service class (using a service class work action set).

This exercise shows both methods.

Additional Information: There are other actions that can be applied, such as collecting statistics for activities of a certain type that are not covered in this exercise.

Step 1: Create a work class set

First, create a work class set containing work classes that will represent the specific types of activities you are interested in. This work class set will be used in conjunction with work action sets to perform actions on the selected types of activities. Below is an example that creates a work class set containing work classes of all possible types, but if you were interested only in one activity type, your work class set could be created to only contain that one work class.

```
CREATE WORK CLASS SET all_class_types
(WORK CLASS read_wc WORK TYPE READ,
 WORK CLASS write_wc WORK TYPE WRITE,
 WORK CLASS ddl_wc WORK TYPE DDL,
 WORK CLASS call_wc WORK TYPE CALL,
 WORK CLASS load_wc WORK TYPE LOAD,
 WORK CLASS all_wc WORK TYPE ALL POSITION LAST)
```

Step 2: Enable the activities event monitor

Enable the event monitor for activities that was created in Exercise 1.

```
SET EVENT MONITOR DB2ACTIVITIES STATE 1
```

Step 3: Create a database work action set

If you want to perform a particular action on all activities of a specific type (such as applying a threshold or collecting activity information), use a database work action set.

Create a work action set at the database level that contains work actions for the specific work class representing the type of activities you want isolated. For this example, we want to collect activity data for all DDL, READ and LOAD activities that run on the system and we also want to stop any large read activity from running. For this exercise, a large read activity is any select statement that has an estimated cost (in timerons) of greater than 10000.

```
CREATE WORK ACTION SET db_was FOR DATABASE
USING WORK CLASS SET all_class_types
(WORK ACTION collect_load_wa ON WORK CLASS load_wc
 COLLECT ACTIVITY DATA WITH DETAILS AND VALUES,
 WORK ACTION collect_ddl_wa ON WORK CLASS ddl_wc
 COLLECT ACTIVITY DATA WITH DETAILS AND VALUES,
 WORK ACTION collect_read_wa ON WORK CLASS read_wc
 COLLECT ACTIVITY DATA WITH DETAILS AND VALUES,
 WORK ACTION stop_large_read_wa on WORK CLASS read_wc
 WHEN ESTIMATEDSQLCOST > 10000 STOP EXECUTION )
```

Step 4: Run activities and view work action set statistics

Run the work1.db2 and work3.db2 scripts.

```
db2 -o -tvf work1.db2
db2 -o -tvf work3.db2
```

You can use the WLM_GET_WORK_ACTION_SET_STATS table function to access the work action set statistics in memory to get the number of times specific activity types have been run. Note that running the following query shows only the load_wc, read_wc and ddl_wc work classes since they are the only work classes that have an applicable work action. All the other activities are counted under the "*" :

```
CONNECT TO SAMPLE

SELECT SUBSTR(WORK_ACTION_SET_NAME, 1, 12) AS WORK_ACTION_SET_NAME,
       SUBSTR(WORK_CLASS_NAME, 1, 12) AS WORK_CLASS_NAME,
       LAST_RESET,
       SUBSTR(CHAR(ACT_TOTAL), 1, 10) AS TOTAL_ACTS
FROM TABLE(WLM_GET_WORK_ACTION_SET_STATS('', -2)) AS WASSTATS
ORDER BY WORK_ACTION_SET_NAME, WORK_CLASS_NAME
```

The output will look something like:

WORK_ACTION_SET_NAME	WORK_CLASS_NAME	LAST_RESET	TOTAL_ACTS
DB_WAS	*	2007-08-15-19.02.47.305556	12
DB_WAS	DDL_WC	2007-08-15-19.02.47.305556	12
DB_WAS	LOAD_WC	2007-08-15-19.02.47.305556	1
DB_WAS	READ_WC	2007-08-15-19.02.47.305556	13

4 record(s) selected.

Step 5: View the activity data collected

Information about every individual DDL, READ and LOAD activities was collected by the activities event monitor, due to the specification of the COLLECT ACTIVITY DATA work action that was applied to the ddl_wc, read_wc, and the load_wc work classes in step 3. Below are a couple of examples of how you might want to look at this activity information.

To get some basic information about the activities, you can simply query the activity monitor table with a statement such as the following:

```
SELECT ACTIVITY_ID,
       SUBSTR(ACTIVITY_TYPE, 1, 8) AS ACTIVITY_TYPE,
       VARCHAR(APPL_ID, 30) AS APPL_ID,
       VARCHAR(APPL_NAME, 10) AS APPL_NAME
FROM ACTIVITY_DB2ACTIVITIES
```

The output will look something like:

ACTIVITY_ID	ACTIVITY_TYPE	APPL_ID	APPL_NAME
1	READ_DML	*LOCAL.karenam.070815192410	db2bp
1	READ_DML	*LOCAL.karenam.070815192418	db2bp
2	READ_DML	*LOCAL.karenam.070815192418	db2bp
3	READ_DML	*LOCAL.karenam.070815192418	db2bp
4	READ_DML	*LOCAL.karenam.070815192418	db2bp
1	DDL	*LOCAL.karenam.070815192418	db2bp
2	DDL	*LOCAL.karenam.070815192418	db2bp
3	DDL	*LOCAL.karenam.070815192418	db2bp
2	READ_DML	*LOCAL.karenam.070815192418	db2bp
1	READ_DML	*LOCAL.karenam.070815192418	db2bp

```

2 READ_DML      *LOCAL.karenam.070815192418  db2bp
3 READ_DML      *LOCAL.karenam.070815192418  db2bp
4 READ_DML      *LOCAL.karenam.070815192418  db2bp
6 LOAD          *LOCAL.karenam.070815192418  db2bp
1 DDL           *LOCAL.karenam.070815192418  db2bp
1 DDL           *LOCAL.karenam.070815192418  db2bp
2 DDL           *LOCAL.karenam.070815192418  db2bp
3 DDL           *LOCAL.karenam.070815192418  db2bp
4 DDL           *LOCAL.karenam.070815192418  db2bp
5 READ_DML      *LOCAL.karenam.070815192418  db2bp
10 READ_DML     *LOCAL.karenam.070815192418  db2bp
1 DDL           *LOCAL.karenam.070815192418  db2bp
2 DDL           *LOCAL.karenam.070815192418  db2bp
3 DDL           *LOCAL.karenam.070815192418  db2bp
4 DDL           *LOCAL.karenam.070815192418  db2bp
1 READ_DML     *LOCAL.karenam.070815192426  db2bp

```

26 record(s) selected.

To obtain additional information about each activity, such as activity text and what service class it ran under, you can perform a query similar to this one:

```

SELECT VARCHAR(A.APPL_NAME, 15) as APPL_NAME,
       VARCHAR(A.TPMON_CLIENT_APP, 20) AS CLIENT_APP_NAME,
       VARCHAR(A.APPL_ID, 30) as APPL_ID,
       VARCHAR(A.SERVICE_SUPERCLASS_NAME, 20) as SUPER_CLASS,
       VARCHAR(A.SERVICE_SUBCLASS_NAME, 20) as SUB_CLASS,
       SQLCODE,
       VARCHAR(S.STMT_TEXT, 300) AS STMT_TEXT
FROM ACTIVITY_DB2ACTIVITIES AS A, ACTIVITYSTMT_DB2ACTIVITIES AS S
WHERE A.APPL_ID = S.APPL_ID AND
      A.ACTIVITY_ID = S.ACTIVITY_ID AND
      A.UOW_ID = S.UOW_ID

```

The output will look something like:

```

APPL_NAME      CLIENT_APP_NAME      APPL_ID
SUPER_CLASS
SUB_CLASS      SQLCODE      STMT_TEXT
-----
---
-----
---
-----
---
-----
---
-----
---
-----
db2bp          CLP wasdbsc.db2      *LOCAL.karenam.070815192410
SYSDEFAULTUS
ERCLASS  SYSDEFAULTSUBCLASS      0 SELECT DISTINCT CURRENT SQLID FROM
SYS
IBM.SYSTABLES

db2bp          CLP work1.db2        *LOCAL.karenam.070815192418
SYSDEFAULTUS
ERCLASS  SYSDEFAULTSUBCLASS      0 values(current client_applname)

:
:
db2bp          CLP work1.db2        *LOCAL.karenam.070815192418
SYSDEFAULTUS
ERCLASS  SYSDEFAULTSUBCLASS      0 drop procedure stp2

```

```

db2bp          CLP work3.db2          *LOCAL.karenam.070815192426
SYSDEFAULTUS
ERCLASS  SYSDEFAULTSUBCLASS          -4712 select count(*) from syscat.tables,
sy
scat.tables, syscat.tables, syscat.tables, syscat.tables, syscat.tables
:
:

```

Note that one of the activities has an SQLCODE of -4712. This indicates execution of the activity was stopped due to a threshold violation. The threshold defined for the stop_large_read_wa work action will prevent any SELECT statement with an estimated cost of greater than 10000 from executing.

Additional information: Load activities (not including load from a cursor) do not have an entry in the activity statement event monitor table (activitystmt_db2activities table) which explains why there is no record for the single load activity that is run by the work1.db2 script in the output shown in the last query above. The reason for this is that load activities are not SQL statements. For load from cursor activities, there is an entry for the cursor statement in the activity statement event monitor table because the cursor itself is a separate activity. There is an entry for all load activities in the activities event monitor table (activity_db2activities).

Step 6: Disable work actions

Before moving on to the service class work action set, drop the database work action set.

```
DROP WORK ACTION SET db_was
```

Additional information: Before dropping any concurrency threshold, that threshold must first be disabled. In this case, there are no work actions that represent a concurrency threshold but if there were, the only way to disable it would be by disabling the work action. A work action threshold cannot be manipulated through THRESHOLD SQL statements; they can be manipulated only through WORK ACTION SET SQL statements. Only work actions that represent concurrency thresholds need to be disabled before dropping the subsequent work action set. For this exercise, because there are no work actions that represent a concurrency threshold, there is no need to disable any of the work actions before dropping the work action set.

If you want to apply a particular action, such as a threshold, to all the activities of a certain type running in a service super class, you should consider using a service class work action set. You can create a mapping work action to map specific types of activities to a specific service subclass and then apply a threshold to that service subclass. The following steps demonstrate how service class work action sets might be used

Step 7: Create a service class to and create a workload

Create a service subclass under the work1_sc service super class that was created in Exercise 2 Step 2.

The service super class work1_sc is the service class that the activities will be mapped to through the workloads. The service subclass work1_sc_read is the service class that the read activities will be mapped to through the work action.

```
CREATE SERVICE CLASS work1_sc_read UNDER work1_sc
```

Create a workload so that all activities submitted by the work3.db2 script will be mapped to work1_sc service super class. Note that activities from work1.db2 are already being mapped to work1_sc from one of the previous exercises.

```
CREATE WORKLOAD work3_w1 CURRENT CLIENT_APPLNAME('CLP work3.db2')
  SERVICE CLASS work1_sc
```

Step 8: Create a service class work action set

Create a work action set at the service class level that contains work actions that apply to the specific work classes representing the types of activities you want isolated. For this example, we want to collect activity data for all DDL, read, and load activities that run under the work1_sc service class and we also want to map read activities to a separate service subclass so that we can treat them differently; in this case, a threshold will be applied to the service subclass to stop any large SELECT statements from running.

```
CREATE WORK ACTION SET sc_was FOR SERVICE CLASS work1_sc
  USING WORK CLASS SET all_class_types (
    WORK ACTION collect_load_wa ON WORK CLASS load_wc
      COLLECT ACTIVITY DATA ON ALL DATABASE PARTITIONS WITH DETAILS AND VALUES,
    WORK ACTION collect_ddl_wa ON WORK CLASS ddl_wc
      COLLECT ACTIVITY DATA ON ALL DATABASE PARTITIONS WITH DETAILS AND VALUES,
    WORK ACTION collect_read_wa ON WORK CLASS read_wc
      COLLECT ACTIVITY DATA ON ALL DATABASE PARTITIONS WITH DETAILS AND VALUES,
    WORK ACTION map_read_wa on WORK CLASS read_wc
      MAP ACTIVITY TO work1_sc_read)
```

Step 9: Create a service class threshold

To get an effect similar to the stop_large_read_wa work action that prevented any large SELECT statements from running, create an ESTIMATEDSQLCOST threshold and apply it to the work1_sc_read service subclass.

```
CREATE THRESHOLD stop_large_activities FOR SERVICE CLASS work1_sc_read
  UNDER work1_sc
  ACTIVITIES ENFORCEMENT DATABASE
  WHEN ESTIMATEDSQLCOST >10000 STOP EXECUTION
```

Step 10: Clear the activity tables, reset the statistics, and run activities

Clear out all of the activity tables so that you can start afresh before running the script again. Then call the wlm_collect_stats() stored procedure to reset the statistics

```
DELETE FROM activity_db2activities
DELETE FROM activitystmt_db2activities
DELETE FROM activityvals_db2activities
```

```
CALL wlm_collect_stats()
```

Now, run work1.db2 and work3.db2 scripts once.

```
db2 -o -tvf work1.db2
db2 -o -tvf work3.db2
```

Note the SQL04712 error for activities that caused the threshold to be exceeded.

Step 11: View work action set statistics

Use the WLM_GET_WORK_ACTION_SET_STATS table function to access the work action set statistics in memory to get the number of times specific activity

types have been run. Note that running the following query shows only the load_wc, ddl_wc, and the read_wc work classes since they are the only three work classes that have a work action applied to them. All the other activities end up being counted under "*":

```
CONNECT TO SAMPLE

SELECT SUBSTR(WORK_ACTION_SET_NAME, 1, 12) AS WORK_ACTION_SET_NAME,
       SUBSTR(CHAR(DBPARTITIONNUM), 1, 4) AS PART,
       SUBSTR(WORK_CLASS_NAME, 1, 12) AS WORK_CLASS_NAME, LAST_RESET,
       SUBSTR(CHAR(ACT_TOTAL), 1, 10) AS TOTAL_ACTS
FROM TABLE(WLM_GET_WORK_ACTION_SET_STATS(' ', -2)) AS WASSTATS
ORDER BY WORK_ACTION_SET_NAME, WORK_CLASS_NAME, PART
```

This time, output will look something such as the following:

WORK_ACTION_SET_NAME	PART	WORK_CLASS_NAME	LAST_RESET	TOTAL_ACTS
SC_WAS	0	*	2007-08-15-19.02.54.597999	12
SC_WAS	0	DDL_WC	2007-08-15-19.02.54.597999	12
SC_WAS	0	LOAD_WC	2007-08-15-19.02.54.597999	1
SC_WAS	0	READ_WC	2007-08-15-19.02.54.597999	12

4 record(s) selected.

Step 12: View the activity data collected

Now query the activity tables again to get information about the individual activities. Note the service subclass that the activities were run under.

```
SELECT VARCHAR(A.APPL_NAME, 15) as APPL_NAME,
       VARCHAR(A.TPMON_CLIENT_APP, 20) AS CLIENT_APP_NAME,
       VARCHAR(A.APPL_ID, 30) as APPL_ID,
       VARCHAR(A.SERVICE_SUPERCLASS_NAME, 20) as SUPER_CLASS,
       VARCHAR(A.SERVICE_SUBCLASS_NAME, 20) as SUB_CLASS,
       SQLCODE,
       VARCHAR(S.STMT_TEXT, 300) AS STMT_TEXT
FROM ACTIVITY_DB2ACTIVITIES AS A, ACTIVITYSTMT_DB2ACTIVITIES AS S
WHERE A.APPL_ID = S.APPL_ID AND
      A.ACTIVITY_ID = S.ACTIVITY_ID AND
      A.UOW_ID = S.UOW_ID
```

The output will look something like:

APPL_NAME	CLIENT_APP_NAME	APPL_ID	SUPER_CLASS
SUB_CLASS	SQLCODE	STMT_TEXT	
db2bp	CLP work1.db2	*LOCAL.karenam.070815195555	WORK1_SC
WORK1_SC_READ		0 values(current client_applname)	
db2bp	CLP work1.db2	*LOCAL.karenam.070815195555	WORK1_SC
WORK1_SC_READ		0 select * from org	
:			
:			
db2bp	CLP work1.db2	*LOCAL.karenam.070815195555	WORK1_SC
SYSDEFAULTSUBCLASS		0 drop procedure stp2	
db2bp	CLP work3.db2	*LOCAL.karenam.070815195600	WORK1_SC
WORK1_SC_READ		-4712 select count(*) from syscat.tables, syscat.tables, syscat.tables, syscat.tables, syscat.tables	

Note again, that one of the activities has a SQLCODE of -4712, this time because of the stop_large_activities service class threshold that was created in step 9 that was violated because the estimated cost for that select statement was too large. Also notice that all read activities are being run under the work1_sc_read service subclass.

Step 13: Reset for next exercise

Disable the even monitor, drop the service class threshold and drop the service class work action set.

```
SET EVENT MONITOR DB2ACTIVITIES STATE 0

DROP THRESHOLD STOP_LARGE_ACTIVITIES
ALTER WORK ACTION SET SC_WAS
ALTER WORK ACTION COLLECT_LOAD_WA DISABLE
ALTER WORK ACTION COLLECT_DDL_WA DISABLE
ALTER WORK ACTION COLLECT_READ_WA DISABLE
ALTER WORK ACTION MAP_READ_WA DISABLE;
DROP WORK ACTION SET SC_WAS
```

Clear out all of the activity tables so that you can start afresh, before running the script again.

```
DELETE FROM activity_db2activities
DELETE FROM activitystmt_db2activities
DELETE FROM activityvals_db2activities
```

Disable all of the workloads that have been created so that all activities will run under the default user workload and get mapped to the default service super class.

```
ALTER WORKLOAD work1_w1 DISABLE
ALTER WORKLOAD work2_w1 DISABLE
ALTER WORKLOAD work3_w1 DISABLE
ALTER WORKLOAD workth_w1 DISABLE
```

Call the wlm_collect_stats() stored procedure to reset the statistics.

```
CALL WLM_COLLECT_STATS()
```

Exercise 5: Using histograms for service classes

This exercise demonstrates how to use the COLLECT AGGREGATE ACTIVITY DATA BASE option on a service class to produce histograms of coordinator activity lifetimes, coordinator activity execution times, and coordinator activity queue times.

Estimated time: 25-30 minutes

These three histograms are useful for knowing more than just the average lifetime, execution time, or queue time of the activities run on the system, since they can be used to calculate standard deviations and can reveal outliers. For more information on histograms, see "Histograms in workload management" on page 180.

Histograms are accessed through the statistics event monitor. This exercise reuses the statistics event monitor created in Exercise 1 Step 1.

Additional Information: The statistics event monitor is a write-to-table event monitor and contains logical data groups. The first is the control logical data

group, which every event monitor has, and then there are the logical data groups that are specific to the statistics event monitor type. The specific logical data groups are:

- histogrambin for histogram information
- qstats for threshold queue statistics
- scstats service class statistics
- wcstats for work class statistics
- wlstats for workload statistics

Step 1: Create views for viewing histogram statistics

Create several views to make querying the HISTOGRAMBIN_DB2STATISTICS table easier. The first view lists all of the histogram types available. This exercise reports just the three basic types: lifetime, execution time and queue time.

```
CREATE VIEW HISTOGRAMTYPES AS
  SELECT DISTINCT SUBSTR(HISTOGRAM_TYPE,1,24) AS HISTOGRAM_TYPE
  FROM HISTOGRAMBIN_DB2STATISTICS
```

A second view makes it easier to find out which service classes are having histograms collected for them. The HISTOGRAMBIN_DB2STATISTICS table identifies the service classes for which histograms are being collected using the service class ID. Joining this table with the SERVICECLASSES catalog table permits the service class information to be presented with the service super class name and service subclass name instead of the service class ID.

```
CREATE VIEW HISTOGRAMSERVICECLASSES AS
  SELECT DISTINCT SUBSTR(HISTOGRAM_TYPE,1,24) AS HISTOGRAM_TYPE,
    SUBSTR(PARENTSERVICECLASSNAME,1,24) AS SERVICE_SUPERCLASS,
    SUBSTR(SERVICECLASSNAME,1,24) AS SERVICE_SUBCLASS
  FROM HISTOGRAMBIN_DB2STATISTICS AS H,
    SYSCAT.SERVICECLASSES AS S
  WHERE H.SERVICE_CLASS_ID = S.SERVICECLASSID
```

The third view lists all of the times that a histogram of a given type was collected for a given service class. Such as the histogramserviceclasses view, it also joins the HISTOGRAMBIN_DB2STATISTICS table with the SERVICECLASSES catalog table. The difference is that it includes the STATISTICS_TIMESTAMP column as one of the columns in the view.

```
CREATE VIEW HISTOGRAMTIMES AS
  SELECT DISTINCT SUBSTR(HISTOGRAM_TYPE,1,24) AS HISTOGRAM_TYPE,
    SUBSTR(PARENTSERVICECLASSNAME,1,24) AS SERVICE_SUPERCLASS,
    SUBSTR(SERVICECLASSNAME,1,24) AS SERVICE_SUBCLASS,
    STATISTICS_TIMESTAMP AS TIMESTAMP
  FROM HISTOGRAMBIN_DB2STATISTICS AS H,
    SYSCAT.SERVICECLASSES AS S
  WHERE H.SERVICE_CLASS_ID = S.SERVICECLASSID
```

The fourth and final view will be used to show the histograms themselves. It also demonstrates something that one often needs to do when dealing with histograms, which is to aggregate them over time. This view shows the top of each bin and the number of activities that were counted towards each bin. For the three histograms in this exercise, the BIN_TOP field measures the number of milliseconds in the activity lifetime, execution time or queue time. When BIN_TOP is, say 3000 milliseconds and the BIN_TOP of the previous bin is 2000 milliseconds and the NUMBER_IN_BIN is ten for a lifetime histogram, you know that ten activities had a lifetime that was between 2 and 3 seconds.

```

CREATE VIEW HISTOGRAMS(HISTOGRAM_TYPE,
                      SERVICE_SUPERCLASS,
                      SERVICE_SUBCLASS,
                      BIN_TOP,
                      NUMBER_IN_BIN) AS
SELECT DISTINCT SUBSTR(HISTOGRAM_TYPE,1,24) AS HISTOGRAM_TYPE,
                SUBSTR(PARENTSERVICECLASSNAME,1,24) AS SERVICE_SUPERCLASS,
                SUBSTR(SERVICECLASSNAME,1,24) AS SERVICE_SUBCLASS,
                TOP AS BIN_TOP,
                SUM(NUMBER_IN_BIN) AS NUMBER_IN_BIN
FROM HISTOGRAMBIN_DB2STATISTICS AS H,
     SYSCAT.SERVICECLASSES AS S
WHERE H.SERVICE_CLASS_ID = S.SERVICECLASSID
GROUP BY HISTOGRAM_TYPE, PARENTSERVICECLASSNAME, SERVICECLASSNAME, TOP

```

Step 2: Turn on the collection of histograms

The activity lifetime, queue time, and execution time histograms are collected for a service subclass when the base collect aggregate activity data option is enabled for the subclass. Enable the base aggregate activity data collection for the default subclass under the default user super class using the COLLECT AGGREGATE ACTIVITY DATA clause.

Note that all activities will be run in the default user service class since all the user defined workloads were disabled at the end of the previous exercise.

```

ALTER SERVICE CLASS SYSDEFAULTSUBCLASS
  UNDER SYSDEFAULTUSERCLASS
  COLLECT AGGREGATE ACTIVITY DATA BASE

```

Step 3: Activate the statistics event monitor

Activate the statistics event monitor that was created earlier so that it may receive the aggregate data whenever it is collected.

```

SET EVENT MONITOR DB2STATISTICS STATE 1

```

Step 4: Run activities and send statistics to the statistics event monitor

Now some activities can be run. After the activities have finished, the WLM_COLLECT_STATS stored procedure is called to send the statistics (including the activity lifetime, execution time and queue time histograms for the default user service class) to the active statistics event monitor. These histograms contain data about all activities that executed in the default user service class since aggregate activity statistics were enabled. Calling this stored procedure also resets the statistics. To show changes in database activity over time, three collection intervals are created. In the first interval, run two scripts, work1.db2 and work2.db2, and then collect and reset the statistics.

```

db2 -o -tvf work1.db2
db2 -o -tvf work2.db2

```

```

CONNECT TO SAMPLE

```

```

CALL WLM_COLLECT_STATS()

```

In the second interval, only run the work1.db2 script once and then collect and reset the statistics.

```
db2 -o -tvf work1.db2

CONNECT TO SAMPLE

CALL WLM_COLLECT_STATS()
```

In the third interval, run work1.db2 twice and run work2.db2 script once and then collect and reset the statistics.

```
db2 -o -tvf work1.db2
db2 -o -tvf work2.db2
db2 -o -tvf work1.db2

CONNECT TO SAMPLE

CALL WLM_COLLECT_STATS()
```

Collecting data periodically such as this permits you to watch how work on your system changes over time.

Additional Information: Collecting data periodically does not need to be a manual operation. Using the WLM_COLLECT_INT database configuration parameter, one can set the interval in minutes after which statistics collection and reset automatically occurs.

Step 5: Query views to view statistics

Now that statistics have been collected, the views created earlier can be used to look at the statistics. The HISTOGRAMTYPES view just returns the types of histograms available.

```
SELECT * FROM HISTOGRAMTYPES
```

```
HISTOGRAM_TYPE
-----
CoordActExecTime
CoordActLifetime
CoordActQueueTime
```

3 record(s) selected.

Since the BASE option was used when altering the service class, there are three histograms: lifetime, exectime and queuetime. The HISTOGRAMSERVICECLASSES view permits you to see the service classes for which a histogram was collected. The example below restricts the output to that of the CoordActLifetime histogram only. Since aggregate activity collection was only turned on for the default user service class's default subclass, only that class is shown when selecting from the HISTOGRAMSERVICECLASSES view.

```
SELECT * FROM HISTOGRAMSERVICECLASSES
WHERE HISTOGRAM_TYPE = 'CoordActLifetime'
ORDER BY SERVICE_SUPERCLASS, SERVICE_SUBCLASS
```

```
HISTOGRAM_TYPE      SERVICE_SUPERCLASS      SERVICE_SUBCLASS
-----
-
CoordActLifetime    SYSDEFAULTUSERCLASS     SYSDEFAULTSUBCLASS
```

1 record(s) selected.

The HISTOGRAMTIMES view shows the times when histograms were collected. Since the WLM_COLLECT_STATS procedure was run three times, there are three timestamps for the lifetime histogram shown.

```

SELECT * FROM HISTOGRAMTIMES
  WHERE HISTOGRAM_TYPE = 'CoordActLifetime'
        AND SERVICE_SUPERCLASS = 'SYSDEFAULTUSERCLASS'
        AND SERVICE_SUBCLASS = 'SYSDEFAULTSUBCLASS'
  ORDER BY TIMESTAMP

```

HISTOGRAM_TYPE	SERVICE_SUPERCLASS	SERVICE_SUBCLASS	TIMESTAMP
CoordActLifetime	SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	2007-08-05-20.44.51.519380
CoordActLifetime	SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	2007-08-05-21.04.27.131281
CoordActLifetime	SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	2007-08-05-21.08.27.474168

3 record(s) selected.

The last view, HISTOGRAMS, is for looking at the histograms themselves. Unlike the HISTOGRAMTIMES view that lists each collection interval as its own row, this view aggregates histogram data across multiple intervals to produce a single histogram of a given type for a given service class.

```

SELECT BIN_TOP, NUMBER_IN_BIN FROM HISTOGRAMS
  WHERE HISTOGRAM_TYPE = 'CoordActLifetime'
        AND SERVICE_SUPERCLASS = 'SYSDEFAULTUSERCLASS'
        AND SERVICE_SUBCLASS = 'SYSDEFAULTSUBCLASS'
  ORDER BY BIN_TOP

```


BIN_TOP	NUMBER_IN_BIN
-1	0
1	88
2	0
3	0
5	2
8	6
12	7
19	8
29	12
44	13
68	23
103	11
158	2
241	5
369	0
562	0
858	0
1309	0
1997	0
3046	0
4647	0
7089	0
10813	0
16493	0
25157	0
38373	0
58532	0
89280	0
136181	0
207720	0
316840	0
483283	0
737162	0
1124409	0
1715085	0
2616055	0

3990325	0
6086529	0
9283913	0
14160950	0
21600000	0

41 record(s) selected.

The output from the histograms can then be used as input into a graphing tool to generate a graph. The diagram below shows a graph that was created using a Ruby Graphing Library called Gruff Graphs.

Lifetime histogram for SYSDEFAULTUSERCLASS (CoordActLifetime):



Running the query above should produce output that will not be exactly the same as what is shown above since activity lifetimes depend on the performance of the system. In the output above, there are 41 bins and all of the largest bins are empty. At the top, there is a bin whose BIN_TOP is -1. This bin represents all of those activities whose lifetime was too large to fit in the histogram. Seeing a NUMBER_OF_BIN greater than zero when the BIN_TOP is -1 indicates that you should probably increase the high bin value of your histogram. In the output above, the NUMBER_IN_BIN is 0, so there is no need to make such a change. A large number of activities, 88 in this case, were counted in the bin with a BIN_TOP of 1. This is the lowest bin and it means that 88 activities had a lifetime between 0 and 1 milliseconds. Another piece of information that can be extracted from the histogram is that, since the largest BIN_TOP for which there is a corresponding non-zero NUMBER_IN_BIN is 241, the largest lifetime of any activity in the workloads collected in this histogram was between 158 milliseconds and 241 milliseconds. The COORD_ACT_LIFETIME_TOP column in the SCSTATS_DB2STATISTICS table gives a more precise measurement of the lifetime of the activity with the largest lifetime.

The same query can be repeated with a histogram_type of CoordActExecTime instead of CoordActLifetime. The execution time histogram is expected to be similar but not identical to the lifetime histogram. The reason they are different, even when there is no queuing, is that execution time does not include initialization time or cursor idle time, while lifetime does.

```


SELECT BIN_TOP, NUMBER_IN_BIN FROM HISTOGRAMS
  WHERE HISTOGRAM_TYPE = 'CoordActExecTime'
     AND SERVICE_SUPERCLASS = 'SYSDEFAULTUSERCLASS'
     AND SERVICE_SUBCLASS = 'SYSDEFAULTSUBCLASS'
  ORDER BY BIN_TOP

```

BIN_TOP	NUMBER_IN_BIN
-1	0
1	112
2	0
3	0
5	0
8	5
12	7
19	5
29	12
44	7
68	11
103	11
158	2
241	5
369	0
562	0
858	0
1309	0
1997	0
3046	0
4647	0
7089	0
10813	0
16493	0
25157	0
38373	0
58532	0
89280	0
136181	0
207720	0
316840	0
483283	0
737162	0
1124409	0
1715085	0
2616055	0
3990325	0
6086529	0
9283913	0
14160950	0
21600000	0

41 record(s) selected.

Execution time histogram for SYSDEFAULTUSERCLASS (CoordActExecTime):



Once again, a large number of activities are counted in the first bin and the highest execution time of any activity is at most 241 milliseconds.

Finally, the HISTOGRAMS view will be used to look at the CoordActQueueTime histogram. This is the simplest histogram because there is no queuing, since no queuing thresholds were created or enabled in this exercise.

```

SELECT BIN_TOP, NUMBER_IN_BIN FROM HISTOGRAMS
  WHERE HISTOGRAM_TYPE = 'CoordActQueueTime'
    AND SERVICE_SUPERCLASS = 'SYSDEFAULTUSERCLASS'
    AND SERVICE_SUBCLASS = 'SYSDEFAULTSUBCLASS'
 ORDER BY BIN_TOP


```

BIN_TOP	NUMBER_IN_BIN
-1	0
1	177
2	0
3	0
5	0
8	0
12	0
19	0
29	0
44	0
68	0
103	0
158	0
241	0
369	0
562	0
858	0
1309	0
1997	0
3046	0
4647	0
7089	0
10813	0
16493	0
25157	0
38373	0
58532	0
89280	0
136181	0

207720	0
316840	0
483283	0
737162	0
1124409	0
1715085	0
2616055	0
3990325	0
6086529	0
9283913	0
14160950	0
21600000	0

41 record(s) selected.

Queue time histogram for SYSDEFAULTUSERCLASS (CoordActQueueTime):



Every activity was counted in the 0 to 1 millisecond bin because every activity spent zero milliseconds queuing.

The last several queries looked at activity lifetimes, execution times and queue times broken down into bins but aggregated across multiple intervals. The following query presents the same information from a different perspective. It shows averages instead of histograms and, rather than combining the intervals, it shows each interval individually. It also reports a count of the number of completed activities which shows how many activities completed in each interval. It uses the SCSTATS_DB2STATISTICS table instead of the HISTOGRAMBIN_DB2STATISTICS table.

```

SELECT STATISTICS_TIMESTAMP,
       COORD_ACT_LIFETIME_AVG AS LIFETIMEAVG,
       COORD_ACT_EXEC_TIME_AVG AS EXECTIMEAVG,
       COORD_ACT_QUEUE_TIME_AVG AS QUEUETIMEAVG,
       COORD_ACT_COMPLETED_TOTAL AS COMPLETED_TOTAL
FROM SCSTATS_DB2STATISTICS
WHERE SERVICE_SUPERCLASS_NAME = 'SYSDEFAULTUSERCLASS'
  AND SERVICE_SUBCLASS_NAME = 'SYSDEFAULTSUBCLASS'
ORDER BY STATISTICS_TIMESTAMP

STATISTICS_TIMESTAMP      LIFETIMEAVG EXECTIMEAVG QUEUETIMEAVG
COMPLETED_TOTAL
-----
2007-08-07-14.07.44.511153      508          475          0

```

```

77
2007-08-07-14.07.46.537777          513          508          0
39
2007-08-07-14.07.51.882173          314          253          0
113

```

3 record(s) selected.

The result shows that average lifetimes are slightly higher than average execution times for each interval and all three are just over a half a second or less. The average queue time, as expected, is zero. The counts of the number of completed activities in each interval is as expected because workloads 1 and 2 were run in the first interval which resulted in 77 activities collected, workload 1 ran alone in the second interval which resulted in 39 activities, and workload 1 ran twice and workload 2 ran once in the third interval, which resulted in 113 activities.

Step 6: Reset for the next exercise

The final step is to turn off collection of aggregate activities on the default user service class and drop the views and delete the information in the statistics tables.

```

ALTER SERVICE CLASS SYSDEFAULTSUBCLASS
  UNDER SYSDEFAULTUSERCLASS
  COLLECT AGGREGATE ACTIVITY DATA NONE

```

```

DROP VIEW histograms
DROP VIEW histogramtimes
DROP VIEW histogramserviceclasses
DROP VIEW histogramtypes

```

```

SET EVENT MONITOR DB2STATISTICS STATE 0

```

```

DELETE FROM HISTOGRAMBIN_DB2STATISTICS
DELETE FROM SCSTATS_DB2STATISTICS

```

Exercise 6: Investigating delays with WLM table functions

This exercise demonstrates how you can determine the cause of an application slow down with the DB2 WLM monitoring facilities.

Estimated time: 10-15 minutes

The DB2 WLM monitoring facilities provide information and statistics for work in a database. Once the cause of a slow-down is identified, you can remedy the situation.

Step 1: Run activities

Two applications are used in this exercise, app1.db2 and app2.db2. Both applications perform DML operations on the SAMPLE database. Run the app1.db2 script in one window followed immediately by the app2.db2 script in a second window.

```

db2 -tvf app1.db2
db2 -tvf app2.db2

```

Step 2: View currently active workload occurrences

The app2.db2 script should now be hanging. From a third window, issue table function WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97 to find the states of all applications running on the database. For this example, you can

think of a workload occurrence as the same as an application. This table function shows information for all workload occurrences in a service class. Since we want to see all workload occurrences in the database, we use wildcards represented by "*" as service_superclass_name and service_subclass_name input parameters.

```
CONNECT TO SAMPLE
```

```
SELECT INTEGER(APPLICATION_HANDLE) APPL_HANDLE,
       VARCHAR(CLIENT_APPLNAME, 15) AS APPL_NAME,
       VARCHAR(SYSTEM_AUTH_ID, 20) AS USER_ID
FROM TABLE
(WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97('*', '*', -2))
```

The output will look something such as the following:

APPL_HANDLE	APPL_NAME	USER_ID
12	CLP app1.db2	DB2USR1
17	CLP app2.db2	DB2USR1
18	-	DB2USR1
19	-	DB2USR1

4 record(s) selected.

From the output, we can tell that the application handle for app2.db2 is 17.

Step 3: Find the agent for the application

To find out what the agents for app2.db2 are doing use the

WLM_GET_SERVICE_CLASS_AGENTS_V97 table function. This table function shows information on agents working in a service class. Since we want to see the agents working for application handle 17, we specify this in the application_handle input parameter. For this example, we are not interested in agents for a particular service class, so we specify wildcards for the service_superclass_name and service_subclass_name input parameters.

```
SELECT INTEGER(APPLICATION_HANDLE) AS APPL_HANDLE,
       UOW_ID, ACTIVITY_ID,
       VARCHAR(AGENT_TYPE, 15) AS AGENT_TYPE,
       VARCHAR(AGENT_STATE, 10) AS AGENT_STATE,
       VARCHAR(EVENT_TYPE, 10) AS EVENT_TYPE,
       VARCHAR(EVENT_OBJECT, 10) AS EVENT_OBJ,
       VARCHAR(EVENT_STATE, 10) AS EVENT_STATE
FROM TABLE
(WLM_GET_SERVICE_CLASS_AGENTS_V97('*', '*', 17, -2))
```

The output will look something like

APPL_HANDLE	UOW_ID	ACTIVITY_ID	AGENT_TYPE	AGENT_STATE	EVENT_TYPE
17	1	2	COORDINATOR	ACTIVE	ACQUIRE
LOCK	IDLE				

1 record(s) selected.

From the output, you can see that the coordinator agent for application 17 is idle and waiting to acquire a lock. This is the reason why app2.db2 appears to be hanging.

Step 4: Find the problem application and resolve the problem

Now that we know why the application is hanging, we can remedy the situation. We know the application is waiting on a lock. To find out which lock this application is waiting on and which application is holding the lock, we can use the db2pd tool. First, we need to find out the current transaction number for our hanging application: Issue db2pd -transactions for application handle 17.

```
db2pd -db sample -transactions app=17
```

The output will look something such as the following:

Address	AppHandl [nod-index]	TranHdl	Locks	State	
Tflag	Tflag2	Firstlsn	Lastlsn	LogSpace	
SpaceReserved	TID	AxRegCnt	GXID		
0x07000000302A7080	17	[000-00017]	7	5	READ
0x00000000	0x00000000	0x000000000000	0x000000000000	0	
0	0x00000000AC3	1	0		

From the output, we can tell that application 17 has transaction handle 7. We can now find which locks this transaction is waiting on by issuing the db2pd -locks command for transaction handle 7.

```
db2pd -db sample -locks 7 wait
```

The output will look something such as the following:

Address	TranHdl	Lockname	Type	Mode	Sts
Owner	Dur	HoldCount	Att	ReleaseFlg	
0x07000000304013F0	7		0002001000000000640002D52	Row	.NS W
2	1	0	0x00 0x00000002		

The output shows that the application is waiting on a row lock. The owner of the lock has transaction handle 2. This transaction is holding the lock and causing our hang. The final step is to determine the corresponding application handle for transaction handle 2. Issue db2pd -transactions command for transaction handle 2.

```
db2pd -db sample -transactions 2
```

The output will look something such as the following:

Address	AppHandl [nod-index]	TranHdl	Locks	State	
Tflag	Tflag2	Firstlsn	Lastlsn	LogSpace	
SpaceReserved	TID	AxRegCnt	GXID		
0x07000000302A2080	12	[000-00012]	2	6	WRITE
0x00000000	0x00000000	0x000002EE000C	0x000002EE005E	232	
396	0x00000000ABB	1	0		

From the output, we can see that transaction handle 2 corresponds to application handle 12. Referring back to the results from table function WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97, you can see that application 12 refers to app1.db2. This application is holding a row lock that is needed by app2.db2. To make app2.db2 proceed, you may commit, rollback or terminate the unit of work or process from the window running app1.db2. Alternatively, you may also force off app1.db2 by issuing FORCE APPLICATION on application handle 12.

```
db2 force application (12)
```

Additional Information: Another way to diagnose hanging applications due to lock contention is to use the SNAPSHOT_LOCKWAIT monitor table function. This table function provides information on lock holders and waiters. To use this table function, the DFT_MON_LOCK monitor switch configuration parameter must be

turned on before the database is started. This switch affects all databases on an instance.

Exercise 7: Cancelling an ongoing activity

This exercise demonstrates how to cancel an activity that is currently active using the `WLM_CANCEL_ACTIVITY` procedure.

Estimated time: 5-10 minutes

Step 1: Issue a long running query

From a CLP window, run the following script that issues a long running query

```
db2 -tvf longquery.db2
```

Step 2: Get the application handle

From another CLP window, call the `WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97` to get the application handle, unit of work ID and activity ID of the cursor activity above.

```
SELECT T.APPLICATION_HANDLE, T.UOW_ID, T.ACTIVITY_ID, T.ACTIVITY_TYPE
FROM SYSIBMADM.APPLICATIONS A,
     TABLE(WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97
            (CAST(NULL AS BIGINT), -2)) T
WHERE (A.AGENT_ID = T.APPLICATION_HANDLE) AND
      (A.COORD_NODE_NUM = T.COORD_PARTITION_NUM) AND
      (A.DBPARTITIONNUM = T.DBPARTITIONNUM) AND
      (T.DBPARTITIONNUM = T.COORD_PARTITION_NUM) AND
      (A.TPMON_CLIENT_APP = 'CLP Longquery.db2')
```

By joining the result of the table function with the `APPLICATIONS` administrative view, we can find the cursor activity that is run from within `longquery.db2`. The output would look something such as the following:

```
APPLICATION_HANDLE  UOW_ID    ACTIVITY_ID  ACTIVITY_TYPE
-----
-----
                267          1           1  READ_DML
```

1 record(s) selected.

Step 3: Cancel the activity

From the same CLP window, call the `WLM_CANCEL_ACTIVITY` stored procedure to cancel the cursor activity above, using the application handle, unit of work ID, and activity ID obtained from the previous step:

```
CONNECT TO SAMPLE

CALL WLM_CANCEL_ACTIVITY (267, 1, 1)

CONNECT RESET
```

Note that in your case, the application handle, unit of work ID, and activity ID will be different.

In the first CLP window, you will see the following output returned by the long running query issued by `longquery.db2`.

```
SQL4725N The activity has been cancelled. SQLSTATE=57014
```

Exercise 8: Discovering what types of activities are running on your system

This exercise demonstrates how you can use the DB2 workload manager monitoring table functions and work action sets to discover what types of activities are running on your system.

Estimated time: 15-20 minutes

You might want to know the number of large activities or load utilities that are being run concurrently on your system, for example. Understanding the types of work being run on the system is important as different types of work will have different resource requirements and impacts on system performance.

Step 1: Determining the number of activities of each type that are running on your system

Before starting, you might want to show the number of activities of a certain type that are currently running by using the `WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97` table function:

```
CONNECT TO SAMPLE
```

```
SELECT ACTIVITY_TYPE,  
       COUNT(*) AS NUMBER_RUNNING  
FROM TABLE (  
  WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97(CAST(NULL AS BIGINT), -2)) AS T  
GROUP BY ACTIVITY_TYPE
```

The output from this query will look something like:

ACTIVITY_TYPE	NUMBER_RUNNING
-----	-----
READ_DML	1

To get information about the different types of activities that have run on your system over a given period of time, you can use work class sets and work actions.

Step 2: Create a database work action set with count activity work actions

To count the number of times an activity of a specific type has been run over a period of time, a work action set needs to be created. In this example, because we are interested in the activities that are run on the entire system, the work action set will be created at the database level and is associated with the `all_class_types` work class set that was created in Exercise 4 Step 1. This work class set contains work classes for all types of recognized activities. If we were only interested in the activities being run in a specific service class, we would create a work action set at the service class level. For this example, we are also interested in the information for all types of activities so that the work action set contains a `COUNT ACTIVITY` work action for each work class in the `all_class_types` work class set.

```
CREATE WORK ACTION SET work1_was FOR DATABASE  
  USING WORK CLASS SET all_class_types  
  (WORK ACTION count_read_wa ON WORK CLASS read_wc COUNT ACTIVITY,  
   WORK ACTION count_write_wa ON WORK CLASS write_wc COUNT ACTIVITY,  
   WORK ACTION count_ddl_wa ON WORK CLASS ddl_wc COUNT ACTIVITY,
```

```

WORK ACTION count_call_wa ON WORK CLASS call_wc COUNT ACTIVITY,
WORK ACTION count_load_wa ON WORK CLASS load_wc COUNT ACTIVITY,
WORK ACTION count_all_wa ON WORK CLASS all_wc COUNT ACTIVITY)

```

Additional information: Each time an activity corresponding to a work class has one or more work actions applied to it, a counter for the work class is incremented by one. The COUNT ACTIVITY work action provides an efficient way to ensure that the counter is updated. If you do not want to perform any other action on an activity other than counting the number of activities of that type that have been run, the COUNT ACTIVITY work action is the best approach.

Step 3: Run some activities

Run the work1.db2 script once.

```
db2 -tvf work1.db2
```

Step 4: View work action set statistics

You can use the WLM_GET_WORK_ACTION_SET_STATS table function to access the work action set statistics in memory to get the number of times specific activity types have been run. For example, the following query will tell you the number of activities that were assigned to each of the work classes in the work class set that has a work action associated with it:

```

CONNECT TO SAMPLE

SELECT SUBSTR(WORK_ACTION_SET_NAME, 1, 12) AS WORK_ACTION_SET_NAME,
       SUBSTR(WORK_CLASS_NAME, 1, 12) AS WORK_CLASS_NAME,
       LAST_RESET,
       SUBSTR(CHAR(ACT_TOTAL), 1, 12) AS TOTAL_ACTS
FROM TABLE(WLM_GET_WORK_ACTION_SET_STATS(' ', -2)) AS WASSTATS
ORDER BY WORK_ACTION_SET_NAME, WORK_CLASS_NAME, PART

```

Additional Information: The blank included with the statement means that result is not to be restricted by the argument (in this example, we want the information for all of the work action sets). The value of the last argument, dbpartitionnum, is the wildcard character -2, which means that data from all database partitions is to be returned.

The output from this query will look something like the following where "*" represents all activities that do not fall into any of the defined work classes or that fall into work classes with no work actions.

```

WORK_ACTION_SET_NAME WORK_CLASS_NAME LAST_RESET
TOTAL_ACTS
-----
-
WORK1_WAS             *                2007-08-14-13.55.30.725886 0
WORK1_WAS             ALL_WC          2007-08-14-13.55.30.725886 2
WORK1_WAS             CALL_WC        2007-08-14-13.55.30.725886 4
WORK1_WAS             DDL_WC        2007-08-14-13.55.30.725886 12
WORK1_WAS             LOAD_WC       2007-08-14-13.55.30.725886 1
WORK1_WAS             READ_WC       2007-08-14-13.55.30.725886 12
WORK1_WAS             WRITE_WC      2007-08-14-13.55.30.725886 6

```

7 record(s) selected.

Step 5: Differentiate activities by more than their type and other attributes

You can separate out activities by more than just their types. For example, you might want to know how many large queries are being run.

Alter the work class set to add a new read work class that will represent large queries. For this example, a large query is any query that has a cardinality greater than 40.

```
ALTER WORK CLASS SET all_class_types
  ADD WORK CLASS large_wc WORK TYPE READ FOR CARDINALITY FROM 41 POSITION AT 1
```

Additional Information: Note that we positioned this work class at position 1. If the POSITION AT clause is not specified, the work class is positioned at the bottom of the work class set. When deciding which work class an activity belongs to, the work classes are checked in the order they are positioned and the first work class whose attributes match the activities attributes is the class that the activity gets assigned to. In this case, if the large_wc were positioned at the end of the list, the large activities would have been assigned to the read_wc since it was positioned ahead of large_wc.

Alter the work action set to add a COUNT ACTIVITY work action and apply it to the new work class.

```
ALTER WORK ACTION SET work1_was
  ADD WORK ACTION count_large_reads ON WORK CLASS large_wc COUNT ACTIVITY
```

Step 6: Reset the statistics and run some activities

Call the WLM_COLLECT_STATS stored procedure to reset the statistics that are stored in memory so that you are starting fresh and when you chose to query that workload management statistical information that is stored in memory, it will contain information for the activities that have been run from this point on.

```
CALL WLM_COLLECT_STATS()
```

Run the work1.db2 script once.

```
db2 -tvf work1.db2
```

Step 7: View work action set statistics

Use the WLM_GET_WORK_ACTION_SET_STATS table function again to access the work action set statistics in memory to get the number of times specific activity types have been run.

```
CONNECT TO SAMPLE
```

```
SELECT SUBSTR(WORK_ACTION_SET_NAME, 1, 12) AS WORK_ACTION_SET_NAME,
       SUBSTR(CHAR(DBPARTITIONNUM), 1, 4) AS PART,
       SUBSTR(WORK_CLASS_NAME, 1, 12) AS WORK_CLASS_NAME,
       LAST_RESET,
       SUBSTR(CHAR(ACT_TOTAL), 1, 12) AS TOTAL_ACTS
FROM TABLE(WLM_GET_WORK_ACTION_SET_STATS('', -2)) AS WASSTATS
ORDER BY WORK_ACTION_SET_NAME, WORK_CLASS_NAME, PART
```

The output will look something such as the following:

```
WORK_ACTION_SET_NAME PART WORK_CLASS_NAME LAST_RESET
TOTAL_ACTS
-----
```

WORK1_WAS	0	*	2007-08-14-13.55.35.650685	0
WORK1_WAS	0	ALL_WC	2007-08-14-13.55.35.650685	2
WORK1_WAS	0	CALL_WC	2007-08-14-13.55.35.650685	4
WORK1_WAS	0	DDL_WC	2007-08-14-13.55.35.650685	12
WORK1_WAS	0	LARGE_WC	2007-08-14-13.55.35.650685	4
WORK1_WAS	0	LOAD_WC	2007-08-14-13.55.35.650685	1
WORK1_WAS	0	READ_WC	2007-08-14-13.55.35.650685	8
WORK1_WAS	0	WRITE_WC	2007-08-14-13.55.35.650685	6

8 record(s) selected.

Note that this time four of the activities from the script are considered large activities.

Step 8: Reset for the next exercise

Drop the work action set:

```
ALTER WORK ACTION SET WORK1_WAS
ALTER WORK ACTION COUNT_READ_WA DISABLE
ALTER WORK ACTION COUNT_WRITE_WA DISABLE
ALTER WORK ACTION COUNT_DDL_WA DISABLE
ALTER WORK ACTION COUNT_CALL_WA DISABLE
ALTER WORK ACTION COUNT_LOAD_WA DISABLE
ALTER WORK ACTION COUNT_ALL_WA DISABLE
ALTER WORK ACTION COUNT_LARGE_READS DISABLE;
ALTER WORK ACTION SET WORK1_WAS DISABLE;
DROP WORK ACTION SET WORK1_WAS;
```

Exercise 9: Capturing detailed information about an executing activity

This exercise demonstrates how you can use the `WLM_CAPTURE_ACTIVITY_IN_PROGRESS` procedure to capture detailed information about a currently executing activity for later historical analysis.

Estimated time: 5-10 minutes

Activity information you capture is sent to the active event monitor for activities. Previous tasks showed how the `COLLECT ACTIVITY DATA` clause is used for workloads, service classes, work actions and thresholds to capture detailed activity information. This clause needs to be specified in advance before an activity begins executing and information is sent to the activities event monitor when the activity completes. The `WLM_CAPTURE_ACTIVITY_IN_PROGRESS` procedure permits you to capture information reactively when you notice a problem with an activity already in progress. When this procedure is used, information about an activity is sent to the activities event monitor immediately. Both basic and statement activity information are collected, but not input data.

Step 1: Enable activities event monitor

Enable the existing event monitor for activities you created in Exercise 1.

```
CONNECT TO SAMPLE
```

```
SET EVENT MONITOR DB2ACTIVITIES STATE 1
```

Step 2: Issue a long running query

From the CLP, run the following script that issues a long running query with a problematic cursor:

```
db2 -tvf longquery.db2
```


Step 2: Alter the service class to collect activity data

Enable activity collection by specifying the COLLECT ACTIVITY DATA clause on the WLM object of interest. For this exercise, we want to generate historical data for activities run in the default service subclass of the default user service super class:

```
ALTER SERVICE CLASS SYSDEFAULTSUBCLASS UNDER SYSDEFAULTUSERCLASS
    COLLECT ACTIVITY DATA ON COORDINATOR WITH DETAILS
```

Step 3: Enable the activities event monitor

Since the activities event monitor was created in Exercise 1 Step 1, enable it now if it is not enabled already.

```
SET EVENT MONITOR DB2ACTIVITIES STATE 1
```

Step 4: Run some activities

Run some activities so that activity data is collected to generate historical data on.

```
db2 -tvf work1.db2
db2 -tvf work2.db2
```

Step 5: Disable the activity monitor

It is highly recommended that you turn off the event monitor for activities before generating historical data. If you do not do this, any DML activities that are run as a result of the historical data generator may also be captured and put into the DB2 event monitor activity tables, thereby dramatically increasing the number of actual activities for which activity data is generated.

```
CONNECT TO SAMPLE
```

```
SET EVENT MONITOR DB2ACTIVITIES STATE 0
```

Step 6: Generate historical data

Run the historical data generator script, wlmhist.pl, to generate historical data for activities that are captured in the activities event monitor tables. The format is as follows:

```
wlmhist.pl dbname user password [fromTime toTime workloadid
    serviceClassName serviceSubclassName activityTable activityStmtTable]
```

Use a dash (-) to bypass any optional parameters.

Additional Information: The historical data generator (wlmhist.pl) script will generate only historical data for DML. If you have previously run the historical data generator (wlmhist.pl) script once or more, it is recommended that, before running it again, you clear the activityTable and activityStmtTable tables in order to avoid duplicating data. If you choose not to clear these two tables, be sure to use the fromTime and toTime input parameters to ensure you do not generate historical data for activities that have already had data generated for them.

For this exercise, generate historical data for all activities that have been captured in the activities event monitor.

```
Perl wlmhist.pl sample db2inst1 password
```

You may notice some errors similar to the following:

```
Error running explain [IBM][CLI Driver][DB2/LINUX8664] SQL0418N A
statement contains a use of a parameter marker that is not valid. SQLSTATE=42610
for statement VALUES (TABLE_SCHEMA(:H00002 , :H00003 )) INTO :H00007
```

```
DBD::DB2::db do failed: [IBM][CLI Driver][DB2/LINUX8664] SQL0418N A
statement
contains a use of a parameter marker that is not valid. SQLSTATE=42610
```

When generating historical data, explain is run on the actual statement. In some cases, explain cannot be run on some statements with parameter markers and an error is returned. Any activity that shows such an error will not have historical data generated for it.

Once the tool has completed generating historical data, it will tell you how many activities it has successfully generated historical data for.

Step 7: Generate historical data reports

Run the historical data report script `wlmhistrep.pl` to generate reports based on the data that was generated in step 1. The format is as follows:

```
wlmhistrep.pl dbAlias userId passwd [outputFile report schemaName fromTime toTime submitter]
```

Use a dash (-) to bypass optional parameters.

The **report** parameter can be any combination from the following letters:

- A: Tables hit
- B: Tables not hit
- C: Indexes hit
- D: Indexes not hit
- E: Submitters

If the **userId** parameter you specify is not the same as what was used to run the `wlmhist.pl` script when the `wlmhist` table was created, you must specify the correct **schemaName**. The **fromTime** and **toTime** parameters must be specified in timestamp format (for example 2007-06-06-17.00.00).

For this exercise, generate reports for tables hit and indexes not hit:

```
Perl wlmhistrep.pl sample db2inst1 password - AD
```

The output will look something such as the following:

```
TABLES HIT REPORT FOR DATABASE sample
```

TABLE NAME	TABLE SCHEMA	% HITS	TOTAL HITS
EMPLOYEE	KARENAM	7.14285714	2
INVENTORY	KARENAM	14.28571429	4
ORG	KARENAM	28.57142857	8
SALES	KARENAM	14.28571429	4
SYSROUTINES	SYSIBM	7.14285714	2
SYSTABLES	SYSIBM	21.42857143	6
SYSTABLESPACES	SYSIBM	7.14285714	2

```
INDEXES NOT HIT REPORT FOR DATABASE sample
```

TABLE NAME	TABLE SCHEMA	INDEX NAME	INDEX SCHEMA	INDEX TYPE
EXPLAIN_ARGUMENT	KARENAM	ARG_I1	KARENAM	REG
HMON_ATM_INFO	SYSTOOLS	ATM_UNIQ	SYSTOOLS	REG
CUSTOMER	KARENAM	CUST_CID_XMLIDX	KARENAM	XVIL
CUSTOMER	KARENAM	CUST_NAME_XMLIDX	KARENAM	XVIL
CUSTOMER	KARENAM	CUST_PHONES_XMLIDX	KARENAM	XVIL
CUSTOMER	KARENAM	CUST_PHONET_XMLIDX	KARENAM	XVIL
EXPLAIN_DIAGNOSTIC	KARENAM	EXP_DIAG_DAT_I1	KARENAM	REG
HMON_COLLECTION	SYSTOOLS	HI_OBJ_UNIQ	SYSTOOLS	REG
ADVISE_INDEX	KARENAM	IDX_I1	KARENAM	REG
ADVISE_INDEX	KARENAM	IDX_I2	KARENAM	REG
SYSATTRIBUTES	SYSIBM	INDATTRIBUTES01	SYSIBM	REG
SYSATTRIBUTES	SYSIBM	INDATTRIBUTES02	SYSIBM	REG
:				
:				

Step 8: Reset for the next exercise

Disable activity collection for the default service subclass of the default user service super class, and clean up the activity tables.

```
ALTER SERVICE CLASS SYSDEFAULTSUBCLASS UNDER SYSDEFAULTUSERCLASS
COLLECT ACTIVITY DATA NONE
```

```
DELETE FROM ACTIVITY_DB2ACTIVITIES
DELETE FROM ACTIVITYSTMT_DB2ACTIVITIES
```

Exercise 11: Using extended aggregates for service classes

This exercise demonstrates how to use the COLLECT AGGREGATE ACTIVITY DATA EXTENDED option on a service class to produce histograms of coordinator activity inter-arrival times and estimated costs.

Estimated time: 25-30 minutes

The inter-arrival time is the time between the arrival of one activity into the system and the arrival of the next activity. The estimated cost of an activity is the estimated system resources that will be used in the execution of the activity and it only applies to DML activities. An inter-arrival time histogram can be useful for correlating with a lifetime histogram or other lifetime statistics for determining whether a change in lifetime statistics was the result of a change in the arrival rate of the workload versus being the result of a change in the complexity of the workload (more complex queries) or a change in the system. The estimated cost histogram can be useful for correlating with the inter-arrival time and lifetime histograms to see whether a change in the lifetime histogram could be due to a change in the complexity of the workload load (more complex queries with higher estimated costs being submitted), due to a change in the arrival rate of activities (determined from the inter-arrival time distribution) or due to a change in the system itself, such as the introduction of a new threshold, a change in the priority given to a service class, or a change in hardware.

For more information on histograms, see “Histograms in workload management” on page 180.

Histograms are accessed through the statistics event monitor. This exercise reuses the statistics event monitor created in Exercise 1 Step 1.

Step 1: Create views for viewing histogram statistics

Create several views to make querying the HISTOGRAMBIN_DB2STATISTICS table easier. The first view lists all of the histogram types available. In this exercise, it reports just the three basic types: lifetime, execution time and queue time.

```
CREATE VIEW HISTOGRAMTYPES AS
  SELECT DISTINCT SUBSTR(HISTOGRAM_TYPE,1,24) HISTOGRAM_TYPE
  FROM HISTOGRAMBIN_DB2STATISTICS
```

A second view makes it easier to find out what service classes are having histograms collected for them. The HISTOGRAMBIN_DB2STATISTICS table reports the service class for which the histogram is being collected by giving the service class ID. Joining this table with the SERVICECLASSES catalog table will permit the service class information to be presented with the service super class name and service subclass name instead of the service class ID.

```
CREATE VIEW HISTOGRAMSERVICECLASSES AS
  SELECT DISTINCT SUBSTR(HISTOGRAM_TYPE,1,24) HISTOGRAM_TYPE,
    SUBSTR(PARENTSERVICECLASSNAME,1,24) SERVICE_SUPERCLASS,
    SUBSTR(SERVICECLASSNAME,1,24) SERVICE_SUBCLASS
  FROM HISTOGRAMBIN_DB2STATISTICS H,
    SYSCAT.SERVICECLASSES S
  WHERE H.SERVICE_CLASS_ID = S.SERVICECLASSID
```

The third view lists all of the times that a histogram of a given type was collected for a given service class. Such as the HISTOGRAMSERVICECLASSES view, it joins the HISTOGRAMBIN_DB2STATISTICS table with the SERVICECLASSES catalog table. The difference is in the STATISTICS_TIMESTAMP column which is included as one of the columns in this view.

```
CREATE VIEW HISTOGRAMTIMES AS
  SELECT DISTINCT SUBSTR(HISTOGRAM_TYPE,1,24) HISTOGRAM_TYPE,
    SUBSTR(PARENTSERVICECLASSNAME,1,24) SERVICE_SUPERCLASS,
    SUBSTR(SERVICECLASSNAME,1,24) SERVICE_SUBCLASS,
    STATISTICS_TIMESTAMP TIMESTAMP
  FROM HISTOGRAMBIN_DB2STATISTICS H,
    SYSCAT.SERVICECLASSES S
  WHERE H.SERVICE_CLASS_ID = S.SERVICECLASSID
```

The fourth and final view will be used to show the histograms themselves. It also demonstrates a common task when dealing with histograms, which is to aggregate them over time. This view shows the top of each bin and the number of activities that were counted towards each bin. Of the two histograms covered below, the BIN_TOP field measures the number of milliseconds in the activity inter-arrival time and the number of timerons in the estimated cost. When BIN_TOP is, 3000 milliseconds and the BIN_TOP of the previous bin is 2000 milliseconds and the NUMBER_IN_BIN is ten for an inter-arrival time histogram you know that there were ten activities which each arrived into the system between 2 and 3 seconds after the arrival of the previous activity, for example.

```
CREATE VIEW HISTOGRAMS(HISTOGRAM_TYPE, SERVICE_SUPERCLASS,
  SERVICE_SUBCLASS, BIN_TOP, NUMBER_IN_BIN) AS
  SELECT DISTINCT SUBSTR(HISTOGRAM_TYPE,1,24) HISTOGRAM_TYPE,
    SUBSTR(PARENTSERVICECLASSNAME,1,24) SERVICE_SUPERCLASS,
    SUBSTR(SERVICECLASSNAME,1,24) SERVICE_SUBCLASS,
    TOP AS BIN_TOP,
    SUM(NUMBER_IN_BIN) AS NUMBER_IN_BIN
  FROM HISTOGRAMBIN_DB2STATISTICS H,
    SYSCAT.SERVICECLASSES S
  WHERE H.SERVICE_CLASS_ID = S.SERVICECLASSID
  GROUP BY HISTOGRAM_TYPE, PARENTSERVICECLASSNAME, SERVICECLASSNAME, TOP
```


Step 2: Turn on collection of histograms

Turning on the collection of histograms is done for the default user service class by altering its default subclass to collect aggregate activity data with the EXTENDED option. This provides both the three histograms available in the BASE option (lifetime, execution time, and queue time) and the two histograms available only when using the EXTENDED option (inter-arrival time and estimated cost).

```
ALTER SERVICE CLASS SYSDEFAULTSUBCLASS UNDER SYSDEFAULTUSERCLASS
  COLLECT AGGREGATE ACTIVITY DATA EXTENDED
```

Step 3: Activate statistics event monitor

If not already active, activate the event monitor that was created earlier so that it can receive aggregate data whenever it is collected.

```
SET EVENT MONITOR DB2STATISTICS STATE 1
```

Step 4: Run activities and send statistics to statistics event monitor

First run some activities; after the activities have finished the WLM_COLLECT_STATS stored procedure is called to send the service class statistics to the active statistics event monitor (including the activity lifetime, execution time, queue time, inter-arrival time and estimated cost histograms for the default user service class). These histograms contain data about all activities that executed in the default user service class since aggregate activity statistics were enabled. Calling the stored procedure also resets the statistics. To show changes in database activity over time, three collection intervals are created.

In the first interval, run two scripts, work1.db2 and work2.db2, then collect and reset the statistics.

```
db2 -o- -tvf work1.db2
db2 -o- -tvf work2.db2
```

```
CONNECT TO SAMPLE
```

```
CALL WLM_COLLECT_STATS
```

In the second interval, run only the work1.db2 script once, then collect and reset the statistics.

```
db2 -o- -tvf work1.db2
```

```
CONNECT TO SAMPLE
```

```
CALL WLM_COLLECT_STATS
```

In the third interval, run the work1.db2 script twice and the work2.db2 script once, then collect and reset the statistics.

```
db2 -o- -tvf work1.db2
db2 -o- -tvf work2.db2
db2 -o- -tvf work1.db2
```

```
CONNECT TO SAMPLE
```

```
CALL WLM_COLLECT_STATS
```

Collecting data periodically such as this permits you to watch how work on your system changes over time.

Additional Information: Collecting data periodically does not need to be a manual operation. Using the WLM_COLLECT_INT database configuration parameter, you can set the interval in minutes after which statistics collection and reset occurs.

Step 5: Query views to view statistics

Now that statistics have been collected, the views created earlier can be used to look at the statistics. The HISTOGRAMTYPES view just returns the types of histograms available.

```
SELECT * FROM HISTOGRAMTYPES
```

```
HISTOGRAM_TYPE
-----
CoordActEstCost
CoordActExecTime
CoordActInterArrivalTime
CoordActLifetime
CoordActQueueTime
```

3 record(s) selected.

Since the EXTENDED option was used when altering the service class, there are five histograms.

The HISTOGRAMSERVICECLASSES view permits you to see the service classes for which a histogram was collected. The example below restricts the output to that of the CoordActInterArrivalTime histogram only. Since aggregate activity collection was turned on only for the default user service class's default subclass, only that class is shown when selecting from the HISTOGRAMSERVICECLASSES view.

```
SELECT * FROM HISTOGRAMSERVICECLASSES
WHERE HISTOGRAM_TYPE = 'CoordActInterArrivalTime'
ORDER BY SERVICE_SUPERCLASS, SERVICE_SUBCLASS
```

```
HISTOGRAM_TYPE      SERVICE_SUPERCLASS      SERVICE_SUBCLASS
-----
-
CoordActInterArrivalTime SYSDEFAULTUSERCLASS      SYSDEFAULTSUBCLASS
```

1 record(s) selected.

The HISTOGRAMTIMES view shows the times when histograms were collected. Since the WLM_COLLECT_STATS procedure was run three times, there are three timestamps for the inter-arrival time histogram shown.

```
SELECT * FROM HISTOGRAMTIMES
WHERE HISTOGRAM_TYPE = 'CoordActInterArrivalTime'
AND SERVICE_SUPERCLASS = 'SYSDEFAULTUSERCLASS'
AND SERVICE_SUBCLASS = 'SYSDEFAULTSUBCLASS'
ORDER BY TIMESTAMP
```

```
HISTOGRAM_TYPE      SERVICE_SUPERCLASS      SERVICE_SUBCLASS
TIMESTAMP
-----
-
CoordActInterArrivalTime SYSDEFAULTUSERCLASS      SYSDEFAULTSUBCLASS
2007-08-08-13.41.38.870298
CoordActInterArrivalTime SYSDEFAULTUSERCLASS      SYSDEFAULTSUBCLASS
2007-08-08-13.41.42.802855
CoordActInterArrivalTime SYSDEFAULTUSERCLASS      SYSDEFAULTSUBCLASS
2007-08-08-13.41.53.577835
```

The last view, HISTOGRAMS, is for looking at the histograms themselves. Unlike the HISTOGRAMTIMES view that lists each collection interval as its own row, this view aggregates histogram data across multiple intervals to produce a single histogram of a given type for a given service class.

```
SELECT BIN_TOP, NUMBER_IN_BIN FROM HISTOGRAMS
WHERE HISTOGRAM_TYPE = 'CoordActInterArrivalTime'
AND SERVICE_SUPERCLASS = 'SYSDEFAULTUSERCLASS'
AND SERVICE_SUBCLASS = 'SYSDEFAULTSUBCLASS'
ORDER BY BIN_TOP
```

BIN_TOP	NUMBER_IN_BIN
-1	0
1	10
2	6
3	7
5	14
8	7
12	32
19	2
29	9
44	24
68	11
103	8
158	8
241	9
369	1
562	10
858	5
1309	5
1997	0
3046	0
4647	0
7089	0
10813	2
16493	2
25157	0
38373	0
58532	0
89280	0
136181	0
207720	0
316840	0
483283	0
737162	0
1124409	0
1715085	0
2616055	0
3990325	0
6086529	0
9283913	0
14160950	0
21600000	0

41 record(s) selected.

Running this query produces output that will not be exactly the same as what is shown above since activity inter-arrival times depend on the performance of the system. In the output above, there are 41 bins and all of the largest bins are empty. At the top, there is a bin whose BIN_TOP is -1. This bin represents all of those activities whose inter-arrival time was too large to fit in the histogram. Seeing a NUMBER_OF_BIN greater than zero when the BIN_TOP is -1 indicates that you should probably increase the high bin value of your histogram. In the output above, the NUMBER_IN_BIN is 0, so there is no need to make such a change. The

majority of activities arrived less than 1309 milliseconds apart from each other. Four activities arrived between 7089 milliseconds and 16493 milliseconds apart.

The same query can be repeated with a histogram_type of CoordActEstCost instead of CoordActInterArrivalTime.

```
SELECT BIN_TOP, NUMBER_IN_BIN FROM HISTOGRAMS
WHERE HISTOGRAM_TYPE = 'CoordActEstCost'
AND SERVICE_SUPERCLASS = 'SYSDEFAULTUSERCLASS'
AND SERVICE_SUBCLASS = 'SYSDEFAULTSUBCLASS'
ORDER BY BIN_TOP
```

BIN_TOP	NUMBER_IN_BIN
-1	0
1	39
2	0
3	0
5	0
8	30
12	0
19	30
29	0
44	0
68	0
103	0
158	0
241	0
369	0
562	0
858	0
1309	0
1997	0
3046	0
4647	0
7089	0
10813	0
16493	0
25157	0
38373	0
58532	0
89280	0
136181	0
207720	0
316840	0
483283	0
737162	0
1124409	0
1715085	0
2616055	0
3990325	0
6086529	0
9283913	0
14160950	0
21600000	0

41 record(s) selected.

A histogram such as this is typical for a small workload. With a small workload, there is not much variation in the size of activities, so there are only three different bins that had activities counted towards them. Slightly more than 60% of the activities had a cost estimate between 5 and 19 timerons, with the rest having cost estimates of less than 1 timeron.

Step 6: Reset for other exercises

The final step is to turn off collection of aggregate activities on the default user service class, to drop the views and to delete the information in the statistics tables.

```
ALTER SERVICE CLASS SYSDEFAULTSUBCLASS UNDER SYSDEFAULTUSERCLASS
    COLLECT AGGREGATE ACTIVITY DATA NONE
```

```
DROP VIEW HISTOGRAMS
DROP VIEW HISTOGRAMTIMES
DROP VIEW HISTOGRAMSERVICECLASSES
DROP VIEW HISTOGRAMTYPES
```

```
SET EVENT MONITOR DB2STATISTICS STATE 0
```

```
DELETE FROM HISTOGRAMBIN_DB2STATISTICS
DELETE FROM SCSTATS_DB2STATISTICS
```

Chapter 7. Workload management scenarios

Workload management sample application

Comprehensive workload management features have been integrated into your DB2 data server with DB2 workload manager, giving you finer control over activities, resources and performance, and deeper insight into how your system is running. A workload management sample application is now available on developerWorks®.

The workload management sample application demonstrates how you can use DB2 workload manager features to achieve the following objectives:

Protect the system from runaway queries

Runaway queries are costly and cause poor performance. The workload management sample application identifies queries with the potential to become runaway queries, and then stops these queries from running after they have violated a specified threshold.

Limit concurrent resource consumption by individual applications

The sample application shows how to use DB2 workload manager features to prevent applications that submit large amounts of concurrent work from negatively affecting the performance of other applications.

Achieve a specific response time

Workload management features permit you to achieve a specific response time objective of the form: "transaction X from application Y shall complete within 1 second in 90% of cases," regardless of what other activity is running concurrently on the system. The sample application will demonstrate how to achieve a response time objective.

Consistent response time for short queries

Queries that typically have a response time of less than 1 second should have a relatively consistent response time regardless of what other workloads are running on the system. The sample application uses the query execution time histogram to monitor consistency.

Protect the system during periods of peak demand

Workload management policy features protect the system from capacity overload during bursts of peak demand by queuing work once the system is sufficiently loaded.

Enable concurrent batch extract, transform, and load (ETL) processing and user queries

Workload management features permit you to run ETL jobs (like loading data into tables) while controlling the performance impact for users running queries concurrently.

To obtain the sample application, see Workload management sample on developerWorks.

Scenario: Investigating a workload-related system slowdown

If you notice a system slowdown (for example, some applications take much longer to complete than expected) and are unsure whether the problem is related to the configuration of the workloads, you can use table function data to investigate and, if necessary, correct the problem.

First, create a query that aggregates data across service classes and database partitions using data from the WLM_GET_SERVICE_SUBCLASS_STATS_V97 table function. Set the first and second arguments to empty strings and the third argument to -2 (a wildcard character) to indicate that data is to be gathered for all service classes on all database partitions.

Your query might resemble the following one:

```
SELECT SUBSTR(SERVICE_SUPERCLASS_NAME,1,19) AS SUPERCLASS_NAME,
       SUBSTR(SERVICE_SUBCLASS_NAME,1,18) AS SUBCLASS_NAME,
       SUBSTR(CHAR(SUM(COORD_ACT_COMPLETED_TOTAL)),1,13) AS ACTSCOMPLETED,
       SUBSTR(CHAR(SUM(COORD_ACT_ABORTED_TOTAL)),1,11) AS ACTSABORTED,
       SUBSTR(CHAR(MAX(CONCURRENT_ACT_TOP)),1,6) AS ACTSHW,
       CAST(CASE WHEN SUM(COORD_ACT_COMPLETED_TOTAL) = 0 THEN 0
              ELSE SUM(COORD_ACT_COMPLETED_TOTAL * COORD_ACT_LIFETIME_AVG)
              / SUM(COORD_ACT_COMPLETED_TOTAL) END / 1000 AS DECIMAL(9,3))
       AS ACTAVGLIFETIME
FROM TABLE(WLM_GET_SERVICE_SUBCLASS_STATS_V97('', '', -2)) AS SCSTATS
GROUP BY SERVICE_SUPERCLASS_NAME, SERVICE_SUBCLASS_NAME
ORDER BY SUPERCLASS_NAME, SUBCLASS_NAME
```

SUPERCLASS_NAME	SUBCLASS_NAME	ACTSCOMPLETED	ACTSABORTED	ACTSHW	ACTAVGLIFETIME
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	20	0	1	3.750
SUP1	SUB1	40	0	8	7.223

In the preceding example data, the SUB1 service subclass in the SUP1 service superclass is running more simultaneous activities than usual. To investigate further, you might want to examine the statistics for workloads that map to this service class. Your query might resemble the following one:

```
SELECT SUBSTR(WLSTATS.WORKLOAD_NAME,1,22) AS WL_NAME,
       SUBSTR(CHAR(WLSTATS.DBPARTITIONNUM),1,4) AS PART,
       CONCURRENT_WLO_TOP AS WLO_HIGH_WTRMRK,
       CONCURRENT_WLO_ACT_TOP AS WLO_ACT_HIGH_WTRMRK
FROM TABLE(WLM_GET_WORKLOAD_STATS_V97('', -2)) AS WLSTATS,
       TABLE(WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97('', '', -2)) AS SCWLOS
WHERE WLSTATS.WORKLOAD_NAME = SCWLOS.WORKLOAD_NAME
AND SCWLOS.SERVICE_SUPERCLASS_NAME = 'SUP1'
AND SCWLOS.SERVICE_SUBCLASS_NAME = 'SUB1'
ORDER BY WL_NAME, PART;
```

WL_NAME	PART	WLO_HIGH_WTRMRK	WLO_ACT_HIGH_WTRMRK
LYNNSALES	0	2	8
LYNNSALES	1	0	0
SYSDEFAULTUSERWORKLOAD	0	1	1
SYSDEFAULTUSERWORKLOAD	1	0	0

The output shows that an application in the LYNNSALES workload submitted 8 activities concurrently. Consider adding a threshold to restrict concurrency of coordinator activities for each workload occurrence.

Scenario: Identifying activities that are taking too long to complete

Workload management table functions simplify the task of identifying a specific activity inside the data server and, if necessary, canceling it without having to end the entire application.

Identifying an activity that is taking too long to complete

Following is an example of identifying a long-running query. Assume that a user from the Sales department who is running the SalesReport application complains that the application is taking too long to complete.

After identifying the application handle, use the WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97 table function to look up all activities currently running in this application. For example, if the application handle is 1, your query might resemble the following one:

```
SELECT SUBSTR(CHAR(COORD_PARTITION_NUM),1,5) AS COORD,  
SUBSTR(CHAR(DBPARTITIONNUM),1,4) AS PART,  
SUBSTR(CHAR(UOW_ID),1,5) AS UOWID,  
SUBSTR(CHAR(ACTIVITY_ID),1,5) AS ACTID,  
SUBSTR(CHAR(PARENT_UOW_ID),1,8) AS PARUOWID,  
SUBSTR(CHAR(PARENT_ACTIVITY_ID),1,8) AS PARACTID,  
SUBSTR(CHAR(ACTIVITY_TYPE),1,8) AS ACTTYPE,  
SUBSTR(CHAR(NESTING_LEVEL),1,7) AS NESTING  
FROM TABLE(WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97(1, -2)) AS WLOACTS  
ORDER BY PART, UOWID, ACTID
```

COORD	PART	UOWID	ACTID	PARUOWID	PARACTID	ACTTYPE	NESTING
0	0	2	3	-	-	CALL	0
0	0	2	5	2	3	READ_DML	1

The activity is identified as having a unit of work ID of 2 and an activity ID of 5. You can then use the WLM_GET_SERVICE_CLASS_AGENTS_V97 table function to discover what the agents that work on this activity are doing:

```
SELECT APPLICATION_HANDLE,  
UOW_ID,  
ACTIVITY_ID,  
SUBSTR(REQUEST_TYPE,1,8) AS REQUEST_TYPE,  
SUBSTR(EVENT_TYPE,1,8) AS EVENT_TYPE,  
SUBSTR(EVENT_OBJECT,1,8) AS EVENT_OBJECT  
FROM TABLE(WLM_GET_SERVICE_CLASS_AGENTS_V97(' ', ' ',  
CAST(NULL AS BIGINT),  
-2)) AS AGENTS  
WHERE APPLICATION_HANDLE = 1  
AND UOW_ID = 2  
AND ACTIVITY_ID = 5
```

For example, the activity might be queued, executing, or waiting on a lock. If the activity were queued, the result would be:

APPLICATION_HANDLE	UOW_ID	ACTIVITY_ID	REQUEST_TYPE	EVENT_TYPE	EVENT_OBJECT
1	2	5	OPEN	WAIT	WLM_QUEUE

If the activity were executing, the result would be:

APPLICATION_HANDLE	UOW_ID	ACTIVITY_ID	REQUEST_TYPE	EVENT_TYPE	EVENT_OBJECT
1	2	5	OPEN	PROCESS	REQUEST

If the activity were waiting on a lock, the result would be:

APPLICATION_HANDLE	UOW_ID	ACTIVITY_ID	REQUEST_TYPE	EVENT_TYPE	EVENT_OBJECT
	1	2		OPEN	ACQUIRE
		5			LOCK

When you know what the activity is doing, you can proceed appropriately:

- If the activity is queued, if the user indicates that the query is taking so long that they no longer care about the results, or you think the query is consuming too many resources, you can cancel it.
- If the activity is important and it is queued, consider cancelling some other less important work that is currently running (reducing the concurrency so that activities leave queue), or maybe the user will be satisfied to know that work is not hanging and is just waiting for chance to run.
- If the activity is waiting for a lock, you can use the snapshot monitor to investigate which locks the application is waiting for.
- If the activity is waiting for a lock held by lower priority activity, consider cancelling that activity.

You might also find it useful to know the DML statement that activity 5 is running. Assuming that you have an active activities event monitor, you can run the `WLM_CAPTURE_ACTIVITY_IN_PROGRESS` procedure to capture information about the DML statement and other information about activity 5 while it is running. Unlike the statement event monitor, the `WLM_CAPTURE_ACTIVITY_IN_PROGRESS` procedure permits you to capture information about a specific query, as opposed to every statement running at the time. You can also obtain the statement text by using `MON_GET_ACTIVITY_DETAILS`.

If you decide that you must cancel the activity, you can use the `WLM_CANCEL_ACTIVITY` routine to cancel the activity without having to end the application that issued it:

```
CALL WLM_CANCEL_ACTIVITY (1, 2, 5)
```

The application that issued the activity receives an `SQL4725N` error. Any application that handles negative SQL codes is able to handle this SQL code.

Identifying an activity hang caused by lock contention

Assume that you have a situation in which a user is complaining about an application that is taking too long. Also assume that you have either the application name or the authorization ID of the long-running application. With this information, you can use the `LIST APPLICATIONS` command to obtain the application handle. Assuming that application handle returned by the `LIST APPLICATIONS` command is 2, you can use the `WLM_GET_SERVICE_CLASS_AGENTS_V97` table function to determine which agents are working on this activity. Your query might resemble the following one:

```
SELECT SUBSTR(CHAR(APPLICATION_HANDLE),1,7) AS APPHANDLE,
       SUBSTR(CHAR(DBPARTITIONNUM),1,4) AS PART,
       SUBSTR(CHAR(AGENT_TID),1,9) AS AGENT_TID,
       SUBSTR(CHAR(AGENT_TYPE),1,11) AS AGENTTYPE,
       SUBSTR(CHAR(EVENT_OBJECT),1,11) AS EVENTOBJECT,
       SUBSTR(CHAR(REQUEST_TYPE),1,7) AS REQTYPE,
       SUBSTR(CHAR(UOW_ID),1,6) AS UOW_ID,
       SUBSTR(CHAR(ACTIVITY_ID),1,6) AS ACT_ID
FROM TABLE(WLM_GET_SERVICE_CLASS_AGENTS_V97(' ', ' ', 2, -2)) AS SCDETAILS
ORDER BY APPHANDLE, PART, AGENT_TID
```

```
APPHANDLE PART AGENT_TID AGENTTYPE EVENTOBJECT REQTYPE UOW_ID ACT_ID
```

```

-----
2          0  1          COORDINATOR REQUEST  OPEN  2  1
2          1  3          SUBAGENT   LOCK      -  2  1

```

The results indicate that agent 1 is waiting on a remote response. Looking at the agent on the remote partition that is working on the same activity, the EVENTOBJECT field indicates that the agent is waiting to obtain a lock.

The next step is to determine who owns the lock. You can obtain this information by turning on the monitor switches and using the snapshot monitor table function, as shown in the following example:

```

SELECT AGENT_ID AS WAITING_FOR_LOCK,
       SUBSTR(APPL_ID_HOLDING_LK,1,40) AS HOLDING_LOCK,
       CAST(LOCK_MODE_REQUESTED AS SMALLINT) AS WANTED,
       CAST(LOCK_MODE AS SMALLINT) AS HELD
FROM TABLE(SNAPSHOT_LOCKWAIT('SAMPLE',-1)) AS SLW

```

```

WAITING_FOR_LOCK  HOLDING_LOCK  WANTED HELD
-----
                2 *LOCAL.DB2.060131021547          9  5

```

You can also determine the lock owner by using the following sequence of commands:

```

db2pd -db database alias -locks
db2pd -db database alias -transactions

```

If you want to cancel the long-running activity, you can use the WLM_CANCEL_ACTIVITY procedure. If the successful completion of the long-running application is more important than the successful completion of the lock-owning application, you can force the lock-owning application.

Scenario: Tuning a DB2 workload manager configuration when capacity planning data is available

If you performed capacity planning, you should have information about the types of users and their expected response times. You can use this information to construct, determine the effectiveness of, and tune your DB2 workload manager configuration.

Assume that you performed capacity planning and that the data in the following table represents the results of this exercise for work types and response time goals:

Table 57. Results of capacity planning

Type of work	Application	Goal	Importance	Expected throughput
Order entry	orderentryapp.exe	Obtain an average response time < 1 second	High	10 000 (both inserts and updates) per day
Business intelligence queries	businessobjects.exe	Obtain an average response time < 10 seconds	High	100 queries per day
Batch processing	batchapp.exe	Maximize throughput	Low	5000 updates per day
Other	All other applications	Best effort	Low	100 activities per day

Based on the data in the preceding table, you might create three service classes (ORDER_ENTRY_SC, BI_QUERIES_SC, and BATCH_SC) and three workloads

(ORDER_ENTRY_WL, BI_QUERIES_WL, and BATCH_WL) to assign work to the service classes. After creating the service classes and workloads, you might create a statistics event monitor to collect aggregate activity information, such as the activity lifetime histogram for each service class. Assume that the data in the following table compares the average daily count of activities in each service class (computed from the activity lifetime histogram) with the volumes that were predicted in the capacity planning exercise:

Table 58. Activities each day

Service class	Predicted activities per day	Actual activities per day
ORDER_ENTRY_SC	10 000	9700
BI_QUERIES_SC	100	115
BATCH_SC	5000	5412
SYSDEFAULTUSERCLASS	100	85

The observed data indicates that the capacity planning estimates were accurate. The data in the following table compares the average activity lifetimes (obtained from the activity lifetime histogram) with the response time goals determined during capacity planning and shows that activities in the BI_QUERIES_SC service class are not meeting their response time objectives.

Table 59. Response times

Service class	Response time goal	Actual average lifetime
ORDER_ENTRY_SC	< 1 second	0.8 seconds
BI_QUERIES_SC	< 10 seconds	30 seconds
BATCH_SC		2 seconds
SYSDEFAULTUSERCLASS		10 minutes

With DB2 workload manager, you can use different approaches when addressing the problem of the business intelligence queries not meeting their response time goals:

- Limiting the concurrency of lower-importance service classes
- Allowing the operating system workload manager to provide less processor resource to less-important service classes
- Modifying the agent and I/O prefetcher priorities for the service classes
- Using any combination of the previous three approaches

Assume that processor time is the resource that is causing the business intelligence queries to fail to meet their goals. Also assume that you use the operating system workload manager to give the SYSDEFAULTUSERCLASS service class less processor resources than other service classes. You can then capture aggregate activity information over a period of days to observe whether the changes to the CPU allocation provide the results that you expect. The data in the following table shows another comparison between response time goals and actual average lifetimes computed from the histograms after you made the operating system workload manager changes. All service classes are now meeting their response time objectives and, because of the processor time reallocation, activities in the SYSDEFAULTUSERCLASS service class have had their response times doubled.

Table 60. Response times after reconfiguration

Service class	Response time goal	Actual average lifetime
ORDER_ENTRY_SC	< 1 second	0.6 seconds
BI_QUERIES_SC	< 10 seconds	9.5 seconds
BATCH_SC		1.5 seconds
SYSDEFAULTUSERCLASS		20 minutes

Scenario: Tuning a DB2 workload manager configuration when capacity planning information is unavailable

You can use the DB2 workload manager tools to help design, monitor, and tune a workload management configuration even if you do not have capacity analysis data to use for designing the configuration.

Assume that you do not initially know which workloads and service classes to create because either you do not have full knowledge of the workload on the system or you do not yet know which workloads are required for stable execution results. Also assume that you know that some applications have response time requirements but that you do not yet know how many other applications are competing for resources with such time-critical applications. You can use the workload management monitoring capabilities to determine this.

To set up a wDB2 workload manager configuration using monitoring data as the foundation:

1. Classify those applications that you know are important. You must isolate these applications and give them an appropriate portion of the system resources.
2. For the rest of the workload, collect statistics for the largest activities in the workload because these activities have the greatest impact on a per-activity basis on the system.
3. Analyze the activity information that you collected in step 2.
4. Repeat steps 1 through 3 on that portion of the workload that is still unclassified. Repeat this step until you know that the remaining unclassified work is not worth classification.

The sections that follow provide information about how to perform these steps.

Step 1. Isolate those applications that are known to be important and give them an appropriate portion of resources

Assume that you have two important business intelligence applications, BI1 and BI2 and that you need to minimize the response times for these applications. You can create workloads for these two applications and map them to a service class called MOSTIMPORTANT for which you can assign system resources.

On the AIX operating system, you use the AIX Workload Manager to create a service class called MOSTIMPORTANT, and give this service class a guaranteed set of resources.

On the DB2 data server, you create the required service classes and workloads:

```
CREATE SERVICE CLASS MOSTIMPORTANT OUTBOUND CORRELATOR 'MOSTIMPORTANT'
CREATE WORKLOAD BI1WORKLOAD APPLNAME ('BI1') SERVICE CLASS MOSTIMPORTANT
CREATE WORKLOAD BI2WORKLOAD APPLNAME ('BI2') SERVICE CLASS MOSTIMPORTANT
```

For the purposes of this example, assume that even after you account for the known applications, a significant portion of the system workload is unaccounted for. You therefore need to better understand and possibly control this workload.

Step 2. For the remaining unclassified workload, collect statistics for the largest activities in the workload

A long-running activity has a greater individual impact on the system than a short-running activity has because the long-running activity occupies system resources for a longer period of time. However, collecting information about a long-running activity imposes no greater overhead than would be imposed by collecting information on a short-running activity. As a result, the best way to collect information on the largest proportion of the workload is to collect information on the longest-running activities first.

Start collecting activity information by first deciding on an activity lifetime above which you collect activity information. You can simplify this task by choosing a portion of the unclassified activities to be collected, such as 30%, and then observing the activity lifetime histogram for these activities. Allow the system to run so that the in-memory statistics are updated, then run the `WLM_COLLECT_STATS` procedure to send the statistics to an active statistics event monitor.

Use the following query to obtain the activity lifetime histogram for the `SYSDEFAULTUSERCLASS` service class as a table that represents the proportion of the total activities that fell into each lifetime range. This query is written assuming that the database is not partitioned.

```
WITH TOTAL AS (
SELECT PARENTSERVICECLASSNAME,
      SERVICECLASSNAME,
      HIST.HISTOGRAM_TYPE,
      SUM(NUMBER_IN_BIN) AS NUMBER_IN_BIN
FROM HISTOGRAMBIN_DB2STATISTICS AS HIST,
      SYSCAT.SERVICECLASSES SC
WHERE
      HIST.SERVICE_CLASS_ID = SC.SERVICECLASSID
      AND HIST.TOP >= 0
      AND SC.PARENTSERVICECLASSNAME = 'SYSDEFAULTUSERCLASS'
      AND SC.SERVICECLASSNAME = 'SYSDEFAULTSUBCLASS'
      AND HIST.HISTOGRAM_TYPE = 'COORDACTLIFETIME'
GROUP BY PARENTSERVICECLASSNAME, SERVICECLASSNAME, HISTOGRAM_TYPE)
SELECT CAST(CAST(TOP AS DOUBLE) / 60000 AS DECIMAL(14,3)) AS TOP_IN_MINUTES,
      CAST(100 * CAST(SUM(HIST.NUMBER_IN_BIN) AS DOUBLE) / TOTAL.NUMBER_IN_BIN AS DECIMAL(4,2))
      AS PERCENT_IN_BIN
FROM HISTOGRAMBIN_DB2STATISTICS AS HIST,
      SYSCAT.SERVICECLASSES SC,
      TOTAL
WHERE HIST.SERVICE_CLASS_ID = SC.SERVICECLASSID
      AND HIST.TOP >= 0
      AND SC.PARENTSERVICECLASSNAME = 'SYSDEFAULTUSERCLASS'
      AND SC.SERVICECLASSNAME = 'SYSDEFAULTSUBCLASS'
      AND HIST.HISTOGRAM_TYPE = 'COORDACTLIFETIME'
      AND TOTAL.PARENTSERVICECLASSNAME = SC.PARENTSERVICECLASSNAME
      AND TOTAL.SERVICECLASSNAME = SC.SERVICECLASSNAME
      AND TOTAL.HISTOGRAM_TYPE = HIST.HISTOGRAM_TYPE
GROUP BY TOP, SC.PARENTSERVICECLASSNAME, SC.SERVICECLASSNAME, HIST.HISTOGRAM_TYPE, TOTAL.NUMBER_IN_BIN;
```

TOP_IN_MINUTES	PERCENT_IN_BIN
-----	-----
0.000	0.00

0.000	0.00
0.000	0.00
0.000	0.00
0.000	0.00
0.000	0.00
0.000	0.00
0.000	0.00
0.000	0.00
0.001	0.00
0.001	0.00
0.002	0.00
0.004	0.00
0.006	0.00
0.009	0.00
0.014	0.00
0.021	0.00
0.033	0.00
0.050	0.00
0.077	0.00
0.118	0.00
0.180	0.00
0.274	0.00
0.419	0.00
0.639	0.00
0.975	0.00
1.488	0.00
2.269	0.00
3.462	0.00
5.280	0.00
8.054	0.00
12.286	0.00
18.740	0.00
28.584	10.00
43.600	15.00
66.505	45.00
101.442	23.00
154.731	5.00
236.015	2.00
360.000	0.00

The following figure shows the results of the preceding query plotted as a graph:




Figure 32. Activity lifetime histogram of unclassified activities

In this example, 30% of the activities fall into the 101 minutes or greater lifetime range. To capture information about these activities, create an activity lifetime threshold of 100 minutes with the `CONTINUE` and `COLLECT ACTIVITY DATA` options as shown in the following example. If this threshold is violated, activity information is sent to an active activities event monitor.

```
CREATE THRESHOLD COLLECTLONGESTRUNNING30PERCENT
FOR SERVICE CLASS SYSDEFAULTSUBCLASS UNDER SYSDEFAULTUSERCLASS
ACTIVITIES ENFORCEMENT DATABASE ENABLE
WHEN ACTIVITYTOTALTIME > 100 MINUTES COLLECT ACTIVITY DATA CONTINUE
```

Allow the system to run so that data is collected.

Assuming that the overhead of collecting information on 30% of the longest-running activities is acceptable, you can let the data collection continue for a few hours or a few days. You can use the collected data to determine which users and applications produce the longest running of the 30% of the DML activities that are still unclassified. These activities might include some that are time critical. You might uncover some surprises, such as low-priority applications that run significant numbers of large activities. When you finish collecting and analyzing the data, you can drop the threshold.

Step 3. Analyze the information about activities collected in the previous step

You can analyze the information you collected about activities in the previous step according to the application that submitted them. You might specify the following query:

```
SELECT SUBSTR (APPL_NAME, 1,16) APPLICATION_NAME,
       AVG(TIMESTAMPDIFF(4, CHAR(TIME_COMPLETED - TIME_CREATED)))
       AS AVG_LIFETIME_MINUTES
       COUNT(*) AS ACTIVITY_COUNT
FROM ACTIVITY_DB2ACTIVITIES
GROUP BY APPL_NAME
```

```
ORDER BY APPL_NAME
```

```
APPLICATION_NAME AVG_LIFETIME_MINUTES ACTIVITY_COUNT
=====
MOSTLYSMALL1          120             21
MOSTLYSMALL2          110             15
UNIMPORTANTAPP        150            10213
```

An analysis of the activities according to the submitting application shows that a large number of the longest-running activities were submitted by the UNIMPORTANTAPP application, which is a relatively unimportant application. You can use a workload to isolate this application from the other unclassified applications and map it to a service class called BESTEFFORT, which receives resources only when all other activities have their resource needs met.

According to the preceding results, the remaining applications in the default service class appear to submit few large activities. You might find it worthwhile to repeat the process of collecting activities executing in the default service class without restricting the collection to long-running activities.

Step 4. Repeat steps 1 to 3 on that portion of the workload that is still unclassified until the remaining unclassified work is not worth classification

Now that you have the two important applications running in the MOSTIMPORTANT service class and the unimportant application running in the BESTEFFORT service class, much less work is running in the default user service class. In this situation, it might be inexpensive to collect information about every activity in this service class. Alternatively, you might not need to further subdivide the work and can stop here. Assume that you want to collect information about the remaining activities, in case the remaining workload contains surprises. You can accomplish this task by setting COLLECT ACTIVITY DATA for the default user service class and creating an activities event monitor:

```
ALTER SERVICE CLASS SYSDEFAULTSUBCLASS UNDER SYSDEFAULTUSERCLASS
COLLECT ACTIVITY DATA ON COORDINATOR WITHOUT DETAILS
```

Allow the system to run so that data is collected. You can analyze the results as in step 3.

```
SELECT SUBSTR (APPL_NAME,1,16) APPLICATION_NAME,
       AVG(TIMESTAMPDIFF(4, CHAR(TIME_COMPLETED - TIME_CREATED)))
       AS AVG_LIFETIME_MINUTES
       COUNT(*) AS ACTIVITY_COUNT
FROM ACTIVITY_DB2ACTIVITIES
GROUP BY APPL_NAME
ORDER BY APPL_NAME
```

```
APPLICATION_NAME AVG_LIFETIME_MINUTES ACTIVITY_COUNT
=====
MOSTLYSMALL1          5             1501
MOSTLYSMALL2          7             124
ONLYSMALL             2            10123
```

The results show that the ONLYSMALL application produces the majority of the unclassified activities. Because this application was not included in the results when you collected information about the largest activities, you can assume that ONLYSMALL did not produce any large queries during the period of data collection.

Scenario: Identifying activities with low estimated cost and high runtime

The following example shows how you can use work classes, work action sets, thresholds, and activity collection to identify activities that have a low estimated cost but a high runtime. This situation could indicate that the estimated cost (in timerons) is inaccurate because of out-of-date table and index statistics.

The first step is to create a work class set with a work class that will be used to identify activities with a low estimated cost. For example:

```
CREATE WORK CLASS SET WCS1
(WORK CLASS SMALLDML WORK TYPE DML FOR TIMERONCOST FROM 0 TO 500)
```

Then, you would create a database work action set with a work action that applies an activity-total-time threshold to the SMALLDML work class. The threshold action is CONTINUE and the COLLECT ACTIVITY DATA option is specified so that an activity that violates the threshold is sent to the activities event monitor on completion:

```
CREATE WORK ACTION SET WAS1 FOR DATABASE USING WORK CLASS SET WCS1
(WORK ACTION WA1 ON WORK CLASS SMALLDML WHEN ACTIVITYTOTALTIME > 15 MINUTES
COLLECT ACTIVITY DATA WITH DETAILS CONTINUE)
```

Finally, you would create and activate a threshold violations event monitor and an activities event monitor:

```
CREATE EVENT MONITOR THVIOLATIONS FOR THRESHOLD VIOLATIONS WRITE TO TABLE
SET EVENT MONITOR THVIOLATIONS STATE 1
```

```
CREATE EVENT MONITOR DB2ACTIVITIES FOR ACTIVITIES WRITE TO TABLE
SET EVENT MONITOR DB2ACTIVITIES STATE 1
```

Now when a DML activity with an estimated cost of less than 500 timerons runs for greater than 15 minutes, a threshold violation record is written to the THVIOLATIONS event monitor (indicating that the total time threshold was violated), and details about the DML activity are collected when the activity completes and sent to the DB2ACTIVITIES event monitor. You can use the information collected about the activity in the DB2ACTIVITIES event monitor to investigate further. For example, you can run the EXPLAIN statement on the query and examine the access plan. You should also consider the system load and queuing at the time the activity was collected, as a long lifetime can be a result of insufficient system resources or the activity being queued. The long lifetime does not necessarily indicate out-of-date statistics.

Chapter 8. Reference

Procedures and table functions

WLM_CANCEL_ACTIVITY - Cancel an activity

This procedure cancels a given activity. If the cancel takes place, an error message will be returned to the application that submitted the activity that was cancelled.

Syntax

```
►►—WLM_CANCEL_ACTIVITY—(—application_handle—,—uow_id—,—activity_id—)————►◄
```

The schema is SYSPROC.

Procedure parameters

application_handle

An input argument of type BIGINT that specifies the application handle whose activity is to be cancelled. If the argument is null, no activity will be found and an SQL4702N with SQLSTATE 5U035 is returned.

uow_id

An input argument of type INTEGER that specifies the unit of work ID of the activity that is to be cancelled. If the argument is null, no activity will be found and an SQL4702N with SQLSTATE 5U035 is returned.

activity_id

An input argument of type INTEGER that specifies the activity ID which uniquely identifies the activity within the unit of work that is to be cancelled. If the argument is null, no activity will be found and an SQL4702N with SQLSTATE 5U035 is returned.

Authorization

EXECUTE privilege on the WLM_CANCEL_ACTIVITY procedure.

Example

An administrator can use the WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES table function to find the application handle, unit of work ID and activity ID of an activity. To cancel an activity with application handle 1, unit of work ID 2 and activity ID 3:

```
CALL WLM_CANCEL_ACTIVITY(1, 2, 3)
```

Usage notes

- If no activity can be found, an SQL4702N with SQLSTATE 5U035 is returned.
- If the activity cannot be cancelled because it not in the correct state (not initialized), an SQL4703N (reason code 1) with SQLSTATE 5U016 is returned.
- If the activity is successfully cancelled, an SQL4725N with SQLSTATE 57014 is returned to the cancelled application.

- If, at the time of the cancel, the coordinator is processing a request for a different activity or is idle, the activity is placed into CANCEL_PENDING state and will be cancelled when the coordinator processes the next request for the activity.

WLM_CAPTURE_ACTIVITY_IN_PROGRESS - Collect activity information for activities event monitor

The WLM_CAPTURE_ACTIVITY_IN_PROGRESS procedure gathers information about a specified activity and writes the information to the active activities event monitor.

When you apply this procedure to an activity with child activities, the procedure recursively generates a record for each child activity. This information is collected and sent when you call the procedure; the procedure does not wait until the parent activity completes execution. The record of the activity in the event monitor is marked as a partial record.

Syntax

```
►► WLM_CAPTURE_ACTIVITY_IN_PROGRESS(—application_handle—, —————►
►—uow_id—, —activity_id—)—————►►
```

The schema is SYSPROC.

Procedure parameters

If you do not specify all of the following parameters, no activity is found, and SQL4702N with SQLSTATE 5U035 is returned.

application_handle

An input argument of type BIGINT that specifies the handle of the application whose activity information is to be captured.

uow_id

An input argument of type INTEGER that specifies the unit of work ID of the activity whose information is to be captured.

activity_id

An input argument of type INTEGER that specifies the activity ID that uniquely identifies the activity within the unit of work whose information is to be captured.

Authorization

EXECUTE privilege on the WLM_CAPTURE_ACTIVITY_IN_PROGRESS procedure.

Example

Assume that a user complains that stored procedure MYSCHEMA.MYSLOWSTP seems to be running more slowly than usual. The administrator wants to investigate the cause of the slowdown. Investigating while the stored procedure is running is not practical, so the administrator decides to capture information about the stored procedure activity and all of the activities nested within it.

An event monitor for DB2 activities named DB2ACTIVITIES has been activated. The administrator uses the WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES function to obtain the application handle, unit of work ID and activity ID for the call of this stored procedure. Assuming that the activity is identified by an application handle of 1, a unit of work ID of 2 and an activity ID of 3, the administrator can now issue the call to WLM_CAPTURE_ACTIVITY_IN_PROGRESS as follows:

```
CALL WLM_CAPTURE_ACTIVITY_IN_PROGRESS(1,2,3)
```

After the procedure is completed, the administrator can use the following table function to find out where the activity spent its time. The function retrieves the information from the DB2ACTIVITIES event monitor.

```
CREATE FUNCTION SHOWCAPTUREDACTIVITY(APPHNDL BIGINT,
                                     UOWID INTEGER,
                                     ACTIVITYID INTEGER)
RETURNS TABLE (UOW_ID INTEGER, ACTIVITY_ID INTEGER, STMT_TEXT VARCHAR(40),
                LIFE_TIME DOUBLE)
LANGUAGE SQL
READS SQL DATA
NO EXTERNAL ACTION
DETERMINISTIC
RETURN WITH RAH (LEVEL, APPL_ID, PARENT_UOW_ID, PARENT_ACTIVITY_ID,
                 UOW_ID, ACTIVITY_ID, STMT_TEXT, ACT_EXEC_TIME) AS
(SELECT 1, ROOT.APPL_ID, ROOT.PARENT_UOW_ID,
        ROOT.PARENT_ACTIVITY_ID, ROOT.UOW_ID, ROOT.ACTIVITY_ID,
        ROOTSTMT.STMT_TEXT, ACT_EXEC_TIME
 FROM ACTIVITY_DB2ACTIVITIES ROOT, ACTIVITYSTMT_DB2ACTIVITIES ROOTSTMT
 WHERE ROOT.APPL_ID = ROOTSTMT.APPL_ID AND ROOT.AGENT_ID = APPHNDL
       AND ROOT.UOW_ID = ROOTSTMT.UOW_ID AND ROOT.UOW_ID = UOWID
       AND ROOT.ACTIVITY_ID = ROOTSTMT.ACTIVITY_ID AND ROOT.ACTIVITY_ID = ACTIVITYID
 UNION ALL
 SELECT PARENT.LEVEL +1, CHILD.APPL_ID, CHILD.PARENT_UOW_ID,
        CHILD.PARENT_ACTIVITY_ID, CHILD.UOW_ID,
        CHILD.ACTIVITY_ID, CHILDSTMT.STMT_TEXT, CHILD.ACT_EXEC_TIME
 FROM RAH PARENT, ACTIVITY_DB2ACTIVITIES CHILD,
        ACTIVITYSTMT_DB2ACTIVITIES CHILDSTMT
 WHERE PARENT.APPL_ID = CHILD.APPL_ID AND
       CHILD.APPL_ID = CHILDSTMT.APPL_ID AND
       PARENT.UOW_ID = CHILD.PARENT_UOW_ID AND
       CHILD.UOW_ID = CHILDSTMT.UOW_ID AND
       PARENT.ACTIVITY_ID = CHILD.PARENT_ACTIVITY_ID AND
       CHILD.ACTIVITY_ID = CHILDSTMT.ACTIVITY_ID AND
       PARENT.LEVEL < 64
 )
SELECT UOW_ID, ACTIVITY_ID, SUBSTR(STMT_TEXT,1,40),
       ACT_EXEC_TIME AS
       LIFE_TIME
FROM RAH
```

The following sample query uses the table function:

```
SELECT * FROM TABLE(SHOWCAPTUREDACTIVITY(1, 2, 3))
AS ACTS ORDER BY UOW_ID, ACTIVITY_ID
```

Usage notes

If there is no active activities event monitor, an SQL1633W with SQLSTATE 01H53 is returned.

Activity information is collected only on the coordinator partition for the activity.

WLM_COLLECT_STATS - Collect and reset workload management statistics

The WLM_COLLECT_STATS procedure gathers statistics for service classes, workloads, work classes, and threshold queues and writes them to the statistics event monitor. The procedure also resets the statistics for service classes, workloads, work classes, and threshold queues. If there is no active statistics event monitor, the procedure only resets the statistics.

Syntax

►►—WLM_COLLECT_STATS—(—)—————►►

The schema is SYSPROC.

Authorization

EXECUTE privilege on the WLM_COLLECT_STATS procedure.

Examples

Example 1: Call WLM_COLLECT_STATS to collect and reset statistics.

```
CALL WLM_COLLECT_STATS()
```

The following is an example of output from this query.

```
Return Status = 0
```

Example 2: Call WLM_COLLECT_STATS to collect and reset statistics while another call is in progress.

```
CALL WLM_COLLECT_STATS()
```

The following is an example of output from this query.

```
SQL1632W The collect and reset statistics request was ignored because  
another collect and reset statistics request is already in progress.
```

Usage notes

The WLM_COLLECT_STATS procedure performs the same collection operation (send statistics to the active statistics event monitor) and reset operation that occur automatically on the interval defined by the **wlm_collect_int** database configuration parameter.

If you call the procedure while another collection and reset request is in progress (for example, while another invocation of the procedure is running or automatic collection is occurring), SQL1632W with SQLSTATE 01H53 is returned, and your new request is ignored.

The WLM_COLLECT_STATS procedure only starts the collection and reset process. The procedure might return to the caller before all statistics have been written to the active statistics event monitor. Depending on how quickly the statistics collection and reset occur, the call to the WLM_COLLECT_STATS procedure (which is itself an activity) is counted in the statistics for either the prior collection interval or the new collection interval that has just started.

WLM_GET_ACTIVITY_DETAILS - Return detailed information about a specific activity

Note: This table function has been deprecated and replaced by the MON_GET_ACTIVITY_DETAILS table function.

This function returns basic statistics of one or more service subclasses.

This function returns detailed information about a specific activity identified by its application handle, unit of work ID, and activity ID. This information includes details about any thresholds that the activity has violated.

Syntax

```
►►WLM_GET_ACTIVITY_DETAILS(►►application_handle►►,►►uow_id►►,►►
►►activity_id►►,►►dbpartitionnum►►)►►►►
```

The schema is SYSPROC.

Table function parameters

application_handle

An input argument of type BIGINT that specifies a valid application handle. If the argument is null, no rows are returned from this function. If the argument is null, an SQL171N error is returned.

uow_id

An input argument of type INTEGER that specifies a valid unit of work identifier unique within the application. If the argument is null, no rows are returned from this function. If the argument is null, an SQL171N error is returned.

activity_id

An input argument of type INTEGER that specifies a valid activity ID unique within the unit of work. If the argument is null, no rows are returned from this function. If the argument is null, an SQL171N error is returned.

dbpartitionnum

An input argument of type INTEGER that specifies a valid partition number in the same instance as the currently connected database when calling this function. Specify -1 for the current database partition, or -2 for all database partitions. If a null value is specified, -1 is set implicitly.

Authorization

EXECUTE privilege on the WLM_GET_ACTIVITY_DETAILS function.

Example

Detailed information about an individual activity can be obtained by using the WLM_GET_ACTIVITY_DETAILS table function. This table function returns activity information as name-value pairs for each partition. This example is restricted to showing only an eleven member subset of the name-value pairs for each partition

for an activity identified by an application handle of 1, a unit of work ID of 1 and an activity ID of 5. For a complete list of name-value pairs, see Table 62 on page 283 and Table 63 on page 285.

```

SELECT SUBSTR(CHAR(DBPARTITIONNUM),1,4) AS PART,
       SUBSTR(NAME, 1, 20) AS NAME,
       SUBSTR(VALUE, 1, 30) AS VALUE
FROM TABLE(WLM_GET_ACTIVITY_DETAILS(1, 1, 5, -2)) AS ACTDETAIL
WHERE NAME IN ('APPLICATION_HANDLE',
              'COORD_PARTITION_NUM',
              'LOCAL_START_TIME',
              'UOW_ID',
              'ACTIVITY_ID',
              'PARENT_UOW_ID',
              'PARENT_ACTIVITY_ID',
              'ACTIVITY_TYPE',
              'NESTING_LEVEL',
              'INVOCATION_ID',
              'ROUTINE_ID')
ORDER BY PART

```

The following is an example of output from this query.

PART	NAME	VALUE
0	APPLICATION_HANDLE	1
0	COORD_PARTITION_NUM	0
0	LOCAL_START_TIME	2005-11-25-18.52.49.343000
0	UOW_ID	1
0	ACTIVITY_ID	5
0	PARENT_UOW_ID	1
0	PARENT_ACTIVITY_ID	3
0	ACTIVITY_TYPE	READ_DML
0	NESTING_LEVEL	0
0	INVOCATION_ID	1
0	ROUTINE_ID	0
1	APPLICATION_HANDLE	1
1	COORD_PARTITION_NUM	0
1	LOCAL_START_TIME	2005-11-25-18.52.49.598000
1	UOW_ID	1
1	ACTIVITY_ID	5
1	PARENT_UOW_ID	
1	PARENT_ACTIVITY_ID	
1	ACTIVITY_TYPE	READ_DML
1	NESTING_LEVEL	0
1	INVOCATION_ID	1
1	ROUTINE_ID	0

Usage note

An `ACTIVITY_STATE` of `QUEUED` means that the coordinator activity has made a RPC to the catalog partition to obtain threshold tickets and has not yet received a response. Seeing this state might indicate that the activity has been queued by WLM or, over short periods of time, might just indicate that the activity is in the process of obtaining its tickets. To obtain a more accurate picture of whether or not the activity is really being queued, one can determine which agent is working on the activity (using the `WLM_GET_SERVICE_CLASS_AGENTS` table function) and find out whether this agent's `event_object` at the catalog partition has a value of `WLM_QUEUE`.

Information returned

Table 61. Information returned for WLM_GET_ACTIVITY_DETAILS

Column Name	Data Type	Description
DBPARTITIONNUM	SMALLINT	Partition number from which this record was collected.
NAME	VARCHAR(256)	Element name. See Table 62 and Table 63 on page 285 for possible values.
VALUE	VARCHAR(1024)	Element values. See Table 62 and Table 63 on page 285 for possible values.

Table 62. Elements returned

Element Name	Description
ACTIVITY_ID	Unique activity identifier within an application.
ACTIVITY_STATE	Possible values include: <ul style="list-style-type: none"> • CANCEL_PENDING • EXECUTING • IDLE • INITIALIZING • QP_CANCEL_PENDING • QP_QUEUED • QUEUED • TERMINATING • UNKNOWN
ACTIVITY_TYPE	Possible values include: <ul style="list-style-type: none"> • CALL • DDL • LOAD • OTHER • READ_DML • WRITE_DML
APPLICATION_HANDLE	A system-wide unique ID for the application. On a single-partitioned database, this identifier consists of a 16 bit counter. On a multi-partitioned database, this identifier consists of the coordinating partition number concatenated with a 16 bit counter. In addition, this identifier will be the same on every partition where the application may make a secondary connection.
COORD_PARTITION_NUM	The coordinator partition of the activity.
DATABASE_WORK_ACTION_SET_ID	If this activity has been mapped to a work action set that has been applied to the database, this column contains the ID of the work action set. This column contains 0 if the activity has not been mapped to a work action set that has been applied to the database.
DATABASE_WORK_CLASS_ID	If this activity has been mapped to a work action set that has been applied to the database, this column contains the ID of the work class of this activity. This column contains 0 if the activity has not been mapped to a work action set that has been applied to the database.
EFFECTIVE_ISOLATION	The effective isolation level for this activity.
EFFECTIVE_LOCK_TIMEOUT	The effective lock timeout value for this activity.

Table 62. Elements returned (continued)

Element Name	Description
EFFECTIVE_QUERY_DEGREE	The effective value of query degree for this activity.
ENTRY_TIME	The time that this activity arrived into the system.
INVOCATION_ID	This distinguishes one particular invocation of this activity from others at the same nesting level. Returns zero if the activity is not nested.
LAST_REFERENCE_TIME	Every time a request occurs in this activity, this field is updated.
LOCAL_START_TIME	The time that this activity began doing work on the partition. It is in local time. This field can be an empty string when an activity has entered the system but is in a queue and has not started executing.
NESTING_LEVEL	This represents the nesting level of this activity. Nesting level is the depth to which this activity is nested within its top-most parent activity.
PACKAGE_NAME	If the activity is a SQL statement, this represents the name of its package.
PACKAGE_SCHEMA	If the activity is a SQL statement, this represents the schema name of its package.
PACKAGE_VERSION_ID	If the activity is a SQL statement, this represents the version of its package.
PARENT_ACTIVITY_ID	Unique activity identifier within a unit of work for the parent of the activity whose ID is ACTIVITY_ID. Returns an empty string if the activity has no parent activity.
PARENT_UOW_ID	Unique unit of work identifier within an application. Refers to the original unit of work this activity's parent activity started in. Returns an empty string if the activity has no parent activity or when at a remote partition.
QP_QUERY_ID	The query ID assigned to this activity by Query Patroller if the activity is a query. A query ID of 0 indicates that Query Patroller did not assign a query ID to this activity.
QUERY_COST_ESTIMATE	Estimated cost, in timerons, for a query, as determined by the SQL compiler.
ROUTINE_ID	Routine unique identifier. Returns zero if the activity is not part of a routine.
ROWS_FETCHED	This is the number of rows read from the table. This reports only those values for the database partition for which this record is recorded. On DPF systems, these values may not reflect the correct totals for the whole activity. When the statement monitor switch is not turned on, this element is not collected and -1 is written instead.
ROWS_MODIFIED	This is the number of rows inserted, updated, or deleted. This reports only those values for the database partition for which this record is recorded. On DPF systems, these values may not reflect the correct totals for the whole activity. When the statement monitor switch is not turned on, this element is not collected and -1 is written instead.
SECTION_NUMBER	If the activity is a SQL statement, this represents its section number.

Table 62. Elements returned (continued)

Element Name	Description
SERVICE_CLASS_ID	Unique identifier of the service class to which this activity belongs.
SERVICE_CLASS_WORK_ACTION_SET_ID	If this activity has been mapped to a work action set that has been applied to a service class, this column contains the ID of the work action set. This column contains 0 if the activity has not been mapped to a work action set that has been applied to a service class.
SERVICE_CLASS_WORK_CLASS_ID	If this activity has been mapped to a work action set that has been applied to a service class, this column contains the ID of the work class of this activity. This column contains 0 if the activity has not been mapped to a work action set that has been applied to a service class.
STMT_PKG_CACHE_ID	Statement package cache identifier.
STMT_TEXT	If the activity is dynamic SQL or it is static SQL for which the statement text is available, this field contains the first 1024 characters of the statement text. It is an empty string otherwise.
SYSTEM_CPU_TIME	The total system CPU time (in seconds and microseconds) used by the database manager agent process, the unit of work, or the statement. When either the statement monitor switch or the timestamp switch is not turned on, this element is not collected and -1 is written instead.
UOW_ID	Unique unit of work identifier within an application. Refers to the original unit of work this activity started in.
USER_CPU_TIME	The total user CPU time (in seconds and microseconds) used by the database manager agent process, the unit of work, or the statement. When either the statement monitor switch or the timestamp switch is not turned on, this element is not collected and -1 is written instead.
UTILITY_ID	If the activity is a utility, this is its utility ID. Otherwise, this field is 0.

Important: The WLM_GET_ACTIVITY_DETAILS table function shows only the thresholds that are currently being applied to an activity.

The following elements are returned only if the corresponding thresholds apply to the activity.

Table 63. Elements returned if applicable

Element Name	Description
ACTIVITYTOTALTIME_THRESHOLD_ID	The ID of the ACTIVITYTOTALTIME threshold that was applied to the activity.
ACTIVITYTOTALTIME_THRESHOLD_VALUE	A timestamp that is computed by adding the ACTIVITYTOTALTIME threshold duration to the activity entry time. If the activity is still executing when this timestamp is reached, the threshold will be violated.
ACTIVITYTOTALTIME_THRESHOLD_VIOLATED	'Yes' indicates that the activity violated the ACTIVITYTOTALTIME threshold. 'No' indicates that the activity has not yet violated the threshold.

Table 63. Elements returned if applicable (continued)

Element Name	Description
CONCURRENTDBCOORDACTIVITIES_DB_THRESHOLD_ID	The ID of the CONCURRENTDBCOORDACTIVITIES_DB threshold that was applied to the activity.
CONCURRENTDBCOORDACTIVITIES_DB_THRESHOLD_QUEUED	'Yes' indicates that the activity was queued by the CONCURRENTDBCOORDACTIVITIES_DB threshold. 'No' indicates that the activity was not queued.
CONCURRENTDBCOORDACTIVITIES_DB_THRESHOLD_VALUE	The upper bound of the CONCURRENTDBCOORDACTIVITIES_DB threshold that was applied to the activity.
CONCURRENTDBCOORDACTIVITIES_DB_THRESHOLD_VIOLATED	'Yes' indicates that the activity violated the CONCURRENTDBCOORDACTIVITIES_DB threshold. 'No' indicates that the activity has not yet violated the threshold.
CONCURRENTDBCOORDACTIVITIES_SUBCLASS_THRESHOLD_ID	The ID of the CONCURRENTDBCOORDACTIVITIES_SUBCLASS threshold that was applied to the activity.
CONCURRENTDBCOORDACTIVITIES_SUBCLASS_THRESHOLD_QUEUED	'Yes' indicates that the activity was queued by the CONCURRENTDBCOORDACTIVITIES_SUBCLASS threshold. 'No' indicates that the activity was not queued.
CONCURRENTDBCOORDACTIVITIES_SUBCLASS_THRESHOLD_VALUE	The upper bound of the CONCURRENTDBCOORDACTIVITIES_SUBCLASS threshold that was applied to the activity.
CONCURRENTDBCOORDACTIVITIES_SUBCLASS_THRESHOLD_VIOLATED	'Yes' indicates that the activity violated the CONCURRENTDBCOORDACTIVITIES_SUBCLASS threshold. 'No' indicates that the activity has not yet violated the threshold.
CONCURRENTDBCOORDACTIVITIES_SUPERCLASS_THRESHOLD_ID	The ID of the CONCURRENTDBCOORDACTIVITIES_SUPERCLASS threshold that was applied to the activity.
CONCURRENTDBCOORDACTIVITIES_SUPERCLASS_THRESHOLD_QUEUED	'Yes' indicates that the activity was queued by the CONCURRENTDBCOORDACTIVITIES_SUPERCLASS threshold. 'No' indicates that the activity was not queued.
CONCURRENTDBCOORDACTIVITIES_SUPERCLASS_THRESHOLD_VALUE	The upper bound of the CONCURRENTDBCOORDACTIVITIES_SUPERCLASS threshold that was applied to the activity.
CONCURRENTDBCOORDACTIVITIES_SUPERCLASS_THRESHOLD_VIOLATED	'Yes' indicates that the activity violated the CONCURRENTDBCOORDACTIVITIES_SUPERCLASS threshold. 'No' indicates that the activity has not yet violated the threshold.
CONCURRENTDBCOORDACTIVITIES_WORK_ACTION_SET_THRESHOLD_ID	The ID of the CONCURRENTDBCOORDACTIVITIES_WORK_ACTION_SET threshold that was applied to the activity.
CONCURRENTDBCOORDACTIVITIES_WORK_ACTION_SET_THRESHOLD_QUEUED	'Yes' indicates that the activity was queued by the CONCURRENTDBCOORDACTIVITIES_WORK_ACTION_SET threshold. 'No' indicates that the activity was not queued.

Table 63. Elements returned if applicable (continued)

Element Name	Description
CONCURRENTDBCOORDACTIVITIES_WORK_ACTION_SET_THRESHOLD_VALUE	The upper bound of the CONCURRENTDBCOORDACTIVITIES_WORK_ACTION_SET threshold that was applied to the activity.
CONCURRENTDBCOORDACTIVITIES_WORK_ACTION_SET_THRESHOLD_VIOLATED	'Yes' indicates that the activity violated the CONCURRENTDBCOORDACTIVITIES_WORK_ACTION_SET threshold. 'No' indicates that the activity has not yet violated the threshold.
CONCURRENTWORKLOADACTIVITIES_THRESHOLD_ID	The ID of the CONCURRENTWORKLOADACTIVITIES threshold that was applied to the activity.
CONCURRENTWORKLOADACTIVITIES_THRESHOLD_VALUE	The upper bound of the CONCURRENTWORKLOADACTIVITIES threshold that was applied to the activity.
CONCURRENTWORKLOADACTIVITIES_THRESHOLD_VIOLATED	'Yes' indicates that the activity violated the CONCURRENTWORKLOADACTIVITIES threshold. 'No' indicates that the activity has not yet violated the threshold.
ESTIMATEDSQLCOST_THRESHOLD_ID	The ID of the ESTIMATEDSQLCOST threshold that was applied to the activity.
ESTIMATEDSQLCOST_THRESHOLD_VALUE	The upper bound of the ESTIMATEDSQLCOST threshold that was applied to the activity.
ESTIMATEDSQLCOST_THRESHOLD_VIOLATED	'1' indicates that the activity violated the ESTIMATEDSQLCOST threshold. '0' indicates that the activity has not yet violated the threshold.
SQLROWSRETURNED_THRESHOLD_ID	The ID of the SQLROWSRETURNED threshold that was applied to the activity
SQLROWSRETURNED_THRESHOLD_VALUE	The upper bound of the SQLROWSRETURNED threshold that was applied to the activity.
SQLROWSRETURNED_THRESHOLD_VIOLATED	'Yes' indicates that the activity violated the SQLROWSRETURNED threshold. 'No' indicates that the activity has not yet violated the threshold.
SQLTEMPSPACE_THRESHOLD_ID	The ID of the SQLTEMPSPACE threshold that was applied to the activity.
SQLTEMPSPACE_THRESHOLD_VALUE	The upper bound of the SQLTEMPSPACE threshold that was applied to the activity.
SQLTEMPSPACE_THRESHOLD_VIOLATED	'Yes' indicates that the activity violated the SQLTEMPSPACE threshold. 'No' indicates that the activity has not yet violated the threshold.

WLM_GET_QUEUE_STATS table function - Return threshold queue statistics

The WLM_GET_QUEUE_STATS function returns basic statistics for one or more threshold queues on all active partitions. This function returns one row of statistics for each threshold queue.

Syntax

```
►—WLM_GET_QUEUE_STATS—(—threshold_predicate—,—threshold_domain—,—  
►—threshold_name—,—threshold_id—)
```

The schema is SYSPROC.

Table function parameters

threshold_predicate

An input argument of type VARCHAR(27) that specifies a threshold predicate. The possible values are as follows:

CONCDBC

Concurrent database coordinator activities threshold

DBCONN

Total database partition connections threshold

SCCONN

Total service class partition connections threshold

If the argument is null or an empty string, data is returned for all thresholds that meet the other criteria.

The *threshold_predicate* values match those of the THRESHOLDPREDICATE column in the SYSCAT.THRESHOLDS view.

threshold_domain

An input argument of type VARCHAR(18) that specifies a threshold domain. The possible values are as follows:

DB Database

SB Service subclass

SP Service superclass

WA Work action set

If the argument is null or an empty string, data is returned for all thresholds that meet the other criteria.

The *threshold_domain* values match those of the DOMAIN column in the SYSCAT.THRESHOLDS view.

threshold_name

An input argument of type VARCHAR(128) that specifies a threshold name. If the argument is null or an empty string, data is returned for all thresholds that meet the other criteria. The *threshold_name* values match those of the THRESHOLDNAME column in the SYSCAT.THRESHOLDS view.

threshold_id

An input argument of type INTEGER that specifies a threshold ID. If the argument is null or -1, data is returned for all thresholds that meet the other criteria. The *threshold_id* values match those of the THRESHOLDID column in the SYSCAT.THRESHOLDS view.

Authorization

EXECUTE privilege on the WLM_GET_QUEUE_STATS function.

Example

The following query displays the basic statistics for all the queues on a system, across all partitions:

```
SELECT substr(THRESHOLD_NAME, 1, 6) THRESHNAME,
       THRESHOLD_PREDICATE,
       THRESHOLD_DOMAIN,
       DBPARTITIONNUM PART,
       QUEUE_SIZE_TOP,
       QUEUE_TIME_TOTAL,
       QUEUE_ASSIGNMENTS_TOTAL QUEUE_ASSIGN
FROM table(WLM_GET_QUEUE_STATS('',' ', -1)) as QSTATS
```

Sample output is as follows:

```
THRESHNAME THRESHOLD_PREDICATE      THRESHOLD_DOMAIN  ...
-----
LIMIT1     CONCDBC                DB                ...
LIMIT2     SCCONN                 SP                ...
LIMIT3     DBCONN                 DB                ...
... PART  QUEUE_SIZE_TOP  QUEUE_TIME_TOTAL  QUEUE_ASSIGN
... -----
... 0           12           1238540           734
... 0            4           741249            24
... 0            7           412785            128
```

Usage note

The function does not aggregate data across queues (on a partition) or across partitions (for one or more queues). However, you can use SQL queries to aggregate data, as shown in the previous example.

Information returned

Table 64. Information returned for WLM_GET_QUEUE_STATS

Column name	Data type	Description
THRESHOLD_PREDICATE	VARCHAR(27)	<p>Threshold predicate of the threshold responsible for this queue. The possible values are as follows:</p> <p><i>CONCDBC</i> Concurrent database coordinator activities threshold</p> <p><i>DBCONN</i> Total database partition connections threshold</p> <p><i>SCCONN</i> Total service class partition connections threshold</p> <p>The threshold predicate values match those of the THRESHOLDPREDICATE column in the SYSCAT.THRESHOLDS view.</p>

Table 64. Information returned for WLM_GET_QUEUE_STATS (continued)

Column name	Data type	Description
THRESHOLD_DOMAIN	VARCHAR(18)	Domain of the threshold responsible for this queue. The possible values are as follows: <i>DB</i> Database <i>SB</i> Service subclass <i>SP</i> Service superclass <i>WA</i> Work action set The threshold domain values match those of the DOMAIN column in the SYSCAT.THRESHOLDS view.
THRESHOLD_NAME	VARCHAR(128)	Unique name of the threshold responsible for this queue. The threshold name value matches that of the THRESHOLDNAME column in the SYSCAT.THRESHOLDS view.
THRESHOLD_ID	INTEGER	Unique ID of the threshold responsible for this queue. The threshold ID value matches that of the THRESHOLDID column in the SYSCAT.THRESHOLDS view.
DBPARTITIONNUM	SMALLINT	Partition number from which this record was collected.
SERVICE_SUPERCLASS_NAME	VARCHAR(128)	Name of the service superclass that is the domain for the threshold responsible for this queue. The value of the column is null if the domain of the threshold is not a service superclass or service subclass.
SERVICE_SUBCLASS_NAME	VARCHAR(128)	Name of the service subclass that is the domain for the threshold responsible for this queue. The value of the column is null if the domain of the threshold is not a service subclass.
WORK_ACTION_SET_NAME	VARCHAR(128)	Name of the work action set that is the domain for the threshold responsible for this queue. The value of the column is null if the domain of the threshold is not a work action set.
WORK_CLASS_NAME	VARCHAR(128)	Name of the work class whose work action belongs to the work action set that is the domain for the threshold responsible for this queue. The value of the column is null if the domain of the threshold is not a work action set.
WORKLOAD_NAME	VARCHAR(128)	Name of the workload that is the domain for the threshold responsible for this queue. The value of the column is null if the domain of the threshold is not a workload.

Table 64. Information returned for WLM_GET_QUEUE_STATS (continued)

Column name	Data type	Description
LAST_RESET	TIMESTAMP	Time when statistics were last reset. There are four events that trigger a reset of statistics: <ul style="list-style-type: none"> You call the WLM_COLLECT_STATS procedure. The wlm_collect_int configuration parameter causes a collection and reset. You reactivate the database. You modify the threshold for which queue statistics are being reported and commit the change. The LAST_RESET time stamp is in local time.
QUEUE_SIZE_TOP	INTEGER	Highest number of connections or activities in the queue since the last reset.
QUEUE_TIME_TOTAL	BIGINT	Sum of the times spent in the queue for all connections or activities placed in this queue since the last reset. Units are milliseconds.
QUEUE_ASSIGNMENTS_TOTAL	BIGINT	Number of connections or activities that were assigned to this queue since the last reset.
QUEUE_SIZE_CURRENT	INTEGER	Number of connections or activities in the queue.
QUEUE_TIME_LATEST	BIGINT	Time spent in the queue by the last connection or activity to leave the queue. Units are milliseconds.
QUEUE_EXIT_TIME_LATEST	TIMESTAMP	Time that the last connection or activity left the queue.
THRESHOLD_CURRENT_CONCURRENCY	INTEGER	Number of connections or activities that are currently running according to the threshold.
THRESHOLD_MAX_CONCURRENCY	INTEGER	Maximum number of connections or activities that the threshold allows to be concurrently running.

WLM_GET_SERVICE_CLASS_AGENTS_V97 table function - List agents running in a service class

The WLM_GET_SERVICE_CLASS_AGENTS_V97 function returns the list of agents, fenced mode processes (db2fmp processes), and system entities on a specified partition that are running in a specified service class or on behalf of a specified application. The system entities are non-agent threads and processes, such as page cleaners and prefetchers.

Syntax


```

▶▶—WLM_GET_SERVICE_CLASS_AGENTS_V97—(—service_superclass_name—,—————▶
▶—service_subclass_name—, —application_handle—, —dbpartitionnum—)————▶▶

```

The schema is SYSPROC.

Table function parameters

service_superclass_name

An input argument of type VARCHAR(128) that specifies the name of a service superclass in the currently connected database. If the argument is null or an empty string, data is retrieved for all the superclasses in the database.

service_subclass_name

An input argument of type VARCHAR(128) that refers to a specific subclass within a superclass. If the argument is null or an empty string, data is retrieved for all the subclasses in the database.

application_handle

An input argument of type BIGINT that specifies the application handle for which agent information is to be returned. If the argument is null, data is retrieved for all applications in the database. An application handle of 0 returns the system entities only.

dbpartitionnum

An input argument of type INTEGER that specifies the partition number in the same instance as the currently connected database. Specify -1 for the current database partition, or -2 for all database partitions. If a null value is specified, -1 is set implicitly.

Authorization

EXECUTE privilege on the WLM_GET_SERVICE_CLASS_AGENTS_V97 function.

Example

Example 1

The following query returns a list of agents that are associated with application handle 1 for all database partitions. You can determine the application handle by using the LIST APPLICATIONS command or the WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97 table function.

```

SELECT SUBSTR(CHAR(APPLICATION_HANDLE),1,7) AS APPHANDLE,
       SUBSTR(CHAR(DBPARTITIONNUM),1,4) AS PART,
       SUBSTR(CHAR(AGENT_TID),1,9) AS AGENT_TID,
       SUBSTR(CHAR(AGENT_TYPE),1,11) AS AGENTTYPE,
       SUBSTR(CHAR(AGENT_STATE),1,10) AS AGENTSTATE,
       SUBSTR(CHAR(REQUEST_TYPE),1,12) AS REQTYPE,
       SUBSTR(CHAR(UOW_ID),1,6) AS UOW_ID,
       SUBSTR(CHAR(ACTIVITY_ID),1,6) AS ACT_ID
FROM TABLE(WLM_GET_SERVICE_CLASS_AGENTS_V97(CAST(NULL AS VARCHAR(128)),
      CAST(NULL AS VARCHAR(128)), 1, -2)) AS SCDETAILS
ORDER BY APPHANDLE, PART, AGENT_TID

```

Sample output is as follows:

APPHANDLE	PART	AGENT_TID	AGENTTYPE	AGENTSTATE	REQTYPE	UOW_ID	ACT_ID
1	0	3	COORDINATOR	ACTIVE	FETCH	1	5
1	0	4	SUBAGENT	ACTIVE	SUBSECTION:1	1	5
1	1	2	SUBAGENT	ACTIVE	SUBSECTION:2	1	5

The output shows a coordinator agent and a subagent on partition 0 and a subagent on partition 1 operating on behalf of an activity with UOW ID 1 and activity ID 5. The AGENTTYPE column with a value of COORDINATOR has a value of FETCH for the REQTYPE column (which indicates the main or initial request type). This means that the type of request is a fetch request for the coordinator agent.

Example 2

The following query determines which lock an agent is waiting on:

```
db2 select event_object, event_type, event_state, varchar(event_object_name, 30)
as event_object_name
from table(wlm_get_service_class_agents_v97('','',cast(NULL as bigint), -1)) as t
```

Sample output is as follows:

EVENT_OBJECT	EVENT_TYPE	EVENT_STATE	EVENT_OBJECT_NAME
REQUEST	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
REQUEST	WAIT	IDLE	-
LOCK	ACQUIRE	IDLE	020005000000000000000000000054
ROUTINE	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
WLM_QUEUE	WAIT	IDLE	MYCONCDBCOORDTH
ROUTINE	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-

21 record(s) selected.

Using the same query at a later time shows that the WLM threshold has queued an agent:

EVENT_OBJECT	EVENT_TYPE	EVENT_STATE	EVENT_OBJECT_NAME
REQUEST	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
WLM_QUEUE	WAIT	IDLE	MYCONCDBCOORDTH
ROUTINE	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-
REQUEST	PROCESS	EXECUTING	-

```

REQUEST      PROCESS      EXECUTING      -
REQUEST      PROCESS      EXECUTING      -
REQUEST      PROCESS      EXECUTING      -
REQUEST      PROCESS      EXECUTING      -
REQUEST      PROCESS      EXECUTING      -
REQUEST      PROCESS      EXECUTING      -
REQUEST      PROCESS      EXECUTING      -
REQUEST      PROCESS      EXECUTING      -

```

21 record(s) selected.

Usage note

The parameters are, in effect, ANDed together. That is, if you specify conflicting input parameters, such as a service superclass SUP_A and a subclass SUB_B such that SUB_B is not a subclass of SUP_A, no rows are returned.

Information returned

Table 65. Information returned by WLM_GET_SERVICE_CLASS_AGENTS_V97

Column name	Data type	Description
SERVICE_SUPERCLASS_NAME	VARCHAR (128)	Name of the service superclass from which this record was collected.
SERVICE_SUBCLASS_NAME	VARCHAR (128)	Name of the service subclass from which this record was collected.
APPLICATION_HANDLE	BIGINT	System-wide unique ID for the application. On a single-partitioned database, this identifier consists of a 16-bit counter. On a multi-partitioned database, this identifier consists of the coordinating partition number concatenated with a 16-bit counter. In addition, this identifier is the same on every partition where the application makes a secondary connection.
DBPARTITIONNUM	SMALLINT	Partition number from which this record was collected.
ENTITY	VARCHAR (32)	One of the following values: <ul style="list-style-type: none"> • If the type of entity is an agent, the value is db2agent. • If the type of entity is a fenced mode process, the value is db2fmp (<i>pid</i>) where <i>pid</i> is the process ID of the fenced mode process. • Otherwise, the value is the name of the system entity.
WORKLOAD_NAME	VARCHAR (128)	Name of the workload from which this record was collected.
WORKLOAD_OCCURRENCE_ID	INTEGER	ID of the workload occurrence. This ID does not uniquely identify the workload occurrence unless it is coupled with the coordinator database partition number and the workload name.
UOW_ID	INTEGER	Unique ID of the unit of work that this activity started in.
ACTIVITY_ID	INTEGER	Unique activity ID within a unit of work.
PARENT_UOW_ID	INTEGER	Unique ID of the unit of work that the parent activity of the activity started in. The value of the column is null if this activity has no parent.
PARENT_ACTIVITY_ID	INTEGER	Unique activity ID within a unit of work for the parent of the activity whose ID is the same as activity_id. The value of this column is null if this activity has no parent.

Table 65. Information returned by WLM_GET_SERVICE_CLASS_AGENTS_V97 (continued)

Column name	Data type	Description
AGENT_TID	BIGINT	Thread ID of the agent or system entity. If this ID is unavailable, the value of the column is null.
AGENT_TYPE	VARCHAR (32)	Agent type. The agent types are as follows: <ul style="list-style-type: none"> • COORDINATOR • OTHER • PDBSUBAGENT • SMPSUBAGENT If the value is COORDINATOR, the agent ID might change in concentrator environments.
SMP_COORDINATOR	INTEGER	Indication of whether the agent is an SMP coordinator: 1 for yes and 0 for no.
AGENT_SUBTYPE	VARCHAR (32)	Agent subtype. The possible subtypes are as follows: <ul style="list-style-type: none"> • DSS • OTHER • RPC • SMP
AGENT_STATE	VARCHAR (32)	Indication of whether an agent is associated or active. The possible values are: <ul style="list-style-type: none"> • ASSOCIATED • ACTIVE
EVENT_TYPE	VARCHAR (32)	Type of event last processed by this agent. The possible values are as follows: <ul style="list-style-type: none"> • ACQUIRE • PROCESS • WAIT
EVENT_OBJECT	VARCHAR (32)	Object of the event last processed by this agent. The possible values are as follows: <ul style="list-style-type: none"> • COMPRESSION_DICTIONARY_BUILD • IMPLICIT_REBIND • INDEX_RECREATE • LOCK • LOCK_ESCALATION • QP_QUEUE • REMOTE_REQUEST • REQUEST • ROUTINE • WLM_QUEUE
EVENT_STATE	VARCHAR (32)	State of the event last processed by this agent. The possible values are as follows: <ul style="list-style-type: none"> • EXECUTING • IDLE
REQUEST_ID	VARCHAR (64)	Request ID. This value is unique only in combination with the value of <i>application_handle</i> . You can use this combination to distinguish between one request that is taking a long time and multiple requests; for example, to distinguish between one long fetch and multiple fetches.

Table 65. Information returned by WLM_GET_SERVICE_CLASS_AGENTS_V97 (continued)

Column name	Data type	Description
REQUEST_TYPE	VARCHAR (32)	Type of request. The possible values are as follows: <ul style="list-style-type: none"> • For coordinator agents: <ul style="list-style-type: none"> – CLOSE – COMMIT – COMPILE – DESCRIBE – EXCSQLSET – EXECIMMD – EXECUTE – FETCH – INTERNAL <i>number</i>, where <i>number</i> is the value of the internal constant – OPEN – PREPARE – REBIND – REDISTRIBUTE – REORG – ROLLBACK – RUNSTATS • For subagents with an AGENT_SUBTYPE of DSS or SMP: <ul style="list-style-type: none"> – If the subsection number is nonzero, the subsection number in the form SUBSECTION:<i>subsection number</i>; otherwise, returns NULL.

Table 65. Information returned by WLM_GET_SERVICE_CLASS_AGENTS_V97 (continued)

Column name	Data type	Description
REQUEST_TYPE (continued)	VARCHAR (32)	<ul style="list-style-type: none"> • For subagents with an AGENT_SUBTYPE of RPC: <ul style="list-style-type: none"> - ABP - CATALOG - INTERNAL - REORG - RUNSTATS - WLM • For subagents with a SUBTYPE of OTHER: <ul style="list-style-type: none"> - ABP - APP_RBSVPT - APP_RELSVPT - BACKUP - CLOSE - EXTERNAL_RBSVPT - EVMON - FORCE - FORCE_ALL - INTERNAL <i>number</i>, where <i>number</i> is the value of the internal constant - INTERRUPT - NOOP (if there is no request) - QP - REDISTRIBUTE - STMT_RBSVPT - STOP_USING - UPDATE_DBM_CFG - WLM
NESTING_LEVEL	INTEGER	Nesting level of the activity whose ID is activity_id. Nesting level is the depth to which this activity is nested within its topmost parent activity.
INVOCATION_ID	INTEGER	Invocation ID, which distinguishes one particular invocation of an activity from others at the same nesting level.
ROUTINE_ID	INTEGER	Unique ID for a routine. The value of this column is null if the activity is not part of a routine.
EVENT_OBJECT_NAME	VARCHAR (1024)	Event object name. If the value of EVENT_OBJECT is LOCK, the value of this column is the name of the lock that the agent is waiting on. If the value of EVENT_OBJECT is WLM_QUEUE, the value of the column is the name of the WLM threshold that the agent is queued on. Otherwise, the value is NULL.
APPLICATION_NAME	VARCHAR (128)	Reserved for future use.
APPLICATION_ID	VARCHAR (128)	Reserved for future use.
CLIENT_PID	BIGINT	Reserved for future use.
SESSION_AUTH_ID	VARCHAR (128)	Reserved for future use.

Table 65. Information returned by WLM_GET_SERVICE_CLASS_AGENTS_V97 (continued)

Column name	Data type	Description
REQUEST_START_TIME	TIMESTAMP	Reserved for future use.
AGENT_STATE_LAST_UPDATE_TIME	TIMESTAMP	Reserved for future use.
EXECUTABLE_ID	VARCHAR (32) FOR BIT DATA	Binary token generated on the data server that uniquely identifies the section that an agent is working on. You can use the executable ID as input to different monitoring interfaces to obtain data about the section. A NULL value is returned if the agent is not working on a section.

WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97 - List workload occurrences

The WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97 function returns the list of all workload occurrences running in a specified service class on a particular partition. A workload occurrence is a specific database connection whose attributes match the definition of a workload and hence is associated with or assigned to the workload.

Syntax

```

▶▶—WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97—(—service_superclass_name—, —————▶
▶—service_subclass_name—, —dbpartitionnum—)—————▶▶

```

The schema is SYSPROC.

Table function parameters

service_superclass_name

An input argument of type VARCHAR(128) that specifies the name of a service superclass in the currently connected database. If the argument is null or an empty string, the data is retrieved for all the superclasses in the database that match the values of the other parameters.

service_subclass_name

An input argument of type VARCHAR(128) that specifies the name of a service subclass in the currently connected database. If the argument is null or an empty string, the data is retrieved for all the subclasses in the database that match the values of the other parameters.

dbpartitionnum

An input argument of type INTEGER that specifies the number of a partition in the same instance as the currently connected database. Specify -1 for the current database partition, or -2 for all database partitions. If the null value is specified, -1 is set implicitly.

Authorization

EXECUTE privilege on the WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97 function.

Example

If an administrator wants to see what workload occurrences are running on the system as a whole, the administrator can call the `WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97` function by specifying a null value or an empty string for `service_superclass_name` and `service_subclass_name` and `-2` for `dbpartitionnum`:

```
SELECT SUBSTR(SERVICE_SUPERCLASS_NAME,1,19) AS SUPERCLASS_NAME,
       SUBSTR(SERVICE_SUBCLASS_NAME,1,18) AS SUBCLASS_NAME,
       SUBSTR(CHAR(DBPARTITIONNUM),1,4) AS PART,
       SUBSTR(CHAR(COORD_PARTITION_NUM),1,4) AS COORDPART,
       SUBSTR(CHAR(APPLICATION_HANDLE),1,7) AS APPHNDL,
       SUBSTR(WORKLOAD_NAME,1,22) AS WORKLOAD_NAME,
       SUBSTR(CHAR(WORKLOAD_OCCURRENCE_ID),1,6) AS WLO_ID
FROM TABLE(WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97
            (CAST(NULL AS VARCHAR(128)), CAST(NULL AS VARCHAR(128)), -2))
AS SCINFO
ORDER BY SUPERCLASS_NAME, SUBCLASS_NAME, PART, APPHNDL,
        WORKLOAD_NAME, WLO_ID
```

If the system has four database partitions and is currently running two workloads, the previous query produces results such as the following ones:

SUPERCLASS_NAME	SUBCLASS_NAME	PART	COORDPART	...
-----	-----	-----	-----	-----
SYSDEFAULTMAINTENAN	SYSDEFAULTSUBCLASS	0	0	...
SYSDEFAULTSYSTEMCLA	SYSDEFAULTSUBCLASS	0	0	...
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	0	0	...
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	0	0	...
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	1	0	...
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	1	0	...
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	2	0	...
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	2	0	...
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	3	0	...
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	3	0	...
...	APPHNDL	WORKLOAD_NAME	WLO_ID	
...	-----	-----	-----	
...	-	-	-	
...	-	-	-	
...	1	SYSDEFAULTUSERWORKLOAD	1	
...	2	SYSDEFAULTUSERWORKLOAD	2	
...	1	SYSDEFAULTUSERWORKLOAD	1	
...	2	SYSDEFAULTUSERWORKLOAD	2	
...	1	SYSDEFAULTUSERWORKLOAD	1	
...	2	SYSDEFAULTUSERWORKLOAD	2	
...	1	SYSDEFAULTUSERWORKLOAD	1	
...	2	SYSDEFAULTUSERWORKLOAD	2	

Usage note

The parameters are, in effect, ANDed together. That is, if you specify conflicting input parameters, such as a service superclass `SUP_A` and a subclass `SUB_B` such that `SUB_B` is not a subclass of `SUP_A`, no rows are returned.

Note: Statistics reported for the workload occurrence (for example, `coord_act_completed_total`) are reset at the beginning of each unit of work when they are combined with the corresponding workload statistics.

Information returned

Table 66. Information returned for WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97

Column name	Data type	Description
SERVICE_SUPERCLASS_NAME	VARCHAR(128)	Name of the service superclass from which this record was collected.
SERVICE_SUBCLASS_NAME	VARCHAR(128)	Name of the service subclass from which this record was collected.
DBPARTITIONNUM	SMALLINT	Partition number from which this record was collected.
COORD_PARTITION_NUM	SMALLINT	Partition number of the coordinator partition of the specified workload occurrence.
APPLICATION_HANDLE	BIGINT	System-wide unique ID for the application. On a single-partitioned database, this identifier consists of a 16-bit counter. On a multi-partitioned database, this identifier consists of the coordinating partition number concatenated with a 16-bit counter. In addition, this identifier is the same on every partition where the application makes a secondary connection.
WORKLOAD_NAME	VARCHAR(128)	Name of the workload from which this record was collected.
WORKLOAD_OCCURRENCE_ID	INTEGER	ID of the workload occurrence. This ID does not uniquely identify the workload occurrence unless it is coupled with the coordinator database partition number and the workload name.

Table 66. Information returned for WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97 (continued)

Column name	Data type	Description
WORKLOAD_OCCURRENCE_STATE	VARCHAR(32)	<p>Workload occurrence state. The values are as follows:</p> <p><i>DECOUPLED</i> Workload occurrence does not have a coordinator agent assigned (concentrator case).</p> <p><i>DISCONNECTPEND</i> Workload occurrence is disconnecting from the database.</p> <p><i>FORCED</i> Workload occurrence has been forced off the database.</p> <p><i>INTERRUPTED</i> Workload occurrence has been interrupted.</p> <p><i>QUEUED</i> Workload occurrence coordinator agent is queued by Query Patroller or a workload management queuing threshold. In a database partitioning feature (DPF) environment, this state might indicate that the coordinator agent has made an RPC to the catalog partition to obtain threshold tickets and has not yet received a response.</p> <p><i>TRANSIENT</i> Workload occurrence has not yet been mapped to a service superclass.</p> <p><i>UOWEXEC</i> Workload occurrence is processing a request.</p> <p><i>UOWWAIT</i> Workload occurrence is waiting for a request from the client.</p>
UOW_ID	INTEGER	Unique ID of the unit of work that this workload occurrence started in.
SYSTEM_AUTH_ID	VARCHAR(128)	System authorization ID under which the workload occurrence was inserted into the system.
SESSION_AUTH_ID	VARCHAR(128)	Session authorization ID under which the workload occurrence was inserted into the system.
APPLICATION_NAME	VARCHAR(128)	Name of the application that created this workload occurrence.
CLIENT_WRKSTNNAME	VARCHAR(255)	Current value of the CLIENT_WRKSTNNAME special register for this workload occurrence.
CLIENT_ACCTNG	VARCHAR(255)	Current value of the CLIENT_ACCTNG special register for this workload occurrence.
CLIENT_USER	VARCHAR(255)	Current value of the CLIENT_USERID special register for this workload occurrence.

Table 66. Information returned for WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97 (continued)

Column name	Data type	Description
CLIENT_APPLNAME	VARCHAR(255)	Current value of the CLIENT_APPLNAME special register for this workload occurrence.
COORD_ACT_COMPLETED_TOTAL	INTEGER	Number of coordinator activities at any nesting level that were completed so far in the current unit of work of this workload occurrence. This statistic is updated every time that an activity in this workload occurrence is completed and is reset at the beginning of each unit of work.
COORD_ACT_ABORTED_TOTAL	INTEGER	Number of coordinator activities that were aborted so far in the current unit of work of this workload occurrence. This statistic is updated every time that an activity in this workload occurrence is aborted and is reset at the beginning of each unit of work.
COORD_ACT_REJECTED_TOTAL	INTEGER	Number of coordinator activities that were rejected so far in the current unit of work of this workload occurrence. Activities are counted as rejected when they are prevented from executing by either a prevent execution work action or a predictive threshold. This statistic is updated every time that an activity in this workload occurrence is rejected and is reset at the beginning of each unit of work.
CONCURRENT_ACT_TOP	INTEGER	Highest number of concurrent activities at any nesting level in either executing state (which includes idle and waiting) or queued state that has been reached for this workload occurrence in the current unit of work. This statistic is reset at the beginning of each unit of work.
ADDRESS	VARCHAR(255)	IP address or secure domain name that created this workload occurrence. Secure domain names are shown converted to IP addresses.

WLM_GET_SERVICE_SUBCLASS_STATS_V97 table function - Return statistics of service subclasses

The WLM_GET_SERVICE_SUBCLASS_STATS_V97 function returns basic statistics for one or more service subclasses.

Syntax

```

▶▶—WLM_GET_SERVICE_SUBCLASS_STATS_V97—(—service_superclass_name—, —————▶
▶—service_subclass_name—, —dbpartitionnum—)—————▶▶

```

The schema is SYSPROC.

Table function parameters

service_superclass_name

An input argument of type VARCHAR(128) that specifies the name of a service superclass in the currently connected database. If the argument is null or an empty string, the data is retrieved for all of the superclasses in the database.

service_subclass_name

An input argument of type VARCHAR(128) that specifies the name of a service subclass in the currently connected database. If the argument is null or an empty string, the data is retrieved for all of the subclasses in the database.

dbpartitionnum

An input argument of type INTEGER that specifies a valid partition number in the same instance as the currently connected database. Specify -1 for the current database partition, or -2 for all database partitions. If the null value is specified, -1 is set implicitly.

Authorization

EXECUTE privilege on the WLM_GET_SERVICE_SUBCLASS_STATS_V97 function.

Examples

Example 1: Because every activity must be mapped to a DB2 service class before being run, you can monitor the global state of the system by using the service class statistics table functions and querying all of the service classes on all partitions. In the following example, a null value is passed for *service_superclass_name* and *service_subclass_name* to return statistics for all service classes, and the value -2 is specified for *dbpartitionnum* to return statistics for all partitions:

```
SELECT SUBSTR(SERVICE_SUPERCLASS_NAME,1,19) AS SUPERCLASS_NAME,
       SUBSTR(SERVICE_SUBCLASS_NAME,1,18) AS SUBCLASS_NAME,
       SUBSTR(CHAR(DBPARTITIONNUM),1,4) AS PART,
       CAST(COORD_ACT_LIFETIME_AVG / 1000 AS DECIMAL(9,3))
       AS AVGLIFETIME,
       CAST(COORD_ACT_LIFETIME_STDDEV / 1000 AS DECIMAL(9,3))
       AS STDDEVLIFETIME,
       SUBSTR(CAST(LAST_RESET AS VARCHAR(30)),1,16) AS LAST_RESET
FROM TABLE(WLM_GET_SERVICE_SUBCLASS_STATS_V97(CAST(NULL AS VARCHAR(128)),
        CAST(NULL AS VARCHAR(128)), -2)) AS SCSTATS
ORDER BY SUPERCLASS_NAME, SUBCLASS_NAME, PART
```

The statement returns service class statistics such as average activity lifetime and standard deviation in seconds, as shown in the following sample output:

```
SUPERCLASS_NAME  SUBCLASS_NAME  PART  ...
-----
SYSDEFAULTUSERCLASS SYSDEFAULTSUBCLASS 0  ...
SYSDEFAULTUSERCLASS SYSDEFAULTSUBCLASS 1  ...
SYSDEFAULTUSERCLASS SYSDEFAULTSUBCLASS 2  ...
SYSDEFAULTUSERCLASS SYSDEFAULTSUBCLASS 3  ...
... AVGLIFETIME STDDEVLIFETIME LAST_RESET
... -----
...      691.242          34.322 2006-07-24-11.44
...      644.740          22.124 2006-07-24-11.44
...      612.431          43.347 2006-07-24-11.44
...      593.451          28.329 2006-07-24-11.44
```

Example 2: The same table function can also give the highest value for average concurrency of coordinator activities running in the service class on each partition:

```
SELECT SUBSTR(SERVICE_SUPERCLASS_NAME,1,19) AS SUPERCLASS_NAME,
       SUBSTR(SERVICE_SUBCLASS_NAME,1,18) AS SUBCLASS_NAME,
       SUBSTR(CHAR(DBPARTITIONNUM),1,4) AS PART,
       CONCURRENT_ACT_TOP AS ACTTOP,
       CONCURRENT_WLO_TOP AS CONNTOP
FROM TABLE(WLM_GET_SERVICE_SUBCLASS_STATS_V97(CAST(NULL AS VARCHAR(128)),
        CAST(NULL AS VARCHAR(128)), -2)) AS SCSTATS
ORDER BY SUPERCLASS_NAME, SUBCLASS_NAME, PART
```

Sample output is as follows:

SUPERCLASS_NAME	SUBCLASS_NAME	PART	ACTTOP	CONNTOP
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	0	10	7
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	1	0	0
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	2	0	0
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	3	0	0

By checking the average execution times and numbers of activities in the output of this table function, you can get a good high-level view of the load on each partition for a specific database. Any significant variations in the high-level gauges returned by this table function might indicate a change in the load on the system.

Example 3: If an activity uses thresholds with REMAP ACTIVITY TO actions, the activity might spend time in more than one service class during its lifetime. You can determine how many activities have passed through a set of service classes by looking at the ACTIVITIES_MAPPED_IN and ACTIVITIES_MAPPED_OUT columns, as shown in the following example:

```
SELECT SUBSTR(SERVICE_SUPERCLASS_NAME,1,19) AS SUPERCLASS_NAME,
       SUBSTR(SERVICE_SUBCLASS_NAME,1,18) AS SUBCLASS_NAME,
       ACTIVITIES_MAPPED_IN AS MAPPED_IN,
       ACTIVITIES_MAPPED_OUT AS MAPPED_OUT
FROM TABLE(WLM_GET_SERVICE_SUBCLASS_STATS_V97(CAST(NULL AS VARCHAR(128)),
        CAST(NULL AS VARCHAR(128)), -2)) AS SCSTATS
ORDER BY SUPERCLASS_NAME, SUBCLASS_NAME
```

Sample output is as follows:

SUPERCLASS_NAME	SUBCLASS_NAME	MAPPED_IN	MAPPED_OUT
SYSDEFAULTUSERCLASS	SYSDEFAULTSUBCLASS	0	0
SUPERCLASS1	SYSDEFAULTSUBCLASS	0	0
SUPERCLASS1	SUBCLASS1	0	7
SUPERCLASS1	SUBCLASS2	7	0

Usage notes

Some statistics are returned only if you set the COLLECT AGGREGATE ACTIVITY DATA and COLLECT AGGREGATE REQUEST DATA parameters for the corresponding service subclass to a value other than NONE.

The WLM_GET_SERVICE_SUBCLASS_STATS_V97 table function returns one row of data per service subclass and per partition. The function does not aggregate data across service classes (on a partition) or across partitions (for one or more service classes). However, you can use SQL queries to aggregate data.

The parameters are, in effect, ANDed together. That is, if you specify conflicting input parameters, such as a superclass named SUPA and a subclass named SUBB such that SUBB is not a subclass of SUPA, no rows are returned.

Information returned

Table 67. Information returned for WLM_GET_SERVICE_SUBCLASS_STATS_V97

Column name	Data type	Description
SERVICE_SUPERCLASS_NAME	VARCHAR(128)	Name of the service superclass from which this record was collected.
SERVICE_SUBCLASS_NAME	VARCHAR(128)	Name of the service subclass from which this record was collected.

Table 67. Information returned for WLM_GET_SERVICE_SUBCLASS_STATS_V97 (continued)

Column name	Data type	Description
DBPARTITIONNUM	SMALLINT	Partition number from which this record was collected.
LAST_RESET	TIMESTAMP	Time when statistics were last reset. There are four events that trigger a reset of statistics: <ul style="list-style-type: none"> You call the WLM_COLLECT_STATS procedure. The wlm_collect_int configuration parameter causes a collection and reset. You reactivate the database. You modify the service subclass for which statistics are being reported and commit the change. The LAST_RESET time stamp is in local time.
COORD_ACT_COMPLETED_TOTAL	BIGINT	The total number of coordinator activities that were submitted since the last reset and that were completed successfully. This count is updated as each activity is completed. If you remap an activity to a different service subclass, that activity counts only toward the total of the subclass in which it is completed.
COORD_ACT_ABORTED_TOTAL	BIGINT	The total number of coordinator activities that were submitted since the last reset and that were completed with errors. This count is updated as each activity aborts. If you remap an activity to a different service subclass, that activity counts only toward the total of the subclass in which it aborts.
COORD_ACT_REJECTED_TOTAL	BIGINT	The total number of coordinator activities that were submitted since the last reset and that were rejected before execution. Activities are counted as rejected when they are prevented from running by either a prevent execution work action or a predictive threshold. This count is updated as each activity is rejected.
CONCURRENT_ACT_TOP	INTEGER	Highest number of concurrent activities at any nesting level in executing state (which includes idle and waiting) that has been reached for this service subclass.

Table 67. Information returned for WLM_GET_SERVICE_SUBCLASS_STATS_V97 (continued)

Column name	Data type	Description
COORD_ACT_LIFETIME_TOP	BIGINT	<p>High watermark for coordinator activity lifetime, evaluated over all nesting levels. If the COLLECT AGGREGATE ACTIVITY DATA parameter of the service class is set to NONE, the value of the column is null. Units are milliseconds.</p> <p>To use this statistic effectively when the service class includes remapped subclasses, you must aggregate the COORD_ACT_LIFETIME_TOP high watermark of the service subclass with that of other subclasses affected by the same remapping threshold or thresholds. You must aggregate these values because an activity can be completed after the subclass has been remapped to a different service subclass. The time that the activity spends in other service subclasses before being remapped is counted only toward the service class in which it is completed.</p>
COORD_ACT_LIFETIME_AVG	DOUBLE	<p>Arithmetic mean of lifetime for coordinator activities at nesting level 0 that were associated with this service subclass since the last reset. If the internally tracked average has overflowed, the value -2 is returned. If the COLLECT AGGREGATE ACTIVITY DATA parameter of the service class is set to NONE, the value of the column is null. Units are milliseconds.</p> <p>The COORD_ACT_LIFETIME_AVG value of a service subclass is unaffected by activities that pass through the subclass but are remapped to a different subclass before they are completed.</p>
COORD_ACT_LIFETIME_STDDEV	DOUBLE	<p>Standard deviation of lifetime for coordinator activities at nesting level 0 that were associated with this service subclass since the last reset. If the COLLECT AGGREGATE ACTIVITY DATA parameter of the service class is set to NONE, the value of the column is null. Units are milliseconds.</p> <p>This standard deviation is computed from the coordinator activity lifetime histogram and may be inaccurate if the histogram is not correctly sized to fit the data. The value of -1 is returned if any values fall into the last histogram bin.</p> <p>The COORD_ACT_LIFETIME_STDDEV value of a service subclass is unaffected by activities that pass through the service subclass but are remapped to a different subclass before they are completed.</p>

Table 67. Information returned for WLM_GET_SERVICE_SUBCLASS_STATS_V97 (continued)

Column name	Data type	Description
COORD_ACT_EXEC_TIME_AVG	DOUBLE	<p>Arithmetic mean of the execution times for coordinator activities at nesting level 0 that were associated with this service subclass since the last reset. If the internally tracked average has overflowed, the value -2 is returned. If the COLLECT AGGREGATE ACTIVITY DATA parameter of the service class is set to NONE, the value of the column is null. Units are milliseconds.</p> <p>The execution time average of a service subclass is unaffected by activities that pass through the subclass but are remapped to a different subclass before they are completed.</p>
COORD_ACT_EXEC_TIME_STDDEV	DOUBLE	<p>Standard deviation of the execution times for coordinator activities at nesting level 0 that were associated with this service subclass since the last reset. Units are milliseconds.</p> <p>This standard deviation is computed from the coordinator activity executetime histogram and may be inaccurate if the histogram is not correctly sized to fit the data. The value of -1 is returned if any values fall into the last histogram bin.</p> <p>The execution time standard deviation of a service subclass is unaffected by activities that pass through the subclass but are remapped to a different subclass before they are completed.</p>
COORD_ACT_QUEUE_TIME_AVG	DOUBLE	<p>Arithmetic mean of the queue time for coordinator activities at nesting level 0 that were associated with this service subclass since the last reset. If the internally tracked average has overflowed, the value -2 is returned. If the COLLECT AGGREGATE ACTIVITY DATA parameter of the service class is set to NONE, the value of the column is null. Units are milliseconds.</p> <p>The queue time average is counted only toward the service subclass in which the activity was queued.</p>

Table 67. Information returned for WLM_GET_SERVICE_SUBCLASS_STATS_V97 (continued)

Column name	Data type	Description
COORD_ACT_QUEUE_TIME_STDDEV	DOUBLE	<p>Standard deviation of the queue time for coordinator activities at nesting level 0 that were associated with this service subclass since the last reset. If the COLLECT AGGREGATE ACTIVITY DATA parameter of the service class is set to NONE, the value of the column is null. Units are milliseconds.</p> <p>This standard deviation is computed from the coordinator activity queuetime histogram and may be inaccurate if the histogram is not correctly sized to fit the data. The value of -1 is returned if any values fall into the last histogram bin.</p> <p>The queue time standard deviation is counted only toward the service subclass in which the activity was queued.</p>
NUM_REQUESTS_ACTIVE	BIGINT	Number of requests that are running in the service subclass at the time that this table function is running.
NUM_REQUESTS_TOTAL	BIGINT	<p>Number of requests that finished running in this service subclass since the last reset. This finished state applies to any request regardless of its membership in an activity. If the COLLECT AGGREGATE ACTIVITY DATA parameter of the service class is set to NONE, the value of the column is null.</p> <p>The NUM_REQUESTS_TOTAL value of a service subclass is unaffected by requests that pass through the service subclass but are not completed in it.</p>
REQUEST_EXEC_TIME_AVG	DOUBLE	<p>Arithmetic mean of the execution times for requests that were associated with this service subclass since the last reset. Units are milliseconds. If the internally tracked average has overflowed, the value -2 is returned. If the COLLECT AGGREGATE REQUEST DATA parameter of this service class is set to NONE, the value of this column is NULL.</p> <p>The execution time average of a service subclass is unaffected by requests that pass through the subclass but are not completed in it.</p>

Table 67. Information returned for WLM_GET_SERVICE_SUBCLASS_STATS_V97 (continued)

Column name	Data type	Description
REQUEST_EXEC_TIME_STDDEV	DOUBLE	<p>Standard deviation of the execution times for requests that were associated with this service subclass since the last reset. Units are milliseconds. If the COLLECT AGGREGATE REQUEST DATA parameter of the service class is set to NONE, the value of this column is NULL.</p> <p>This standard deviation is computed from the request executetime histogram and may be inaccurate if the histogram is not correctly sized to fit the data. The value of -1 is returned if any values fall into the last histogram bin.</p> <p>The execution time standard deviation of a service subclass is unaffected by requests that pass through the subclass but were not completed in it.</p>
REQUEST_EXEC_TIME_TOTAL	BIGINT	<p>Sum of the execution times for requests that were associated with this service subclass since the last reset. Units are milliseconds. If the COLLECT AGGREGATE REQUEST DATA parameter of the service class is set to NONE, the value of this column is NULL.</p> <p>This total is computed from the request execution time histogram and may be inaccurate if the histogram is not correctly sized to fit the data. The value of -1 is returned if any values fall into the last histogram bin.</p> <p>The execution time total of a service subclass is unaffected by requests that pass through the subclass but are not completed in it.</p>
ACT_REMAPPED_IN	BIGINT	Number of activities remapped into this service subclass by a threshold REMAP ACTIVITY action since the last reset.
ACT_REMAPPED_OUT	BIGINT	Number of activities remapped out of this service subclass by a threshold REMAP ACTIVITY action since the last reset.
CONCURRENT_WLO_TOP	INTEGER	Highest number of concurrent occurrences of the specified workload on this partition since the last reset.
UOW_TOTAL_TIME_TOP	BIGINT	Reserved for future use.

WLM_GET_SERVICE_SUPERCLASS_STATS - Return statistics of service superclasses

The WLM_GET_SERVICE_SUPERCLASS_STATS function returns basic statistics for one or more service superclasses.

Syntax

```
►►—WLM_GET_SERVICE_SUPERCLASS_STATS—(—service_superclass_name—,—————►  
►—dbpartitionnum—)—————►
```

The schema is SYSPROC.

Table function parameters

service_superclass_name

An input argument of type VARCHAR(128) that specifies the name of a service superclass in the currently connected database. If the argument is null or an empty string, data is retrieved for all the superclasses in the database.

dbpartitionnum

An input argument of type INTEGER that specifies a valid partition number in the same instance as the currently connected database. Specify -1 for the current database partition, or -2 for all database partitions. If the null value is specified, -1 is set implicitly.

Authorization

EXECUTE privilege on the WLM_GET_SERVICE_SUPERCLASS_STATS function.

Example

The following query displays the basic statistics for all the service superclasses on the system, across all database partitions:

```
SELECT SUBSTR(SERVICE_SUPERCLASS_NAME, 1, 26) SERVICE_SUPERCLASS_NAME,  
       DBPARTITIONNUM,  
       LAST_RESET,  
       CONCURRENT_CONNECTION_TOP CONCURRENT_CONN_TOP  
FROM TABLE(WLM_GET_SERVICE_SUPERCLASS_STATS(' ', -2)) as SCSTATS
```

Sample output is as follows:

SERVICE_SUPERCLASS_NAME	DBPARTITIONNUM	...
-----	-----	...
SYSDEFAULTSYSTEMCLASS	0	...
SYSDEFAULTMAINTENANCECLASS	0	...
SYSDEFAULTUSERCLASS	0	...
...	LAST_RESET	CONCURRENT_CONN_TOP
...	-----	-----
...	2006-09-05-09.38.44.396788	0
...	2006-09-05-09.38.44.396795	0
...	2006-09-05-09.38.44.396796	1

Usage note

The WLM_GET_SERVICE_SUPERCLASS_STATS table function returns one row of data per service superclass and per partition. The function does not aggregate data across service superclasses (on a partition) or across partitions (for one or more service superclasses). However, you can use SQL queries to aggregate data, as shown in the previous example.

Information returned

Table 68. Information returned for WLM_GET_SERVICE_SUPERCLASS_STATS

Column name	Data type	Description
SERVICE_SUPERCLASS_NAME	VARCHAR(128)	Name of the service superclass from which this record was collected.
DBPARTITIONNUM	SMALLINT	Partition number from which this record was collected.
LAST_RESET	TIMESTAMP	Time when statistics were last reset. There are four events that trigger a reset of statistics: <ul style="list-style-type: none">• You call the WLM_COLLECT_STATS procedure.• The wlm_collect_int configuration parameter causes a collection and reset.• You reactivate the database.• You modify the service superclass for which statistics are being reported and commit the change. The LAST_RESET time stamp is in local time.
CONCURRENT_CONNECTION_TOP	INTEGER	Highest number of concurrent coordinator connections in this class since the last reset.

WLM_GET_WORK_ACTION_SET_STATS - Return work action set statistics

The WLM_GET_WORK_ACTION_SET_STATS function returns the statistics for a work action set.

Syntax

```
►►—WLM_GET_WORK_ACTION_SET_STATS—(—work_action_set_name—, —————►  
►—dbpartitionnum—)—————►►
```

The schema is SYSPROC.

Table function parameters

work_action_set_name

An input argument of type VARCHAR(128) that specifies the work action set to return statistics for. If the argument is null or an empty string, statistics are returned for all work action sets.

dbpartitionnum

An input argument of type INTEGER that specifies a valid partition number in the same instance as the currently connected database. Specify -1 for the current database partition, or -2 for all database partitions. If the null value is specified, -1 is set implicitly.

Authorization

EXECUTE privilege on the WLM_GET_WORK_ACTION_SET_STATS function.

Example

Assume that there are three work classes: ReadClass, WriteClass, and LoadClass. There is a work action associated with ReadClass and a work action associated with LoadClass, but there is no work action associated with WriteClass. On partition 0, the number of activities currently running or queued are as follows:

- ReadClass class: eight
- WriteClass class: four
- LoadClass class: two
- Unassigned: three

```
ELECT SUBSTR(WORK_ACTION_SET_NAME,1,18) AS WORK_ACTION_SET_NAME,
       SUBSTR(CHAR(DBPARTITIONNUM),1,4) AS PART,
       SUBSTR(WORK_CLASS_NAME,1,15) AS WORK_CLASS_NAME,
       LAST_RESET,
       SUBSTR(CHAR(WLO_ACT_TOTAL),1,14) AS ACT_TOTAL
FROM TABLE(WLM_GET_WORK_ACTION_SET_STATS
(CAST(NULL AS VARCHAR(128)), -2)) AS WASSTATS
ORDER BY WORK_ACTION_SET_NAME, WORK_CLASS_NAME, PART
```

Sample output is as follows. Because there is no work action associated with the WriteClass work class, the four activities to which it applies are counted in the artificial class denoted by an asterisk (*) in the output. The three activities that were not assigned to any work class are also included in the artificial class.

WORK_ACTION_SET_NAME	PART	WORK_CLASS_NAME	LAST_RESET	ACT_TOTAL
AdminActionSet	0	ReadClass	2005-11-25-18.52.49.343000	8
AdminActionSet	1	ReadClass	2005-11-25-18.52.50.478000	0
AdminActionSet	0	LoadClass	2005-11-25-18.52.49.343000	2
AdminActionSet	1	LoadClass	2005-11-25-18.52.50.478000	0
AdminActionSet	0	*	2005-11-25-18.52.49.343000	7
AdminActionSet	1	*	2005-11-25-18.52.50.478000	0

Information returned

Table 69. Information returned for WLM_GET_WORK_ACTION_SET_STATS

Column name	Data type	Description
WORK_ACTION_SET_NAME	VARCHAR(128)	Name of the work action set. A name is returned only if you enable the work action set.
DBPARTITIONNUM	SMALLINT	Partition number from which this record was collected.
LAST_RESET	TIMESTAMP	Time when statistics were last reset. There are four events that trigger a reset of statistics: <ul style="list-style-type: none"> • You call the WLM_COLLECT_STATS procedure. • The wlm_collect_int configuration parameter causes a collection and reset. • You reactivate the database. • You modify the work action set for which statistics are being reported and commit the change. The LAST_RESET time stamp is in local time.
WORK_CLASS_NAME	VARCHAR(128)	Name of the work class related to the specified work action set. A work class name is returned only if you enable a work action associated with the work class. The asterisk (*) represents an artificial work class created to count all those activities that did not belong to the other work classes for which you associated one or more work actions.

Table 69. Information returned for WLM_GET_WORK_ACTION_SET_STATS (continued)

Column name	Data type	Description
ACT_TOTAL	BIGINT	Number of activities at any nesting level that were assigned to the work class specified by WORK_CLASS_NAME.

WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97 - Return a list of activities

The WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97 function returns the list of all activities that were submitted by the specified application on the specified partition and have not yet been completed.

Syntax

```
►►—WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97—(—application_handle—,—————►
►—dbpartitionnum—)—————►
```

The schema is SYSPROC.

Table function parameters

application_handle

An input argument of type BIGINT that specifies an application handle for which a list of activities is to be returned. If the argument is null, the data is retrieved for all the applications in the database.

dbpartitionnum

An input argument of type INTEGER that specifies a valid partition number in the same instance as the currently connected database. Specify -1 for the current database partition, or -2 for all database partitions. If the null value is specified, -1 is set implicitly.

Authorization

EXECUTE privilege on the WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97 function.

Examples

Example 1: Activities currently running with a known application handle

After you identify the application handle, you can look up all the activities currently running in this application. For example, suppose that an administrator wants to list the activities of an application whose application handle, determined by using the LIST APPLICATIONS command, is 1. The administrator runs the following query:

```
SELECT SUBSTR(CHAR(COORD_PARTITION_NUM),1,5) AS COORD,
       SUBSTR(CHAR(DBPARTITIONNUM),1,4) AS PART,
       SUBSTR(CHAR(UOW_ID),1,5) AS UOWID,
       SUBSTR(CHAR(ACTIVITY_ID),1,5) AS ACTID,
       SUBSTR(CHAR(PARENT_UOW_ID),1,8) AS PARUOWID,
       SUBSTR(CHAR(PARENT_ACTIVITY_ID),1,8) AS PARACTID,
```

```

ACTIVITY_TYPE AS ACTTYPE,
SUBSTR(CHAR(NESTING_LEVEL),1,7) AS NESTING
FROM TABLE(WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97(1, -2)) AS WLOACTS
ORDER BY PART, UOWID, ACTID

```

Sample output from the query is as follows:

```

COORD PART UOWID ACTID PARUOWID PARACTID ACTTYPE NESTING
-----
0 0 2 3 - - CALL 0
0 0 2 5 2 3 READ_DML 1

```

Example 2: Activities currently running on the system

The following query joins the WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97 output with the MON_GET_PKG_CACHE_STMT output on EXECUTABLE_ID to provide statement text for all the SQL activities currently running on the system:

```

SELECT t.application_handle,
       t.uow_id,
       t.activity_id,
       varchar(p.stmt_text, 256) as stmt_text
FROM table(wlm_get_workload_occurrence_activities_v97(NULL, -1)) as t,
     table(mon_get_pkg_cache_stmt(NULL, NULL, NULL, -1)) as p
WHERE t.executable_id = p.executable_id

```

Sample output is as follows:

```

APPLICATION_HANDLE  UOW_ID  ACTIVITY_ID  STMT_TEXT
-----
1 1 1 SELECT * FROM SYSCAT.TABLES
47 1 36 INSERT INTO T1 VALUES(123)

```

Information returned

Table 70. Information returned by WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97

Column name	Data type	Description
APPLICATION_HANDLE	BIGINT	System-wide unique ID for the application. On a single-partitioned database, this identifier consists of a 16-bit counter. On a multi-partitioned database, this identifier consists of the coordinating partition number concatenated with a 16-bit counter. In addition, this identifier is the same on every partition where the application makes a secondary connection.
DBPARTITIONNUM	SMALLINT	Partition number from which this record was collected.
COORD_PARTITION_NUM	SMALLINT	Coordinator partition number of the activity.
LOCAL_START_TIME	TIMESTAMP	Local time that this activity began doing work on the partition. The value of the column is null when an activity has entered the system but is in a queue and has not started running.
UOW_ID	INTEGER	Unique ID of the original unit of work that the activity started in.
ACTIVITY_ID	INTEGER	Unique activity ID within a unit of work.

Table 70. Information returned by WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97 (continued)

Column name	Data type	Description
PARENT_UOW_ID	INTEGER	Unique ID of the original unit of work that the parent activity of the activity started in. The value of this column is null if the activity has no parent activity or is at a remote partition.
PARENT_ACTIVITY_ID	INTEGER	Unique activity ID within a unit of work for the parent of the activity whose ID is specified by ACTIVITY_ID. The value of this column is null if the activity has no parent activity or is at a remote partition.
ACTIVITY_STATE	VARCHAR(32)	Activity state. Possible values are as follows: <p><i>CANCEL_PENDING</i> The activity was cancelled because there was no agent actively working on a request for the activity. The next time that a request is submitted as part of the activity, the activity will be cancelled, and an SQL4725N error will be generated.</p> <p><i>EXECUTING</i> Agents are actively working on a request for the activity.</p> <p><i>IDLE</i> There is no agent actively processing a request for the activity.</p> <p><i>INITIALIZING</i> The activity has been submitted but has not yet started running. During the initializing state, predictive thresholds are applied to the activity to determine whether the activity will be allowed to run.</p>

Table 70. Information returned by WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97 (continued)

Column name	Data type	Description
ACTIVITY_STATE (continued)	VARCHAR(32)	<p>Activity state. Possible values are as follows:</p> <p><i>QP_CANCEL_PENDING</i> This state is the same as the CANCEL_PENDING state except that the activity was cancelled by Query Patroller rather than by the WLM_CANCEL_ACTIVITY procedure.</p> <p><i>QP_QUEUED</i> The activity is queued by Query Patroller.</p> <p><i>QUEUED</i> The activity is queued by a workload management queuing threshold. In a database partitioning feature (DPF) environment, this state might mean that the coordinator agent has made an RPC to the catalog partition to obtain threshold tickets and has not yet received a response. This state might indicate that the activity has been queued by a workload management queuing threshold or, if not much time has elapsed, can indicate that the activity is in the process of obtaining its tickets. To obtain a more accurate picture of whether the activity is being queued, determine what agent is working on the activity, and find out whether the EVENT_OBJECT value of the object at the catalog partition has a value of WLM_QUEUE.</p> <p><i>TERMINATING</i> The activity has finished running and is being removed from the system.</p>
ACTIVITY_TYPE	VARCHAR(32)	<p>Activity type. Possible values are as follows:</p> <ul style="list-style-type: none"> • CALL • DDL • LOAD • OTHER • READ_DML • WRITE_DML <p>Refer to "Identify types of work with work classes" in <i>Workload Manager Guide and Reference</i> for a description of the different types of SQL statements that are associated with each activity type.</p>

Table 70. Information returned by WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97 (continued)

Column name	Data type	Description
NESTING_LEVEL	INTEGER	Depth to which this activity is nested within its topmost parent activity.
INVOCATION_ID	INTEGER	Invocation ID, which distinguishes one particular invocation of this activity from others at the same nesting level.
ROUTINE_ID	INTEGER	Unique ID of the routine.
UTILITY_ID	INTEGER	One of the following values: <ul style="list-style-type: none"> • If the activity is a utility, the value is the ID of the utility. • If the activity is not a utility, the value is null.
SERVICE_CLASS_ID	INTEGER	Unique ID of the service class to which this activity belongs.
DATABASE_WORK_ACTION_SET_ID	INTEGER	One of the following values: <ul style="list-style-type: none"> • If this activity has been categorized into a work class of database scope, the value is the ID of the work class set of which this work class is a member. • If this activity has not been categorized into a work class of database scope, the value is null.
DATABASE_WORK_CLASS_ID	INTEGER	One of the following values: <ul style="list-style-type: none"> • If this activity has been categorized into a work class of database scope, the value is the ID of the work class. • If this activity has not been categorized into a work class of database scope, the value is null.
SERVICE_CLASS_WORK_ACTION_SET_ID	INTEGER	One of the following values: <ul style="list-style-type: none"> • If this activity has been categorized into a work class of service class scope, the value is the ID of the work action set associated with the work class set to which the work class belongs. • If this activity has not been categorized into a work class of service class scope, the value is null.
SERVICE_CLASS_WORK_CLASS_ID	INTEGER	One of the following values: <ul style="list-style-type: none"> • If this activity has been categorized into a work class of service class scope, the value is the ID of the work class assigned to this activity. • If this activity has not been categorized into a work class of service class scope, the value is null.

Table 70. Information returned by WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97 (continued)

Column name	Data type	Description
EXECUTABLE_ID	VARCHAR(32) FOR BIT DATA	An opaque binary token generated on the data server that uniquely identifies the section. You can use the executable ID as input to different monitoring interfaces to obtain data about the section. For non-SQL activities, such as LOAD, a NULL value is returned.
TOTAL_CPU_TIME	BIGINT	Reserved for future use.
ROWS_READ	BIGINT	Reserved for future use.
ROWS_RETURNED	BIGINT	Reserved for future use.
QUERY_COST_ESTIMATE	BIGINT	Reserved for future use.
DIRECT_READS	BIGINT	Reserved for future use.
DIRECT_WRITES	BIGINT	Reserved for future use.

WLM_GET_WORKLOAD_STATS_V97 table function - Return workload statistics

The WLM_GET_WORKLOAD_STATS_V97 function returns one row of workload statistics for every combination of workload name and database partition number.

Syntax

►► WLM_GET_WORKLOAD_STATS_V97 (—workload_name—, —dbpartitionnum—) ◀◀

The schema is SYSPROC.

Table function parameters

workload_name

An input argument of type VARCHAR(128) that specifies a workload for which the statistics are to be returned. If the argument is NULL or an empty string, statistics are returned for all workloads.

dbpartitionnum

An input argument of type INTEGER that specifies the number of a partition in the same instance as the currently connected database. Specify -1 for the current database partition, or -2 for all database partitions. If a null value is specified, -1 is set implicitly.

Authorization

EXECUTE privilege on the WLM_GET_WORKLOAD_STATS_V97 function.

Example

The following query displays statistics for workloads:

```
SELECT SUBSTR(WORKLOAD_NAME,1,18) AS WL_DEF_NAME,
       SUBSTR(CHAR(DBPARTITIONNUM),1,4) AS PART,
       COORD_ACT_LIFETIME_TOP,
```

```

        COORD_ACT_LIFETIME_AVG,
        COORD_ACT_LIFETIME_STDDEV
FROM TABLE(WLM_GET_WORKLOAD_STATS_V97(CAST(NULL AS VARCHAR(128)), -2)) AS WLSTATS
ORDER BY WL_DEF_NAME, PART

```

Sample output from the query is as follows:

```

WL_DEF_NAME          PART COORD_ACT_LIFETIME_TOP ...
-----
SYSDEFAULTADMWORKL  0                -1 ...
SYSDEFAULTUSERWORK  0                -1 ...
WL1                  0                2 ...
... COORD_ACT_LIFETIME_AVG  COORD_ACT_LIFETIME_STDDEV
... -----
... -1.000000000000000E+000  -1.000000000000000E+000
... -1.000000000000000E+000  -1.000000000000000E+000
... +2.560000000000000E+000  +6.000000000000001E-002

```

Usage note

The function does not aggregate data across workloads, partitions, or service classes. However, you can use SQL queries to aggregate data.

Information returned

Table 71. Information returned by WLM_GET_WORKLOAD_STATS_V97

Column name	Data type	Description
WORKLOAD_NAME	VARCHAR(128)	Name of the workload from which this record was collected.
DBPARTITIONNUM	SMALLINT	Partition number from which this record was collected
LAST_RESET	TIMESTAMP	Time when statistics were last reset. There are four events that trigger a reset of statistics: <ul style="list-style-type: none"> You call the WLM_COLLECT_STATS procedure. The <code>wlm_collect_int</code> configuration parameter causes a collection and reset. You reactivate the database. You modify the workload for which statistics are being reported and commit the change. The LAST_RESET timestamp is in local time.
CONCURRENT_WLO_TOP	INTEGER	Highest number of concurrent occurrences of the specified workload on this partition since the last reset.
CONCURRENT_WLO_ACT_TOP	INTEGER	Highest number of concurrent activities (both coordinator and nested) in either executing state (which includes idle and waiting) or queued state that has been reached in any occurrence of this workload since the last reset. The value of the column is updated by each workload occurrence at the end of its unit of work.
COORD_ACT_COMPLETED_TOTAL	BIGINT	Total number of coordinator activities at any nesting level that were assigned to any occurrence of this workload that were completed since the last reset. The value of this column is updated by each workload occurrence at the end of its unit of work.

Table 71. Information returned by WLM_GET_WORKLOAD_STATS_V97 (continued)

Column name	Data type	Description
COORD_ACT_ABORTED_TOTAL	BIGINT	The total number of coordinator activities at any nesting level that were assigned to any occurrence of this workload that were aborted before completion since the last reset. The value of this column is updated by each workload occurrence at the end of its unit of work.
COORD_ACT_REJECTED_TOTAL	BIGINT	The total number of coordinator activities at any nesting level that were assigned to any occurrence of this workload that were rejected before execution since the last reset. The value of this column is updated by each workload occurrence at the end of its unit of work. Activities are counted as rejected when they are prevented from executing by either a prevent execution work action or a predictive threshold. Unlike the column of the same name in the WLM_GET_SERVICE_SUBCLASS_STATS_V97 function, this WLM_GET_WORKLOAD_STATS_V97 column also includes the number of rejections that occur before an activity can be assigned to a service class. For example, such a rejection occurs when an activity violates the ConcurrentWorkloadOccurrences threshold.
WLO_COMPLETED_TOTAL	BIGINT	Number of workload occurrences to be completed since last reset.
COORD_ACT_LIFETIME_TOP	BIGINT	High watermark for coordinator activity lifetime, collected over all nesting levels. Units are milliseconds. If the COLLECT AGGREGATE ACTIVITY DATA parameter of the service class is set to NONE, the value of the column is null.
COORD_ACT_LIFETIME_AVG	DOUBLE	Arithmetic mean of lifetime for completed or aborted coordinator activities at nesting level 0 that are associated with this workload. Units are milliseconds. If the internally tracked average has overflowed, the value -2 is returned. If the COLLECT AGGREGATE ACTIVITY DATA parameter of the workload is set to NONE, the value of the column is null.
COORD_ACT_LIFETIME_STDDEV	DOUBLE	Standard deviation of lifetime for completed or aborted coordinator activities at nesting level 0 that are associated with this workload. Units are milliseconds. If the COLLECT AGGREGATE ACTIVITY DATA parameter of the workload is set to NONE, the value of the column is null. This standard deviation is computed from the coordinator activity lifetime histogram and may be inaccurate if the histogram is not correctly sized to fit the data. If any values fall into the last histogram bin, the value -1 is returned.

Table 71. Information returned by WLM_GET_WORKLOAD_STATS_V97 (continued)

Column name	Data type	Description
COORD_ACT_EXEC_TIME_AVG	DOUBLE	Arithmetic mean of the execution times for completed or aborted coordinator activities at nesting level 0 that are associated with this workload. Units are milliseconds. If the internally tracked average has overflowed, the value -2 is returned. If the COLLECT AGGREGATE ACTIVITY DATA parameter of the workload is set to NONE, the value of the column is null.
COORD_ACT_EXEC_TIME_STDDEV	DOUBLE	Standard deviation of the execution times for completed or aborted coordinator activities at nesting level 0 that are associated with this workload. Units are milliseconds. This standard deviation is computed from the coordinator activity execution time histogram and may be inaccurate if the histogram is not correctly sized to fit the data. If any values fall into the last histogram bin, the value -1 is returned. If the COLLECT AGGREGATE ACTIVITY DATA parameter of the workload is set to NONE, the value of the column is null.
COORD_ACT_QUEUE_TIME_AVG	DOUBLE	Arithmetic mean of the queue time for completed or aborted coordinator activities at nesting level 0 that are associated with this workload. Units are milliseconds. If the internally tracked average has overflowed, the value -2 is returned. If the COLLECT AGGREGATE ACTIVITY DATA parameter of the service class is set to NONE, the value of the column is null.
COORD_ACT_QUEUE_TIME_STDDEV	DOUBLE	Standard deviation of the queue time for completed or aborted coordinator activities at nesting level 0 that are associated with this workload. Units are milliseconds. If the COLLECT AGGREGATE ACTIVITY DATA parameter of the workload is set to NONE, the value of the column is null. This standard deviation is computed from the coordinator activity queue time histogram and may be inaccurate if the histogram is not correctly sized to fit the data. If any values fall into the last histogram bin, the value -1 is returned.
UOW_TOTAL_TIME_TOP	BIGINT	Reserved for future use.

WLM_SET_CLIENT_INFO procedure - Set client information

The WLM_SET_CLIENT_INFO procedure sets client information associated with the current connection at the DB2 server.

By using this procedure, you can set the client's user ID, application name, workstation name, accounting information, or workload information at the DB2 server. Calling this procedure changes the stored values of the relevant transaction processor (TP) monitor client information fields and special register settings for this connection.

The client information fields are used at the DB2 server for determining the identity of the application or user currently using the connection. The client

information fields for a connection are considered during DB2 workload evaluation and also displayed in any DB2 audit records or application snapshots generated for this connection.

Unlike the `sqleseti` API, this procedure does not set client information at the client but instead sets the corresponding client attributes on the DB2 server. Therefore, you cannot use the `sqlqry` API to query the client information that is set at the DB2 server using this procedure.

The data values provided with the procedure are converted to the appropriate database code page before being stored in the related TP monitor fields or special registers. Any data value which exceeds the maximum supported size after conversion to the database code page is truncated before being stored at the server. The truncated values are returned by both the TP monitor fields and the special registers when those stored values are queried.

The `WLM_SET_CLIENT_INFO` procedure is not under transaction control, and client information changes made by the procedure are independent of committing or rolling back units of work. However, because workload reevaluation occurs at the beginning of the next unit of work for each application, you must issue either a `COMMIT` or a `ROLLBACK` statement to make client information changes effective.

If your applications must be portable across IBM® data servers, use one of the database application programming interfaces to set client information at the server instead of using the `WLM_SET_CLIENT_INFO` procedure. The procedure is available only for DB2 Database for Linux, UNIX, and Windows.

Syntax

```
▶▶ WLM_SET_CLIENT_INFO (—client_userid—, —client_wrkstname—, —————▶
▶ —client_applname—, —client_acctstr—, —client_workload—) —————▶▶
```

The schema is `SYSPROC`.

Procedure parameters

client_userid

An input argument of type `VARCHAR(255)` that specifies the user ID for the client. If you specify `NULL`, the value remains unchanged. If you specify an empty string, which is the default value, the user ID for the client is reset to the default value, which is blank.

client_wrkstname

An input argument of type `VARCHAR(255)` that specifies the workstation name for the client. If you specify `NULL`, the value remains unchanged. If you specify an empty string, which is the default value, the workstation name for the client is reset to the default value, which is blank.

client_applname

An input argument of type `VARCHAR(255)` that specifies the application name for the client. If you specify `NULL`, the value remains unchanged. If you specify an empty string, which is the default value, the application name for the client is reset to the default value, which is blank.

client_acctstr

An input argument of type `VARCHAR(255)` that specifies the accounting string

for the client. If you specify NULL, the value remains unchanged. If you specify an empty string, which is the default value, the accounting string for the client is reset to the default value, which is blank.

client_workload

An input argument of type VARCHAR(255) that specifies the workload assignment mode for the client. If you specify NULL, the value remains unchanged. The values are as follows:

SYSDEFAULTADMWORKLOAD

Specifies that the database connection will be assigned to SYSDEFAULTADMWORKLOAD, enabling users with ACCESSCTRL, DATAACCESS, DBADM, SECADM, or WLMADM authority to bypass the normal workload evaluation.

AUTOMATIC

Specifies that the database connection will be assigned to a workload chosen by the workload evaluation that is performed automatically by the server.

Note: The *client_workload* argument is case sensitive.

Authorization

EXECUTE privilege on the WLM_SET_CLIENT_INFO procedure.

Examples

The following procedure call sets the user ID, workstation name, application name, accounting string, and workload assignment mode for the client:

```
CALL SYSPROC.WLM_SET_CLIENT_INFO('db2user', 'machine.torolab.ibm.com',
    'auditor', 'Accounting department', 'AUTOMATIC')
```

The following procedure call sets the user ID to db2user2 for the client without setting the other client attributes:

```
CALL SYSPROC.WLM_SET_CLIENT_INFO('db2user2', NULL, NULL, NULL, NULL)
```

The following procedure call resets the user ID for the client to blank without modifying the values of the other client attributes:

```
CALL SYSPROC.WLM_SET_CLIENT_INFO('', NULL, NULL, NULL, NULL)
```

Workload management monitor elements

The following monitor elements provide information about activities, threshold violations, and workload management statistics.

activate_timestamp - Activate timestamp monitor element

The time when an event monitor was activated.

Table 72. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Activity	event_activity	-
Activity	event_activitystmt	-
Activity	event_activityvals	-
Threshold Violations	event_thresholdviolations	-

Usage

Use this element to correlate information returned by the above event types.

activity_collected - Activity collected monitor element

This element indicates whether or not activity event monitor records are to be collected for a violated threshold.

Table 73. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Threshold violations	event_thresholdviolations	-

Usage

Use this element to determine whether to expect an activity event for the activity that violated the threshold to be written to the activity event monitor.

When an activity finishes or aborts and the activity event monitor is active at the time, if the value of this monitor element is 'Y', the activity that violated this threshold will be collected. If the value of this monitor element is 'N', it will not be collected.

activity_id - Activity ID monitor element

Counter which uniquely identifies an activity for an application within a given unit of work.

Table 74. Table Function Monitoring Information

Table Function	Monitor Element Collection Level
MON_GET_ACTIVITY_DETAILS table function - Get complete activity details (reported in DETAILS XML document)	Always collected

Table 75. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Locking	-	-
Activities	event_activity	-
Activities	event_activitystmt	-
Activities	event_activityvals	-
Threshold violations	event_thresholdviolations	-

Usage

Use this element in conjunction with other activity history elements for analysis of the behavior of an activity.

To uniquely identify an activity outside its unit of work, use the combination of **activity_id** and **uow_id** plus one of the following: **appl_id** or **agent_id**.

activity_secondary_id - Activity secondary ID monitor element

The value for this element is incremented each time an activity record is written for the same activity. For example, if an activity record is written once as a result of having called the WLM_CAPTURE_ACTIVITY_IN_PROGRESS procedure and a second time when the activity ends, the element would have a value of 0 for the first record and 1 for the second record.

Table 76. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Activities	event_activity	-
Activities	event_activitystmt	-
Activities	event_activityvals	-

Usage

Use this element with **activity_id**, **uow_id**, and **appl_id** monitor elements to uniquely identify activity records when information about the same activity has been written to the activities event monitor multiple times.

For example, information about an activity would be sent to the activities event monitor twice in the following case:

- the WLM_CAPTURE_ACTIVITY_IN_PROGRESS stored procedure was used to capture information about the activity while it was running
- information about the activity was collected when the activity completed, because the COLLECT ACTIVITY DATA clause was specified on the service class with which the activity is associated

activity_type - Activity type monitor element

The type of the activity.

Table 77. Table Function Monitoring Information

Table Function	Monitor Element Collection Command and Level
MON_GET_ACTIVITY_DETAILS table function - Get complete activity details (reported in DETAILS XML document)	Always collected

Table 78. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Activities	event_activity	-

Usage

The possible values are:

- LOAD
- READ_DML
- WRITE_DML
- DDL
- CALL

- OTHER

act_cpu_time_top – Activity CPU time top monitor element

The high watermark for processor time used by activities at all nesting levels in a service class, workload, or work class.

The monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the service class or workload in which the activity runs is set to NONE. Activities contribute towards this high watermark only when request metrics are enabled.

For service classes, when you remap activities between service subclasses with a REMAP ACTIVITY action, only the act_cpu_time_top high watermark of the service subclass where an activity completes is updated, provided that a new high watermark is reached. The act_cpu_time_top high watermarks of other service subclasses an activity is mapped to but does not complete in are unaffected.

Table 79. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_scstats	-
Statistics	event_wcstats	-
Statistics	event_wlstats	-

Usage

Use this element to determine the highest amount of processor time used by an activity on a partition for a service class, workload, or work class during the time interval collected.

act_exec_time - Activity execution time monitor element

Time spent executing at this partition, in microseconds. For cursors, the execution time is the combined time for the open, the fetches, and the close. The time when the cursor is idle is not counted towards execution time. For routines, execution time is the start to end of routine invocation. The lifetimes of any cursors left open by routine (to return a result set) after the routine finishes are not counted towards the routine execution time. For all other activities, execution time is the difference between start time and stop time. In all cases, execution time does not include time spent initializing or queued.

Table 80. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Activities	event_activity	-

Usage

This element can be used alone to know the elapsed time spent executing the activity by DB2 on each partition. This element can also be used together with **time_started** and **time_completed** monitor elements on the coordinator partition to compute the idle time for cursor activities. You can use the following formula:

Cursor idle time = (time_completed - time_started) - act_exec_time

act_remapped_in – Activities remapped in monitor element

Count of the number of activities to be remapped into this service subclass since the last reset.

Table 81. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_scstats	-

Usage

Use this count to determine whether the remapping of activities into the service subclass is occurring as desired.

act_remapped_out – Activities remapped out monitor element

Count of the number of activities to be remapped out of this service subclass since the last reset.

Table 82. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_scstats	-

Usage

Use this count to determine whether the remapping of activities out of the service subclass is occurring as desired.

act_rows_read_top – Activity rows read top monitor element

The high watermark for the number of rows read by activities at all nesting levels in a service class, workload, or work class.

The monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the service class or workload in which the activity runs is set to NONE. Activities contribute towards this high watermark only when request metrics are enabled.

For service classes, when you remap activities between service subclasses with a REMAP ACTIVITY action only the act_rows_read_top high watermark of the service subclass where an activity completes is updated, provided that a new high watermark is reached. The act_rows_read_top high watermarks of service subclasses an activity is mapped to but does not complete in are unaffected.

Table 83. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_scstats	-
Statistics	event_wcstats	-
Statistics	event_wlstats	-

Usage

Use this element to determine the highest number of rows read by an activity on a partition for a service class, workload, or work class during the time interval collected.

act_total - Activities total monitor element

Total number of activities at any nesting level that had work actions corresponding to the specified work class applied to them since the last reset.

Table 84. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_wcstats	-

Usage

Every time an activity has one or more work actions associated with a work class applied to it, a counter for the work class is updated. This counter is exposed using the **act_total** monitor element. The counter can be used to judge the effectiveness of the work action set (for example, how many activities have a actions applied). It can also be used to understand the different types of activities on the system.

agg_temp_tablespace_top - Aggregate temporary table space top monitor element

The high watermark in KB for the aggregate temporary table space usage of DML activities at all nesting levels in a service class. The aggregate is computed by summing the temporary table space usage across all activities in the service subclass, and this high watermark represents the highest value reached by this aggregate since the last reset. The monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the service class is set to NONE. An AGGSQLTEMPSPACE threshold must be defined and enabled for at least one service subclass in the same superclass as the subclass to which this record belongs, otherwise a value of 0 is returned.

Table 85. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_scstats	-

Usage

Use this element to determine the highest aggregate DML activity system temporary table space usage reached on a partition for a service subclass in the time interval collected.

arm_correlator - Application response measurement correlator monitor element

Identifier of a transaction in the Application Response Measurement (ARM) standard.

Table 86. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Activities	event_activity	-

Usage

This element can be used to link an activity collected by the activities event monitor to the applications associated with the activity, if such applications also support the Application Response Measurement (ARM) standard.

bin_id - Histogram bin identifier monitor element

The identifier of a histogram bin. The **bin_id** is unique within a histogram.

Table 87. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_histogrambin	-

Usage

Use this element to distinguish bins within the same histogram.

bottom - Histogram bin bottom monitor element

The exclusive bottom end of the range of a histogram bin. The value of this monitor element is also the top inclusive end of the range of the previous histogram bin, if there is one.

Table 88. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_histogrambin	-

Usage

Use this element with the corresponding **top** element to determine the range of a bin within a histogram.

concurrent_act_top - Concurrent activity top monitor element

The high watermark for the concurrent activities (at any nesting level) in a service subclass since the last reset.

Table 89. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_scstats	-

Usage

Use this element to know the highest concurrency of activities (including nested activities) reached on a partition for a service subclass in the time interval collected.

concurrent_connection_top - Concurrent connection top monitor element

High watermark for concurrent coordinator connections in this service class since the last reset. This field has the same value in every subclass of the same superclass.

Table 90. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_scstats	-

Usage

This element may be useful in determining where to place thresholds on connection concurrency by showing where the current high watermark is. It is also useful for verifying that such a threshold is configured correctly and doing its job.

concurrent_wlo_act_top - Concurrent WLO activity top monitor element

High watermark for concurrent activities (at any nesting level) of any occurrence of this workload since the last reset.

Table 91. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_wlstats	-

Usage

Use this element to know the highest number of concurrent activities reached on a partition for any occurrence of this workload in the time interval collected.

concurrent_wlo_top - Concurrent workload occurrences top monitor element

The high watermark for the concurrent occurrences of a workload since the last reset.

Table 92. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_wlstats	-
Statistics	event_scstats	-

Usage

Use this element to know the highest concurrency of workload occurrences reached on a partition for a workload in the time interval collected.

coord_act_aborted_total - Coordinator activities aborted total monitor element

The total number of coordinator activities at any nesting level that completed with errors since the last reset. For service classes, the value is updated when the activity completes. For workloads, the value is updated by each workload occurrence at the end of its unit of work.

For service classes, if you remap an activity to a different subclass with a REMAP ACTIVITY action before it aborts, then this activity counts only towards the total of the subclass it aborts in.

Table 93. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_scstats	-
Statistics	event_wlstats	-

Usage

Use this element to understand if activities on the system are completing successfully. Activities may be aborted due to cancellation, errors or reactive thresholds.

coord_act_completed_total - Coordinator activities completed total monitor element

The total number of coordinator activities at any nesting level that completed successfully since the last reset. For service classes, the value is updated when the activity completes. For workloads, the value is updated by each workload occurrence at the end of its unit of work.

For service classes, if you remap an activity to a different subclass with a REMAP ACTIVITY action before it completes, then this activity counts only towards the total of the subclass it completes in.

Table 94. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_wlstats	-
Statistics	event_scstats	-

Usage

This element can be used to determine the throughput of activities in the system or to aid in calculating average activity lifetime across multiple partitions.

coord_act_est_cost_avg - Coordinator activity estimated cost average monitor element

Arithmetic mean of the estimated costs for coordinator DML activities at nesting level 0 associated with this service subclass or work class since the last reset. If the internally tracked average has overflowed, the value -2 is returned. For service subclasses, this monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the service subclass is set to NONE or BASE. For work classes, this monitor elements returns -1 if no COLLECT AGGREGATE ACTIVITY

DATA EXTENDED work action is specified for the work class. For workloads, this monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the workload is set to NONE or BASE. Units are milliseconds.

For service classes, the estimated cost of an activity is counted only towards the service subclass in which the activity enters the system. When you remap activities between service subclasses with a REMAP ACTIVITY action, the coord_act_est_cost_avg mean of the service subclass you remap an activity to is unaffected.

Table 95. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_scstats	-
Statistics	event_wcstats	-
Statistics	event_wlstats	-

Usage

Use this statistic to determine the arithmetic mean of the estimated costs of coordinator DML activities at nesting level 0 that are associated this service subclass, workload, or work class that completed or aborted since the last statistics reset.

This average can also be used to determine whether or not the histogram template used for the activity estimated cost histogram is appropriate. Compute the average activity estimated cost from the activity estimated cost histogram. Compare the computed average with this monitor element. If the computed average deviates from the true average reported by this monitor element, consider altering the histogram template for the activity estimated cost histogram, using a set of bin values that are more appropriate for your data.

coord_act_exec_time_avg - Coordinator activities execution time average monitor element

Arithmetic mean of execution times for coordinator activities at nesting level 0 associated with this service subclass or work class since the last reset. If the internally tracked average has overflowed, the value -2 is returned. For service subclasses, this monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the service subclass is set to NONE. For work classes, this monitor elements returns -1 if no COLLECT AGGREGATE ACTIVITY DATA work action is specified for the work class. For workloads, this monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the workload is set to NONE. Units are milliseconds.

For service classes, when you remap activities between service subclasses with a REMAP ACTIVITY action, the coord_act_exec_time_avg mean of service subclasses an activity is mapped to but does not complete in is unaffected.

Table 96. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_scstats	-
Statistics	event_wcstats	-
Statistics	event_wlstats	-

Usage

Use this statistic to determine the arithmetic mean of execution time for coordinator activities associated with a service subclass, workload, or work class that completed or aborted.

This average can also be used to determine whether or not the histogram template used for the activity execution time histogram is appropriate. Compute the average activity execution time from the activity execution time histogram. Compare the computed average with this monitor element. If the computed average deviates from the true average reported by this monitor element, consider altering the histogram template for the activity execution time histogram, using a set of bin values that are more appropriate for your data.

coord_act_interarrival_time_avg - Coordinator activity arrival time average monitor element

Arithmetic mean of the time between arrivals of coordinator activities at nesting level 0 associated with this service subclass or work class since the last reset. If the internally tracked average has overflowed, the value -2 is returned. For service subclasses, this monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the service subclass is set to NONE or BASE. For work classes, this monitor elements returns -1 if no COLLECT AGGREGATE ACTIVITY DATA EXTENDED work action is specified for the work class. For workloads, this monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the workload is set to NONE or BASE. Units are milliseconds.

For service classes, the inter-arrival time mean is calculated for service subclasses through which activities enter the system. When you remap activities between service subclasses with a REMAP ACTIVITY action, the coord_act_interarrival_time_avg of the service subclass you remap an activity to is unaffected.

Table 97. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_scstats	-
Statistics	event_wcstats	-
Statistics	event_wlstats	-

Usage

Use this statistic to determine the arithmetic mean between arrivals of coordinator activities at nesting level 0 associated with this service subclass, workload, or work class.

The inter-arrival time can be used to determine arrival rate, which is the inverse of inter-arrival time. This average can also be used to determine whether or not the histogram template used for the activity inter-arrival time histogram is appropriate. Compute the average activity inter-arrival time from the activity inter-arrival time histogram. Compare the computed average with this monitor element. If the computed average deviates from the true average reported by this monitor element, consider altering the histogram template for the activity inter-arrival time histogram, using a set of bin values that are more appropriate for your data.

coord_act_lifetime_avg - Coordinator activity lifetime average monitor element

Arithmetic mean of lifetime for coordinator activities at nesting level 0 associated with this service subclass, workload, or work class since the last reset. If the internally tracked average has overflowed, the value -2 is returned. For service subclasses, this monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the service subclass is set to NONE. For work classes, this monitor elements returns -1 if no COLLECT AGGREGATE ACTIVITY DATA work action is specified for the work class. For workloads, this monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the workload is set to NONE. Units are milliseconds.

For service classes, when you remap activities between service subclasses with a REMAP ACTIVITY action, only the the coord_act_lifetime_avg mean of the final service class where the activity completes is affected.

Table 98. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_scstats	-
Statistics	event_wcstats	-
Statistics	event_wlstats	-

Usage

Use this statistic to determine the arithmetic mean of the lifetime for coordinator activities associated with a service subclass, workload, or work class that completed or aborted.

This statistic can also be used to determine whether or not the histogram template used for the activity lifetime histogram is appropriate. Compute the average activity lifetime from the activity lifetime histogram. Compare the computed average with this monitor element. If the computed average deviates from the true average reported by this monitor element, consider altering the histogram template for the activity lifetime histogram, using a set of bin values that are more appropriate for your data.

coord_act_lifetime_top - Coordinator activity lifetime top monitor element

High watermark for coordinator activity lifetime, counted at all nesting levels. Units are milliseconds. For service classes, this monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the service class is set to NONE. For work classes, this monitor element returns -1 if no COLLECT AGGREGATE ACTIVITY DATA work action is specified for the work class. For workloads, this monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the workload is set to NONE.

To effectively use this statistic with service classes when you also remap activities between service subclasses with a REMAP ACTIVITY action, you must aggregate the coord_act_lifetime_top high watermark of any given service subclass with that of other subclasses affected by the same remapping threshold or thresholds. This is because an activity will complete after it has been remapped to a different service

subclass by a remapping threshold, and the time the activity spends in other service subclasses before being remapped is counted only towards the service class in which it completes.

Table 99. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_wcstats	-
Statistics	event_scstats	-
Statistics	event_wlstats	-

Usage

This element can be used to help determine whether or not thresholds on activity lifetime are being effective and can also help to determine how to configure such thresholds.

coord_act_queue_time_avg - Coordinator activity queue time average monitor element

Arithmetic mean of queue time for coordinator activities at nesting level 0 associated with this service subclass or work class since the last reset. If the internally tracked average has overflowed, the value -2 is returned. For service subclasses, this monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the service subclass is set to NONE. For work classes, this monitor elements returns -1 if no COLLECT AGGREGATE ACTIVITY DATA work action is specified for the work class. For workloads, this monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the workload is set to NONE. Units are milliseconds.

For service classes, the queue time counts only towards the service subclass in which the activity completes or is aborted. When you remap activities between service subclasses with a REMAP ACTIVITY action, the coord_act_queue_time_avg mean of service subclasses an activity is mapped to but does not complete in is unaffected.

Element identifier

coord_act_queue_time_avg

Element type

information

-->

Table 100. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_scstats	-
Statistics	event_wcstats	-
Statistics	event_wlstats	-

Usage

Use this statistic to determine the arithmetic mean of the queue time for coordinator activities associated with a service subclass, workload, or work class that completed or aborted.

This statistic can also be used to determine whether or not the histogram template used for the activity queue time histogram is appropriate. Compute the average activity queue time from the activity queue time histogram. Compare the computed average with this monitor element. If the computed average deviates from the true average reported by this monitor element, consider altering the histogram template for the activity queue time histogram, using a set of bin values that are more appropriate for your data.

coord_act_rejected_total - Coordinator activities rejected total monitor element

The total number of coordinator activities at any nesting level that were rejected instead of being allowed to execute since the last reset. This counter is updated when an activity is prevented from executing by either a predictive threshold or a prevent execution work action. For service classes, the value is updated when the activity completes. For workloads, the value is updated by each workload occurrence at the end of its unit of work.

Table 101. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_scstats	-
Statistics	event_wlstats	-

Usage

This element can be used to help determine whether or not predictive thresholds and work actions that prevent execution are being effective and whether or not they are too restrictive.

coord_partition_num - Coordinator partition number monitor element

The coordinator partition of the unit of work or activity. In a multi-partition system, the coordinator partition is the partition where the application connected to the database.

Table 102. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Unit of work	-	-
Activities	event_activity	-
Threshold violations	event_thresholdviolations	-

Usage

This element allows the coordinator partition to be identified for activities or units of work that have records on partitions other than the coordinator.

cost_estimate_top - Cost estimate top monitor element

The high watermark for the estimated cost of DML activities at all nesting levels in a service subclass or work class. For service subclasses, this monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the service subclass is set to NONE. For work classes, this monitor elements returns -1 if no COLLECT AGGREGATE ACTIVITY DATA work action is specified for the work class.

For service classes, the estimated cost of DML activities is counted only towards the service subclass in which the activity enters the system. When you remap activities between service subclasses with a REMAP ACTIVITY action, the cost_estimate_top of the service subclass you remap an activity to is unaffected.

Table 103. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_scstats	-
Statistics	event_wcstats	-
Statistics	event_wlstats	-

Usage

Use this element to determine the highest DML activity estimated cost reached on a partition for a service class, workload, or work class in the time interval collected.

coord_act_lifetime_avg - Coordinator activity lifetime average monitor element

Arithmetic mean of lifetime for coordinator activities at nesting level 0 associated with this service subclass, workload, or work class since the last reset. If the internally tracked average has overflowed, the value -2 is returned. For service subclasses, this monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the service subclass is set to NONE. For work classes, this monitor elements returns -1 if no COLLECT AGGREGATE ACTIVITY DATA work action is specified for the work class. For workloads, this monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the workload is set to NONE. Units are milliseconds.

For service classes, when you remap activities between service subclasses with a REMAP ACTIVITY action, only the the coord_act_lifetime_avg mean of the final service class where the activity completes is affected.

Table 104. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_scstats	-
Statistics	event_wcstats	-
Statistics	event_wlstats	-

Usage

Use this statistic to determine the arithmetic mean of the lifetime for coordinator activities associated with a service subclass, workload, or work class that completed or aborted.

This statistic can also be used to determine whether or not the histogram template used for the activity lifetime histogram is appropriate. Compute the average activity lifetime from the activity lifetime histogram. Compare the computed average with this monitor element. If the computed average deviates from the true average reported by this monitor element, consider altering the histogram template for the activity lifetime histogram, using a set of bin values that are more appropriate for your data.

coord_act_queue_time_avg - Coordinator activity queue time average monitor element

Arithmetic mean of queue time for coordinator activities at nesting level 0 associated with this service subclass or work class since the last reset. If the internally tracked average has overflowed, the value -2 is returned. For service subclasses, this monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the service subclass is set to NONE. For work classes, this monitor elements returns -1 if no COLLECT AGGREGATE ACTIVITY DATA work action is specified for the work class. For workloads, this monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the workload is set to NONE. Units are milliseconds.

For service classes, the queue time counts only towards the service subclass in which the activity completes or is aborted. When you remap activities between service subclasses with a REMAP ACTIVITY action, the coord_act_queue_time_avg mean of service subclasses an activity is mapped to but does not complete in is unaffected.

Element identifier

coord_act_queue_time_avg

Element type

information

-->

Table 105. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_scstats	-
Statistics	event_wcstats	-
Statistics	event_wlstats	-

Usage

Use this statistic to determine the arithmetic mean of the queue time for coordinator activities associated with a service subclass, workload, or work class that completed or aborted.

This statistic can also be used to determine whether or not the histogram template used for the activity queue time histogram is appropriate. Compute the average activity queue time from the activity queue time histogram. Compare the computed average with this monitor element. If the computed average deviates from the true average reported by this monitor element, consider altering the histogram template for the activity queue time histogram, using a set of bin values that are more appropriate for your data.

coord_act_exec_time_avg - Coordinator activities execution time average monitor element

Arithmetic mean of execution times for coordinator activities at nesting level 0 associated with this service subclass or work class since the last reset. If the internally tracked average has overflowed, the value -2 is returned. For service subclasses, this monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the service subclass is set to NONE. For work classes, this

monitor elements returns -1 if no COLLECT AGGREGATE ACTIVITY DATA work action is specified for the work class. For workloads, this monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the workload is set to NONE. Units are milliseconds.

For service classes, when you remap activities between service subclasses with a REMAP ACTIVITY action, the coord_act_exec_time_avg mean of service subclasses an activity is mapped to but does not complete in is unaffected.

Table 106. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_scstats	-
Statistics	event_wcstats	-
Statistics	event_wlstats	-

Usage

Use this statistic to determine the arithmetic mean of execution time for coordinator activities associated with a service subclass, workload, or work class that completed or aborted.

This average can also be used to determine whether or not the histogram template used for the activity execution time histogram is appropriate. Compute the average activity execution time from the activity execution time histogram. Compare the computed average with this monitor element. If the computed average deviates from the true average reported by this monitor element, consider altering the histogram template for the activity execution time histogram, using a set of bin values that are more appropriate for your data.

request_exec_time_avg - Request execution time average monitor element

Arithmetic mean of the execution times for requests associated with this service subclass since the last reset. If the internally tracked average has overflowed, the value -2 is returned. This monitor element returns -1 when COLLECT AGGREGATE REQUEST DATA for the service subclass is set to NONE. Units are milliseconds.

When you remap activities between service subclasses with a REMAP ACTIVITY action, the request_exec_time_avg mean counts the partial request in each subclass involved in remapping.

Table 107. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_scstats	-

Usage

Use this statistic to quickly understand the average amount of time that is spent processing each request on a database partition in this service subclass.

This average can also be used to determine whether or not the histogram template used for the request execution time histogram is appropriate. Compute the average

request execution time from the request execution time histogram. Compare the computed average with this monitor element. If the computed average deviates from the true average reported by this monitor element, consider altering the histogram template for the request execution time histogram, using a set of bin values that are more appropriate for your data.

coord_act_est_cost_avg - Coordinator activity estimated cost average monitor element

Arithmetic mean of the estimated costs for coordinator DML activities at nesting level 0 associated with this service subclass or work class since the last reset. If the internally tracked average has overflowed, the value -2 is returned. For service subclasses, this monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the service subclass is set to NONE or BASE. For work classes, this monitor element returns -1 if no COLLECT AGGREGATE ACTIVITY DATA EXTENDED work action is specified for the work class. For workloads, this monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the workload is set to NONE or BASE. Units are milliseconds.

For service classes, the estimated cost of an activity is counted only towards the service subclass in which the activity enters the system. When you remap activities between service subclasses with a REMAP ACTIVITY action, the coord_act_est_cost_avg mean of the service subclass you remap an activity to is unaffected.

Table 108. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_scstats	-
Statistics	event_wcstats	-
Statistics	event_wlstats	-

Usage

Use this statistic to determine the arithmetic mean of the estimated costs of coordinator DML activities at nesting level 0 that are associated this service subclass, workload, or work class that completed or aborted since the last statistics reset.

This average can also be used to determine whether or not the histogram template used for the activity estimated cost histogram is appropriate. Compute the average activity estimated cost from the activity estimated cost histogram. Compare the computed average with this monitor element. If the computed average deviates from the true average reported by this monitor element, consider altering the histogram template for the activity estimated cost histogram, using a set of bin values that are more appropriate for your data.

coord_act_interarrival_time_avg - Coordinator activity arrival time average monitor element

Arithmetic mean of the time between arrivals of coordinator activities at nesting level 0 associated with this service subclass or work class since the last reset. If the internally tracked average has overflowed, the value -2 is returned. For service subclasses, this monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the service subclass is set to NONE or BASE. For work

classes, this monitor elements returns -1 if no COLLECT AGGREGATE ACTIVITY DATA EXTENDED work action is specified for the work class. For workloads, this monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the workload is set to NONE or BASE. Units are milliseconds.

For service classes, the inter-arrival time mean is calculated for service subclasses through which activities enter the system. When you remap activities between service subclasses with a REMAP ACTIVITY action, the coord_act_interarrival_time_avg of the service subclass you remap an activity to is unaffected.

Table 109. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_scstats	-
Statistics	event_wcstats	-
Statistics	event_wlstats	-

Usage

Use this statistic to determine the arithmetic mean between arrivals of coordinator activities at nesting level 0 associated with this service subclass, workload, or work class.

The inter-arrival time can be used to determine arrival rate, which is the inverse of inter-arrival time. This average can also be used to determine whether or not the histogram template used for the activity inter-arrival time histogram is appropriate. Compute the average activity inter-arrival time from the activity inter-arrival time histogram. Compare the computed average with this monitor element. If the computed average deviates from the true average reported by this monitor element, consider altering the histogram template for the activity inter-arrival time histogram, using a set of bin values that are more appropriate for your data.

db_work_action_set_id - Database work action set ID monitor element

If this activity has been categorized into a work class of database scope, this monitor element shows the ID of the work action set associated with the work class set to which the work class belongs. Otherwise, this monitor element shows the value of 0.

Table 110. Table Function Monitoring Information

Table Function	Monitor Element Collection Command and Level
WLM_GET_ACTIVITY_DETAILS_COMPLETE (reported in DETAILS XML document)	Always collected

Table 111. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Activities	event_activity	-

Usage

This element can be used with the **db_work_class_id** element to uniquely identify the database work class of the activity, if one exists.

db_work_class_id - Database work class ID monitor element

If this activity has been categorized into a work class of database scope, this monitor element displays the ID of the work class. Otherwise, this monitor element displays the value of 0.

Table 112. Table Function Monitoring Information

Table Function	Monitor Element Collection Command and Level
WLM_GET_ACTIVITY_DETAILS_COMPLETE table function - Get complete activity details (reported in DETAILS XML document)	Always collected

Table 113. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Activities	event_activity	-

Usage

This element can be used with the **db_work_action_set_id** element to uniquely identify the database work class of the activity, if one exists.

destination_service_class_id – Destination service class ID monitor element

The ID of the service subclass to which an activity was remapped when the threshold violation record to which this element belongs was generated. This element has a value of zero for any threshold action other than REMAP ACTIVITY.

Element identifier

destination_service_class_id

Element type

Information

Table 114. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Threshold violations	event_thresholdviolations	-

Usage

Use this element to trace the path of an activity through the service classes to which it was remapped. This element can also be used to compute aggregates of how many activities were mapped into a given service subclass.

histogram_type - Histogram type monitor element

The type of the histogram, in string format.

There are six histogram types.

CoordActQueueTime

A histogram of the time non-nested activities spend queued (for example, in a threshold queue), measured on the coordinator partition.

CoordActExecTime

A histogram of the time non-nested activities spend executing at the coordinator partition. Execution time does not include time spent initializing or queued. For cursors, execution time includes only the time spent on open, fetch and close requests. When an activity is remapped between service subclasses, the execution time histogram is updated only for the service subclass in which the activity completes execution.

CoordActLifetime

A histogram of the elapsed time from when a non-nested activity is identified by the database manager until the activity completes execution, as measured on the coordinator partition. When you remap activities between service subclasses, the lifetime histogram is updated only for the service subclass in which the activity completes execution.

CoordActInterArrivalTime

A histogram of the time interval between the arrival of non-nested coordinator activities. The inter-arrival time mean is calculated for service subclasses through which activities enter the system. When you remap activities between service subclasses, the inter-arrival time histogram of the service subclass you remap an activity to is unaffected.

CoordActEstCost

A histogram of the estimated cost of non-nested DML activities. The estimated cost of an activity is counted only towards the service subclass in which the activity enters the system.

ReqExecTime

A histogram of request execution times, which includes requests on the coordinator partition, and any subrequests on both coordinator and non-coordinator partitions (like RPC requests or SMP subagent requests). Requests included may or may not be associated with an activity: Both PREPARE and OPEN requests are included in this histogram, for example, but while OPEN requests are always associated with a cursor activity, PREPARE requests are not part of any activity. The execution time histogram of a service subclass involved in remapping counts the portion of the execution time spent by the partial request in the service subclass.

Table 115. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_histogrambin	-

Usage

Use this element to identify the type of histogram. Several histograms can belong to the same statistics record, but only one of each type.

last_wlm_reset - Time of last reset monitor element

This element, in the form of a local timestamp, shows the time at which the last statistics event record of this type was created.

Table 116. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_scstats	-
Statistics	event_wlstats	-
Statistics	event_wcstats	-
Statistics	event_qstats	-

Usage

Use the **wlm_last_reset** and **statistics_timestamp** monitor elements to determine a period of time over which the statistics in an event monitor statistics record were collected. The collection interval begins at the **wlm_last_reset** time and ends at **statistics_timestamp**.

num_threshold_violations - Number of threshold violations monitor element

The number of threshold violations that have taken place in this database since it was last activated.

Table 117. Snapshot Monitoring Information

Snapshot Level	Logical Data Grouping	Monitor Switch
Database	dbase	Basic

For snapshot monitoring, this counter can be reset.

Table 118. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Database	event_db	-

Usage

This element can be used to help determine whether or not thresholds are effective for this particular application or whether the threshold violations are excessive.

num_remaps - Number of remaps monitor element

Count of the number of times this activity has been remapped. If **num_remaps** is greater than zero, the **service_class_id** of this activity record is the ID of the last service class to which the activity was remapped.

Table 119. Table Function Monitoring Information

Table Function	Monitor Element Collection Command and Level
WLM_GET_ACTIVITY_DETAILS table function - Get complete activity details (reported in DETAILS XML document)	Always collected

Table 120. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Activities	event_activity	-

Usage

Use this information to verify whether the activity was remapped the expected number of times.

number_in_bin - Number in bin monitor element

This element holds the count of the number of activities or requests that fall within the histogram bin.

Table 121. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_histogrambin	-

Usage

Use this element to represent the height of a bin in the histogram.

parent_activity_id - Parent activity ID monitor element

The unique ID of the activity's parent activity within the parent activity's unit of work. If there is no parent activity, the value of this monitor element is 0.

Table 122. Table Function Monitoring Information

Table Function	Monitor Element Collection Command and Level
MON_GET_ACTIVITY_DETAILS table function - Get complete activity details (reported in DETAILS XML document)	Always collected

Table 123. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Activities	event_activity	-

Usage

Use this element along with the **parent_uow_id** element and **appl_id** element to uniquely identify the parent activity of the activity described in this activity record.

parent_uow_id - Parent unit of work ID monitor element

The unique unit of work identifier within an application handle. The ID of the unit of work in which the activity's parent activity originates. If there is no parent activity, the value is 0.

Table 124. Table Function Monitoring Information

Table Function	Monitor Element Collection Command and Level
MON_GET_ACTIVITY_DETAILS table function - Get complete activity details (reported in DETAILS XML document)	Always collected

Table 125. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Activities	event_activity	-

Usage

Use this element along with the `parent_activity_id` element and `appl_id` element to uniquely identify the parent activity of the activity described in this activity record.

prep_time - Preparation time monitor element

Time in milliseconds required to prepare an SQL statement if the activity is an SQL statement.

Table 126. Table Function Monitoring Information

Table Function	Monitor Element Collection Level
MON_GET_PKG_CACHE_STMT table function - Get SQL statement activity metrics in the package cache	ACTIVITY METRICS BASE

Table 127. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Activities	event_activity	-

Usage

This element can be used to identify how much of the activity's total lifetime was spent preparing the SQL statement, if this was an SQL activity.

queue_assignments_total - Queue assignments total monitor element

The number of times any connection or activity was assigned to this threshold queue since the last reset.

Table 128. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_qstats	-

Usage

This element can be used to determine the number of times any connection or activity was queued in this particular queue in a given period of time determined by the statistics collection interval. This can help to determine the effectiveness of queuing thresholds.

queue_size_top - Queue size top monitor element

Highest queue size that has been reached since the last reset.

Table 129. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_qstats	-

Usage

Use this element to gauge the effectiveness of queuing thresholds and to detect when queuing is excessive.

queue_time_total - Queue time total monitor element

Sum of the times spent in the queue for all connections or activities placed in this queue since the last reset. Units are milliseconds.

Table 130. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_qstats	-

Usage

Use this element to gauge the effectiveness of queuing thresholds and to detect when queuing is excessive.

request_exec_time_avg - Request execution time average monitor element

Arithmetic mean of the execution times for requests associated with this service subclass since the last reset. If the internally tracked average has overflowed, the value -2 is returned. This monitor element returns -1 when COLLECT AGGREGATE REQUEST DATA for the service subclass is set to NONE. Units are milliseconds.

When you remap activities between service subclasses with a REMAP ACTIVITY action, the request_exec_time_avg mean counts the partial request in each subclass involved in remapping.

Table 131. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_scstats	-

Usage

Use this statistic to quickly understand the average amount of time that is spent processing each request on a database partition in this service subclass.

This average can also be used to determine whether or not the histogram template used for the request execution time histogram is appropriate. Compute the average request execution time from the request execution time histogram. Compare the computed average with this monitor element. If the computed average deviates from the true average reported by this monitor element, consider altering the histogram template for the request execution time histogram, using a set of bin values that are more appropriate for your data.

rows_fetched - Rows fetched monitor element

The number of rows read from the table.

This monitor element is an alias of the **rows_read** monitor element.

Note: This monitor element reports only the values for the database partition for which this information is recorded. On DPF systems, these values may not reflect the correct totals for the whole activity.

Element identifier

rows_fetched

Element type

counter

-->

Table 132. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Activities	event_activity	Statement

Usage

See the **rows_read** monitor element for details.

rows_modified - Rows modified monitor element

The number of rows inserted, updated, or deleted.

This monitor element is an alias of the **rows_written** monitor element.

Table 133. Table Function Monitoring Information

Table Function	Monitor Element Collection Level
MON_GET_CONNECTION table function - Get connection metrics	REQUEST METRICS BASE
MON_GET_CONNECTION_DETAILS table function - Get detailed connection metrics (reported in DETAILS XML document)	REQUEST METRICS BASE
MON_GET_SERVICE_SUBCLASS table function - Get service subclass metrics	REQUEST METRICS BASE
MON_GET_SERVICE_SUBCLASS_DETAILS table function - Get detailed service subclass metrics (reported in DETAILS XML document)	REQUEST METRICS BASE
MON_GET_UNIT_OF_WORK table function - Get unit of work metrics	REQUEST METRICS BASE
MON_GET_UNIT_OF_WORK_DETAILS table function - Get detailed unit of work metrics (reported in DETAILS XML document)	REQUEST METRICS BASE
MON_GET_WORKLOAD table function - Get workload metrics	REQUEST METRICS BASE

Table 133. Table Function Monitoring Information (continued)

Table Function	Monitor Element Collection Level
MON_GET_WORKLOAD_DETAILS table function - Get detailed workload metrics (reported in DETAILS XML document)	REQUEST METRICS BASE
MON_GET_ACTIVITY_DETAILS table function - Get complete activity details (reported in DETAILS XML document)	ACTIVITY METRICS BASE
MON_GET_PKG_CACHE_STMT table function - Get SQL statement activity metrics in the package cache	ACTIVITY METRICS BASE

Table 134. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Activities	event_activity (reported in the details_xml document)	ACTIVITY METRICS BASE
Statistics	event_scstats (reported in the details_xml document)	REQUEST METRICS BASE
Statistics	event_wlstats (reported in the details_xml document)	REQUEST METRICS BASE
Unit of work	Reported in the in the system_metrics document.	-
Activities	event_activity	Statement

Usage

See the **rows_written** monitor element for details.

rows_returned - Rows returned monitor element

The number of rows that have been selected and returned to the application. This element has a value of 0 for partial activity records (for example, if an activity is collected while it is still executing or when a full activity record could not be written to the event monitor due to memory limitations).

This monitor element is an alias of the **fetch_count** monitor element.

Table 135. Table Function Monitoring Information

Table Function	Monitor Element Collection Level
MON_GET_CONNECTION table function - Get connection metrics	REQUEST METRICS BASE
MON_GET_CONNECTION_DETAILS table function - Get detailed connection metrics (reported in DETAILS XML document)	REQUEST METRICS BASE
MON_GET_SERVICE_SUBCLASS table function - Get service subclass metrics	REQUEST METRICS BASE
MON_GET_SERVICE_SUBCLASS_DETAILS table function - Get detailed service subclass metrics (reported in DETAILS XML document)	REQUEST METRICS BASE

Table 135. Table Function Monitoring Information (continued)

Table Function	Monitor Element Collection Level
MON_GET_UNIT_OF_WORK table function - Get unit of work metrics	REQUEST METRICS BASE
MON_GET_UNIT_OF_WORK_DETAILS table function - Get detailed unit of work metrics (reported in DETAILS XML document)	REQUEST METRICS BASE
MON_GET_WORKLOAD table function - Get workload metrics	REQUEST METRICS BASE
MON_GET_WORKLOAD_DETAILS table function - Get detailed workload metrics (reported in DETAILS XML document)	REQUEST METRICS BASE
MON_GET_ACTIVITY_DETAILS table function - Get complete activity details (reported in DETAILS XML document)	ACTIVITY METRICS BASE
MON_GET_PKG_CACHE_STMT table function - Get SQL statement activity metrics in the package cache	ACTIVITY METRICS BASE

Table 136. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Activities	event_activity (reported in the details_xml document)	ACTIVITY METRICS BASE
Statistics	event_scstats (reported in the details_xml document)	REQUEST METRICS BASE
Statistics	event_wlstats (reported in the details_xml document)	REQUEST METRICS BASE
Unit of work	Reported in the in the system_metrics document.	-
Activities	event_activity	-

Usage

This element can be used to help determine thresholds for rows returned to the application or can be used to verify that such a threshold is configured correctly and doing its job.

rows_returned_top - Actual rows returned top monitor element

The high watermark for the actual rows returned of DML activities at all nesting levels in a service class or work class. For service classes, this monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the service class is set to NONE. For work classes, this monitor element returns -1 if no COLLECT AGGREGATE ACTIVITY DATA work action is specified for the work class. For workloads, this monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the workload is set to NONE.

For service classes, when you remap activities between service subclasses with a REMAP ACTIVITY action, only the rows_returned_top high watermark of the

service subclass where an activity completes is updated. High watermarks of service subclasses an activity is mapped to but does not complete in are unaffected.

Table 137. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_scstats	-
Statistics	event_wcstats	-
Statistics	event_wlstats	-

Usage

Use this element to know the highest DML activity actual rows returned reached on a partition for a service class, workload, or work class in the time interval collected.

sc_work_action_set_id - Service class work action set ID monitor element

If this activity has been categorized into a work class of service class scope, this monitor element displays the ID of the work action set associated with the work class set to which the work class belongs. Otherwise, this monitor element displays the value of 0.

Table 138. Table Function Monitoring Information

Table Function	Monitor Element Collection Command and Level
WLM_GET_ACTIVITY_DETAILS_COMPLETE table function - Get complete activity details (reported in DETAILS XML document)	Always collected

Table 139. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Activities	event_activity	-

Usage

This element can be used with the `sc_work_class_id` element to uniquely identify the service class work class of the activity, if one exists.

sc_work_class_id - Service class work class ID monitor element

If this activity has been categorized into a work class of service class scope, this monitor element displays the ID of the work class assigned to this activity. Otherwise, this monitor element displays the value of 0.

Table 140. Table Function Monitoring Information

Table Function	Monitor Element Collection Command and Level
WLM_GET_ACTIVITY_DETAILS_COMPLETE table function - Get complete activity details (reported in DETAILS XML document)	Always collected

Table 141. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Activities	event_activity	-

Usage

This element can be used with the `sc_work_action_set_id` element to uniquely identify the service class work class of the activity, if one exists.

section_env - Section environment monitor element

A handle that gives details of an activity's section.

Table 142. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Activities	event_activitystmt	-

Usage

This element is to be used with future IBM tools for extracting section information for the activity described in this record

service_class_id - Service class ID monitor element

Unique ID of service subclass. For a workload, this ID represents the service subclass that the workload is mapped to. For a unit of work, this ID represents the service subclass ID of the workload that the connection issuing the unit of work is associated with.

Table 143. Table Function Monitoring Information

Table Function	Monitor Element Collection Level
MON_GET_SERVICE_SUBCLASS table function - Get service subclass metrics	Always collected
MON_GET_SERVICE_SUBCLASS_DETAILS table function - Get detailed service subclass metrics	Always collected
MON_GET_UNIT_OF_WORK table function - Get unit of work metrics	Always collected
MON_GET_UNIT_OF_WORK_DETAILS table function - Get detailed unit of work metrics	Always collected
MON_GET_WORKLOAD table function - Get workload metrics	Always collected
MON_GET_ACTIVITY_DETAILS table function - Get complete activity details (reported in DETAILS XML document)	Always collected

Table 144. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Activities	event_activity (reported in the details_xml document)	ACTIVITY METRICS BASE
Statistics	event_scstats (reported in the details_xml document)	REQUEST METRICS BASE

Table 144. Event Monitoring Information (continued)

Event Type	Logical Data Grouping	Monitor Switch
Locking	-	-
Unit of work	-	-
Statistics	event_histogrambin	-
Statistics	event_scstats	-

Usage

The value of this element matches a value from column SERVICECLASSID of view SYSCAT.SERVICECLASSES. Use this element to look up the service subclass name, or link information about a service subclass from different sources. For example, join service class statistics with histogram bin records.

service_subclass_name - Service subclass name monitor element

The name of a service subclass.

Table 145. Table Function Monitoring Information

Table Function	Monitor Element Collection Level
MON_GET_SERVICE_SUBCLASS_DETAILS table function - Get detailed service subclass metrics (reported in DETAILS XML document)	Always collected
MON_GET_SERVICE_SUBCLASS table function - Get service subclass metrics	Always collected
MON_GET_UNIT_OF_WORK table function - Get unit of work metrics	Always collected
MON_GET_UNIT_OF_WORK_DETAILS table function - Get detailed unit of work metrics (reported in DETAILS XML document)	Always collected

Table 146. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_scstats (reported in the details_xml document)	REQUEST METRICS BASE
Locking	-	-
Unit of work	-	-
Activities	event_activity	-
Statistics	event_scstats	-
Statistics	event_qstats	-

Usage

Use this element in conjunction with other activity elements for analysis of the behavior of an activity or with other statistics elements for analysis of a service class or threshold queue.

service_superclass_name - Service superclass name monitor element

The name of a service superclass.

Table 147. Table Function Monitoring Information

Table Function	Monitor Element Collection Level
MON_GET_SERVICE_SUBCLASS table function - Get service subclass metrics	Always collected
MON_GET_SERVICE_SUBCLASS_DETAILS table function - Get detailed service subclass metrics (reported in DETAILS XML document)	Always collected
MON_GET_UNIT_OF_WORK table function - Get unit of work metrics	Always collected
MON_GET_UNIT_OF_WORK_DETAILS table function - Get detailed unit of work metrics (reported in DETAILS XML document)	Always collected

Table 148. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_scstats (reported in the details_xml document)	REQUEST METRICS BASE
Unit of work	-	-
Activities	event_activity	-
Statistics	event_scstats	-
Statistics	event_qstats	-

Usage

Use this element in conjunction with other activity elements for analysis of the behavior of an activity or with other statistics elements for analysis of a service class or threshold queue.

source_service_class_id - Source service class ID monitor element

The ID of the service subclass from which an activity was remapped when the threshold violation record to which this element belongs was generated. This element has a value of zero when the threshold action is anything other than a REMAP ACTIVITY action.

Element identifier

source_service_class_id

Element type

Information

Table 149. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Threshold violations	event_thresholdviolations	-

Usage

Use this element to trace the path of an activity through the service classes to which it was remapped. It can also be used to compute aggregates of how many activities were mapped out of a given service subclass.

statistics_timestamp - Statistics timestamp monitor element

The time at which this statistics record was generated.

Table 150. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_scstats	-
Statistics	event_wlstats	-
Statistics	event_wcstats	-
Statistics	event_qstats	-
Statistics	event_histogrambin	-

Usage

Use this element to determine when this statistics record was generated.

Use this element along with the **last_wlm_reset** element to identify the time interval over which the statistics in this statistics record were generated.

This monitor element can also be used to group together all statistics records that were generated for the same collection interval.

temp_tablespace_top - Temporary table space top monitor element

The high watermark in KB for the temporary table space usage of DML activities at all nesting levels in a service class or work class. For service classes, this monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the service class is set to NONE. For work classes, this monitor elements returns -1 if no COLLECT AGGREGATE ACTIVITY DATA work action is specified for the work class. For workloads, this monitor element returns -1 when COLLECT AGGREGATE ACTIVITY DATA for the workload is set to NONE.

For service classes, when you remap activities between service subclasses with a REMAP ACTIVITY action, only the temp_tablespace_top high watermark of the service subclass where an activity completes is changed. High watermarks of service subclasses an activity is mapped to but does not complete in are unaffected.

Table 151. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_scstats	-
Statistics	event_wcstats	-
Statistics	event_wlstats	-

Usage

Use this element to determine the highest DML activity system temporary table space usage reached on a partition for a service class, workload, or work class in the time interval collected.

This element is only updated by activities that have a temporary table space threshold applied to them. If no temporary table space threshold is applied to an activity, a value of 0 is returned.

threshold_action - Threshold action monitor element

The action of the threshold to which this threshold violation record applies. Possible values include Stop, Continue and Remap.

Table 152. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Threshold violations	event_thresholdviolations	-

Usage

Use this element to determine whether the activity that violated the threshold was stopped when the violation occurred, was allowed to continue executing, or was remapped to another service subclass. If the activity was stopped, the application that submitted the activity will have received an SQL4712N error. If the activity was remapped to another service subclass, agents working for the activity on the partition will be moving to the target service subclass of the threshold.

threshold_domain - Threshold domain monitor element

The domain of the threshold responsible for this queue.

Possible values are

- Database
- Work Action Set
- Service Superclass
- Service Subclass
- Workload

Table 153. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_qstats	-

Usage

This element can be used for distinguishing the queue statistics of thresholds that have the same predicate but different domains.

threshold_maxvalue - Threshold maximum value monitor element

For non-queuing thresholds, this monitor element represents the value that was exceeded to cause this threshold violation. For queuing thresholds, this monitor element represents the level of concurrency that caused the queuing. The level of

concurrency that caused the violation of the queuing threshold is the sum of **threshold_maxvalue** and **threshold_queuesize** monitor elements.

Table 154. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Threshold violations	event_thresholdviolations	-

Usage

For activity thresholds, this element provides a historical record of what the threshold's maximum value was at the time the threshold was violated. This is useful when the threshold's maximum value has changed since the time of the violation and the old value is no longer available from the SYSCAT.THRESHOLDS view.

threshold_name - Threshold name monitor element

The unique name of the threshold responsible for this queue.

Table 155. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_qstats	-

Usage

Use this element to uniquely identify the queuing threshold whose statistics this record represents.

threshold_predicate - Threshold predicate monitor element

Identifies the type of threshold that was violated or for which statistics were collected.

Table 156. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Threshold violations	event_thresholdviolations	-
Statistics	event_qstats	-

Usage

Use this monitor element in conjunction with other statistics or threshold violation monitor elements for analysis of a threshold violation.

threshold_queuesize - Threshold queue size monitor element

The size of the queue for a queuing threshold. An attempt to exceed this size causes a threshold violation. For a non-queuing threshold, this value is 0.

Table 157. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Threshold violations	event_thresholdviolations	-

Usage

Use this element to determine the number of activities or connections in the queue for this threshold at the time the threshold was violated.

thresholdid - Threshold ID monitor element

Identifies the threshold to which a threshold violation record applies or for which queue statistics were collected.

Table 158. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Threshold violations	event_thresholdviolations	-
Statistics	event_qstats	-

Usage

Use this monitor element in conjunction with other activity history monitor elements for analysis of a threshold queue or for analysis of the activity that violated a threshold.

time_completed - Time completed monitor element

The time at which the activity described by this activity record finished executing. This element is a local timestamp.

This field has a value of "0000-00-00-00.00.00.000000" when a full activity record could not be written to a table event monitor due to memory limitations or if the activity was captured while it was in progress.

Element identifier

time_completed

Element type

information

-->

Table 159. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Activities	event_activity	-

Usage

Use this element in conjunction with other activity history elements for analysis of the behavior of an activity.

time_created - Time created monitor element

The time at which a user submitted the activity described by this activity record. This element is a local timestamp.

Table 160. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Activities	event_activity	-

Usage

Use this element in conjunction with other activity history elements for analysis of the behavior of an activity.

time_of_violation - Time of violation monitor element

The time at which the threshold violation described in this threshold violation record occurred. This element is a local timestamp.

Table 161. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Threshold violations	event_thresholdviolations	-

Usage

Use this element in conjunction with other threshold violations monitor elements for analysis of a threshold violation.

time_started - Time started monitor element

The time at which the activity described by this activity record began executing. This element is a local timestamp.

Table 162. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Activities	event_activity	-

Usage

Use this element in conjunction with other activity history elements for analysis of the behavior of an activity.

top - Histogram bin top monitor element

The inclusive top end of the range of a histogram bin. The value of this monitor element is also the bottom exclusive end of the range of the next histogram bin.

Table 163. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_histogrambin	-

Usage

Use this element with the corresponding **bottom** element to determine the range of a bin within a histogram.

uow_id - Unit of work ID monitor element

The unit of work identifier. The unit of work ID is unique within an application handle.

Table 164. Table Function Monitoring Information

Table Function	Monitor Element Collection Level
MON_GET_UNIT_OF_WORK table function - Get unit of work metrics	Always collected
MON_GET_UNIT_OF_WORK_DETAILS table function - Get detailed unit of work metrics	Always collected
MON_GET_ACTIVITY_DETAILS table function - Get complete activity details	Always collected

Table 165. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Locking	-	-
Unit of work	-	-
Activities	event_activity	-
Activities	event_activitystmt	-
Activities	event_activityvals	-
Threshold violations	event_thresholdviolations	-

Usage

Use this element in conjunction with other activity history elements for analysis of the behavior of an activity.

You can also use this element with the **activity_id** and **appl_id** monitor elements to uniquely identify an activity.

uow_total_time_top - UOW total time top monitor element

Reserved for future use.

Table 166. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_wlstats	-
Statistics	event_scstats	-

Usage

Reserved for future use.

wlo_completed_total - Workload occurrences completed total monitor element

The number of workload occurrences to complete since last reset.

Table 167. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_wlstats	-

Usage

Use this element to determine how many occurrences of a given workload are driving work into the system.

work_action_set_id - Work action set ID monitor element

The ID of the work action set to which this statistics record applies.

Table 168. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_histogrambin	-
Statistics	event_wcstats	-

Usage

Use this element in conjunction with other activity history elements for analysis of the behavior of an activity or with other statistics elements for analysis of a work class.

work_action_set_name - Work action set name monitor element

The name of the work action set to which the statistics shown as part of this event are associated.

Table 169. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_qstats	-
Statistics	event_wcstats	-

Usage

Use this element along with the **work_class_name** element to uniquely identify the work class whose statistics are being shown in this record or to uniquely identify the work class which is the domain of the threshold queue whose statistics are shown in this record.

work_class_id - Work class ID monitor element

The identifier of the work class to which this statistics record applies.

Table 170. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_wcstats	-
Statistics	event_histogrambin	-

Usage

Use this element in conjunction with other statistics elements for analysis of a work class.

work_class_name - Work class name monitor element

The name of the work class to which the statistics shown as part of this event are associated.

Table 171. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Statistics	event_qstats	-
Statistics	event_wcstats	-

Usage

Use this element along with the **work_action_set_name** element to uniquely identify the work class whose statistics are being shown in this record or to uniquely identify the work class which is the domain of the threshold queue whose statistics are shown in this record.

workload_id - Workload ID monitor element

An integer that uniquely identifies a workload.

Table 172. Table Function Monitoring Information

Table Function	Monitor Element Collection Level
MON_GET_WORKLOAD table function - Get workload metrics	Always collected
MON_GET_WORKLOAD_DETAILS table function - Get detailed workload metrics	Always collected

Table 173. Snapshot Monitoring Information

Snapshot Level	Logical Data Grouping	Monitor Switch
Application	appl_info	Basic

Table 174. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Locking	-	-
Unit of work	-	-
Activities	event_activity (reported in the details_xml document)	ACTIVITY METRICS BASE
Statistics	event_scstats (reported in the details_xml document)	REQUEST METRICS BASE
Statistics	event_wlstats (reported in the details_xml document)	REQUEST METRICS BASE
Unit of work	Reported in the in the system_metrics document.	-
Statistics	event_wlstats	-
Statistics	event_histogrambin	-
Activities	event_activity	-

Usage

Use this ID to uniquely identify the workload to which this activity, application, histogram bin, or workload statistics record belongs.

workload_name - Workload name monitor element

Name of the workload.

Table 175. Table Function Monitoring Information

Table Function	Monitor Element Collection Level
MON_GET_UNIT_OF_WORK table function - Get unit of work metrics	Always collected
MON_GET_UNIT_OF_WORK_DETAILS table function - Get detailed unit of work metrics	Always collected
MON_GET_WORKLOAD table function - Get workload metrics	Always collected
MON_GET_WORKLOAD_DETAILS table function - Get detailed workload metrics	Always collected

Table 176. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Locking	-	-
Unit of work	-	-
Activities	event_activity (reported in the details_xml document)	ACTIVITY METRICS BASE
Statistics	event_scstats (reported in the details_xml document)	REQUEST METRICS BASE
Statistics	event_wlstats (reported in the details_xml document)	REQUEST METRICS BASE
Unit of work	Reported in the in the system_metrics document.	-
Statistics	event_wlstats	-

Usage

In the statistics event monitor and workload table functions, the workload name identifies the workload for which statistics or metrics are being collected and reported. In the unit of work event monitor and unit of work table functions, the workload name identifies the workload that the unit of work was associated with.

Use the workload name to identify units of work or sets of information that apply to a particular workload of interest.

workload_occurrence_id - Workload occurrence identifier monitor element

The ID of the workload occurrence to which this activity belongs.

Table 177. Table Function Monitoring Information

Table Function	Monitor Element Collection Level
MON_GET_UNIT_OF_WORK table function - Get unit of work metrics	Always collected
MON_GET_UNIT_OF_WORK_DETAILS table function - Get detailed unit of work metrics	Always collected

Table 178. Event Monitoring Information

Event Type	Logical Data Grouping	Monitor Switch
Unit of work	-	-
Activities	event_activity	-

Usage

Use this to identify the workload occurrence that submitted the activity.

Commands

SET WORKLOAD command

Specifies the workload to which the database connection is to be assigned. This command can be issued prior to connecting to a database or it can be used to reassign the current connection once the connection has been established. If the connection has been established, the workload reassignment will be performed at the beginning of the next unit of work.

Authorization

None, but see usage notes

Required connection

None

Command syntax

```

➤➤ SET WORKLOAD TO {AUTOMATIC | SYSDEFAULTADMWORKLOAD} ➤➤

```

Command parameters

AUTOMATIC

Specifies that the database connection will be assigned to a workload chosen by the workload evaluation that is performed automatically by the server.

SYSDEFAULTADMWORKLOAD

Specifies that the database connection will be assigned to the SYSDEFAULTADMWORKLOAD, allowing users with *accessctrl*, *dataaccess*, *wlmadm*, *secadm* or *dbadm* authority to bypass the normal workload evaluation.

Examples

To assign the connection to the SYSDEFAULTADMWORKLOAD:

```
SET WORKLOAD TO SYSDEFAULTADMWORKLOAD
```

To reset the workload assignment so that it uses the workload that is chosen by the workload evaluation performed by the server:

```
SET WORKLOAD TO AUTOMATIC
```

Usage notes

If the session authorization ID of the database connection does not have *accessctrl*, *dataaccess*, *wlmadm*, *secadm* or *dbadm* authority, the connection cannot be assigned to the SYSDEFAULTADMWORKLOAD and an SQL0552N error will be returned. If the SET WORKLOAD TO SYSDEFAULTADMWORKLOAD command is issued prior to connecting to a database, the SQL0552N error will be returned after the database connection has been established, at the beginning of the first unit of work. If the command is issued when the database connection has been established, the SQL0552N error will be returned at the beginning of the next unit of work, when the workload reassignment is supposed to take place.

Configuration parameters

wlm_collect_int - Workload management collection interval configuration parameter

This parameter specifies a collect and reset interval, in minutes, for workload management (WLM) statistics.

Every x *wlm_collect_int* minutes, (where x is the value of the *wlm_collect_int* parameter) all workload management statistics are collected and sent to any active statistics event monitor; then the statistics are reset. If an active event monitor exists, depending on how it was created, the statistics are written either to file or to a table. If it does not exist, the statistics are only reset and not collected.

The collect and reset process is initiated from the catalog partition. The *wlm_collect_int* parameter must be specified on the catalog partition. It is not used on other partitions.

Configuration type

Database

Parameter type

Configurable online

Default [range]

0 [0 (no collection performed), 5 - 32 767]

The workload management statistics collected by a statistics event monitor can be used to monitor both short term and long term system behavior. A small interval can be used to obtain both short term and long term system behavior because the results can be merged together to obtain long term behavior. However, having to manually merge the results from different intervals complicates the analysis. If it's not required, a small interval unnecessarily increases the overhead. Therefore, reduce the interval to capture shorter term behavior, and increase the interval to reduce overhead when only analysis of long term behavior is sufficient.

The interval needs to be customized per database, not for each SQL request, or command invocation, or application. There are no other configuration parameters that need to be considered.

Note: All WLM statistics table functions return statistics that have been accumulated since the last time the statistics were reset. The statistics will be reset regularly on the interval specified by this configuration parameter.

Catalog views

SYSCAT.HISTOGRAMTEMPLATEBINS

Each row represents a histogram template bin.

Table 179. SYSCAT.HISTOGRAMTEMPLATEBINS Catalog View

Column Name	Data Type	Nullable	Description
TEMPLATENAME	VARCHAR (128)	Y	Name of the histogram template.
TEMPLATEID	INTEGER		Identifier for the histogram template.
BINID	INTEGER		Identifier for the histogram template bin.
BINUPPERVALUE	BIGINT		The upper value for a single bin in the histogram template.

SYSCAT.HISTOGRAMTEMPLATES

Each row represents a histogram template.

Table 180. SYSCAT.HISTOGRAMTEMPLATES Catalog View

Column Name	Data Type	Nullable	Description
TEMPLATEID	INTEGER		Identifier for the histogram template.
TEMPLATENAME	VARCHAR (128)		Name of the histogram template.
CREATE_TIME	TIMESTAMP		Time at which the histogram template was created.
ALTER_TIME	TIMESTAMP		Time at which the histogram template was last altered.
NUMBINS	INTEGER		Number of bins in the histogram template, including the last bin that has an unbounded top value.
REMARKS	VARCHAR (254)	Y	User-provided comments, or the null value.

SYSCAT.HISTOGRAMTEMPLATEUSE

Each row represents a relationship between a workload management object that can use histogram templates and a histogram template.

Table 181. SYSCAT.HISTOGRAMTEMPLATEUSE Catalog View

Column Name	Data Type	Nullable	Description
TEMPLATENAME	VARCHAR (128)	Y	Name of the histogram template.
TEMPLATEID	INTEGER		Identifier for the histogram template.

Table 181. SYSCAT.HISTOGRAMTEMPLATEUSE Catalog View (continued)

Column Name	Data Type	Nullable	Description
HISTOGRAMTYPE	CHAR (1)		The type of information collected by histograms based on this template. <ul style="list-style-type: none"> • C = Activity estimated cost histogram • E = Activity execution time histogram • I = Activity interarrival time histogram • L = Activity life time histogram • Q = Activity queue time histogram • R = Request execution time histogram
OBJECTTYPE	CHAR (1)		The type of WLM object. <ul style="list-style-type: none"> • b = Service class • k = Work action • w = Workload
OBJECTID	INTEGER		Identifier of the WLM object.
SERVICECLASSNAME	VARCHAR (128)	Y	Name of the service class.
PARENTSERVICECLASSNAME	VARCHAR (128)	Y	The name of the parent service class of the service subclass that uses the histogram template.
WORKACTIONNAME	VARCHAR (128)	Y	The name of the work action that uses the histogram template.
WORKACTIONSETNAME	VARCHAR (128)	Y	The name of the work action set containing the work action that uses the histogram template.
WORKLOADNAME	VARCHAR (128)	Y	The name of the workload that uses the histogram template.

SYSCAT.SERVICECLASSES

Each row represents a service class.

Table 182. SYSCAT.SERVICECLASSES Catalog View

Column Name	Data Type	Nullable	Description
SERVICECLASSNAME	VARCHAR (128)		Name of the service class.
PARENTSERVICECLASSNAME	VARCHAR (128)	Y	Service class name of the parent service superclass.
SERVICECLASSID	SMALLINT		Identifier for the service class.
PARENTID	SMALLINT		Identifier for the parent service class for this service class. 0 if this service class is a super service class.
CREATE_TIME	TIMESTAMP		Time when the service class was created.
ALTER_TIME	TIMESTAMP		Time when the service class was last altered.
ENABLED	CHAR (1)		State of the service class. <ul style="list-style-type: none"> • N = Disabled • Y = Enabled

Table 182. SYSCAT.SERVICECLASSES Catalog View (continued)

Column Name	Data Type	Nullable	Description
AGENTPRIORITY	SMALLINT		Thread priority of the agents in the service class relative to the normal priority of DB2 threads. <ul style="list-style-type: none"> • -20 to 20 (Linux and UNIX) • -6 to 6 (Windows) • -32768 = not set
PREFETCHPRIORITY	CHAR (1)		Prefetch priority of the agents in the service class. <ul style="list-style-type: none"> • H = High • L = Low • M = Medium • Blank = not set
BUFFERPOOLPRIORITY	CHAR (1)		Bufferpool priority of the agents in the service class <ul style="list-style-type: none"> • H = High • L = Low • M = Medium • Blank = Not set
INBOUNDCORRELATOR	VARCHAR (128)	Y	For future use.
OUTBOUNDCORRELATOR	VARCHAR (128)	Y	String used to associate the service class with an operating system workload manager service class.
COLLECTAGGACTDATA	CHAR (1)		Specifies what aggregate activity data should be captured for the service class by the applicable event monitor. <ul style="list-style-type: none"> • B = Collect base aggregate activity data • E = Collect extended aggregate activity data • N = None
COLLECTAGGREQDATA	CHAR (1)		Specifies what aggregate activity data should be captured for the service class by the applicable event monitor. <ul style="list-style-type: none"> • B = Collect base aggregate request data • N = None
COLLECTACTDATA	CHAR (1)		Specifies what activity data should be collected by the applicable event monitor. <ul style="list-style-type: none"> • D = Activity data with details • N = None • S = Activity data with details and section environment • V = Activity data with details and values • W = Activity data without details • X = Activity data with details, section environment, and values

Table 182. SYSCAT.SERVICECLASSES Catalog View (continued)

Column Name	Data Type	Nullable	Description
COLLECTACTPARTITION	CHAR (1)		Specifies where activity data is collected. <ul style="list-style-type: none"> • C = Database partition of the coordinator of the activity • D = All database partitions
COLLECTREQMETRICS	CHAR (1)		Specifies the monitoring level for requests submitted by a connection that is associated with the service superclass. <ul style="list-style-type: none"> • B = Collect base request metrics • E = Collect extended request metrics • N = None
REMARKS	VARCHAR (254)	Y	User-provided comments, or the null value.

SYSCAT.THRESHOLDS

Each row represents a threshold.

Table 183. SYSCAT.THRESHOLDS Catalog View

Column Name	Data Type	Nullable	Description
THRESHOLDNAME	VARCHAR (128)		Name of the threshold.
THRESHOLDID	INTEGER		Identifier for the threshold.
ORIGIN	CHAR (1)		Origin of the threshold. <ul style="list-style-type: none"> • U = Threshold was created by a user • W = Threshold was created through a work action set
THRESHOLDCLASS	CHAR (1)		Classification of the threshold. <ul style="list-style-type: none"> • A = Aggregate threshold • C = Activity threshold
THRESHOLDPREDICATE	VARCHAR (15)		Type of the threshold. Possible values are: <ul style="list-style-type: none"> • AGGTEMPSPACE • CONCDBC • CONCWCN • CONCWOC • CONNIDLETIME • CPUTIME • CPUTIMEINSC • DBCONN • ESTSQLCOST • ROWSREAD • ROWSREADINSC • ROWSRET • SCCONN • TEMPSPACE • TOTALTIME
THRESHOLDPREDICATEID	SMALLINT		Identifier for the threshold predicate.

Table 183. SYSCAT.THRESHOLDS Catalog View (continued)

Column Name	Data Type	Nullable	Description
DOMAIN	CHAR (2)		Domain of the threshold. <ul style="list-style-type: none"> • DB = Database • SB = Service subclass • SP = Service superclass • WA = Work action set • WD = Workload definition
DOMAINID	INTEGER		Identifier for the object with which the threshold is associated. This can be a service class, work action or workload unique ID. If this is a database threshold, this value is 0.
ENFORCEMENT	CHAR (1)		Scope of enforcement for the threshold. <ul style="list-style-type: none"> • D = Database • P = Database partition • W = Workload occurrence
QUEUING	CHAR (1)		<ul style="list-style-type: none"> • N = The threshold is not queuing • Y = The threshold is queuing
MAXVALUE	BIGINT		Upper bound specified by the threshold.
QUEUESIZE	INTEGER		If QUEUING is 'Y', the size of the queue. -1 otherwise.
COLLECTACTDATA	CHAR (1)		Specifies what activity data should be collected by the applicable event monitor. <ul style="list-style-type: none"> • D = Activity data with details • N = None • S = Activity data with details and section environment • V = Activity data with details and values • W = Activity data without details • X = Activity data with details, section environment, and values
COLLECTACTPARTITION	CHAR (1)		Specifies where activity data is collected. <ul style="list-style-type: none"> • C = Database partition of the coordinator of the activity • D = All database partitions
EXECUTION	CHAR (1)		Indicates whether execution continues, stops, or remaps to a different service subclass after the threshold has been exceeded. <ul style="list-style-type: none"> • C = Execution continues • R = Execution is remapped • S = Execution stops
REMAPSCID	SMALLINT		Target service subclass ID of the REMAP ACTIVITY action.
VIOLATIONRECORDLOGGED	CHAR (1)		Indicates whether a record is written to the event monitor upon threshold violation. <ul style="list-style-type: none"> • N = No • Y = Yes

Table 183. SYSCAT.THRESHOLDS Catalog View (continued)

Column Name	Data Type	Nullable	Description
CHECKINTERVAL	INTEGER		The interval, in seconds, in which the threshold condition is checked if THRESHOLDPREDICATE is: <ul style="list-style-type: none"> • 'CPUTIME' • 'CPUTIMEINSC' • 'ROWSREAD' • 'ROWSREADINSC' Otherwise, -1.
ENABLED	CHAR (1)		<ul style="list-style-type: none"> • N = This threshold is disabled. • Y = This threshold is enabled.
CREATE_TIME	TIMESTAMP		Time at which the threshold was created.
ALTER_TIME	TIMESTAMP		Time at which the threshold was last altered.
REMARKS	VARCHAR (254)	Y	User-provided comments, or the null value.

SYSCAT.WORKACTIONS

Each row represents a work action that is defined for a work action set.

Table 184. SYSCAT.WORKACTIONS Catalog View

Column Name	Data Type	Nullable	Description
ACTIONNAME	VARCHAR (128)		Name of the work action.
ACTIONID	INTEGER		Identifier for the work action.
ACTIONSETNAME	VARCHAR (128)	Y	Name of the work action set.
ACTIONSETID	INTEGER		Identifier of the work action set to which this work action belongs. This column refers to the ACTIONSETID column in the SYSCAT.WORKACTIONSETS view.
WORKCLASSNAME	VARCHAR (128)	Y	Name of the work class.
WORKCLASSID	INTEGER		Identifier of the work class. This column refers to the WORKCLASSID column in the SYSCAT.WORKCLASSES view.
CREATE_TIME	TIMESTAMP		Time at which the work action was created.
ALTER_TIME	TIMESTAMP		Time at which the work action was last altered.
ENABLED	CHAR (1)		<ul style="list-style-type: none"> • N = This work action is disabled. • Y = This work action is enabled.

Table 184. SYSCAT.WORKACTIONS Catalog View (continued)

Column Name	Data Type	Nullable	Description
ACTIONTYPE	CHAR (1)		<p>The type of action performed on each DB2 activity that matches the attributes in the work class within scope.</p> <ul style="list-style-type: none"> • B = Collect basic aggregate activity data, specifiable only for work action sets that apply to service classes. • C = Allow any DB2 activity under the associated work class to execute and increment the work class counter. • D = Collect activity data with details at the database partition of the coordinator of the activity. • E = Collect extended aggregate activity data, specifiable only for work action sets that apply to service classes. • F = Collect activity data with details, section, and values at the database partition of the coordinator of the activity. • G = Collect activity details and section at the database partition of the coordinator of the activity and collect activity data at all database partitions. • H = Collect activity details, section, and values at the database partition of the coordinator of the activity and collect activity data at all database partitions. • M = Map to a service subclass, specifiable only for work action sets that apply to service classes. • P = Prevent the execution of any DB2 activity under the work class with which this work action is associated. • S = Collect activity data with details and section at the database partition of the coordinator of the activity. • T = The action represents a threshold, specifiable only for work action sets that are associated with a database. • U = Map all activities with a nesting level of zero and all activities nested under these activities to a service subclass, specifiable only for work action sets that apply to service classes. • V = Collect activity data with details and values at the coordinator partition. • W = Collect activity data without details at the coordinator partition. • X = Collect activity data with details at the coordinator partition and collect activity data at all database partitions. • Y = Collect activity data with details and values at the coordinator partition and collect activity data at all database partitions. • Z = Collect activity data without details at all database partitions.

Table 184. SYSCAT.WORKACTIONS Catalog View (continued)

Column Name	Data Type	Nullable	Description
REFOBJECTID	INTEGER	Y	If ACTIONTYPE is 'M' (map) or 'N' (map nested), this value is set to the ID of the service subclass to which the DB2 activity is mapped. If ACTIONTYPE is 'T' (threshold), this value is set to the ID of the threshold to be used. For all other actions, this value is NULL.
REFOBJECTTYPE	VARCHAR (30)		If the ACTIONTYPE is 'M' or 'N', this value is set to 'SERVICE CLASS'; if the ACTIONTYPE is 'T', this value is 'THRESHOLD'; the null value otherwise.

SYSCAT.WORKACTIONSETS

Each row represents a work action set.

Table 185. SYSCAT.WORKACTIONSETS Catalog View

Column Name	Data Type	Nullable	Description
ACTIONSETNAME	VARCHAR (128)		Name of the work action set.
ACTIONSETID	INTEGER		Identifier for the work action set.
WORKCLASSETNAME	VARCHAR (128)	Y	Name of the work class set.
WORKCLASSETID	INTEGER		The identifier of the work class set that is to be mapped to the object specified by the OBJECTID. This column refers to WORKCLASSETID in the SYSCAT.WORKCLASSETS view.
CREATE_TIME	TIMESTAMP		Time at which the work action set was created.
ALTER_TIME	TIMESTAMP		Time at which the work action set was last altered.
ENABLED	CHAR (1)		<ul style="list-style-type: none"> N = This work action set is disabled. Y = This work action set is enabled.
OBJECTTYPE	CHAR (1)		<ul style="list-style-type: none"> b = Service superclass Blank = Database
OBJECTNAME	VARCHAR (128)	Y	Name of the service class.
OBJECTID	INTEGER		The identifier of the object to which the work class set (specified by the WORKCLASSETID) is mapped. If the OBJECTTYPE is blank, the OBJECTID is -1. If the OBJECTTYPE is 'b', the OBJECTID is the ID of the service superclass.
REMARKS	VARCHAR (254)	Y	User-provided comments, or the null value.

SYSCAT.WORKCLASSES

Each row represents a work class defined for a work class set.

Table 186. SYSCAT.WORKCLASSES Catalog View

Column Name	Data Type	Nullable	Description
WORKCLASSNAME	VARCHAR (128)		Name of the work class.
WORKCLASSETNAME	VARCHAR (128)	Y	Name of the work class set.
WORKCLASSID	INTEGER		Identifier for the work class.
WORKCLASSETID	INTEGER		Identifier for the work class set to which this work class belongs. This column refers to the WORKCLASSETID column in the SYSCAT.WORKCLASSETS view.
CREATE_TIME	TIMESTAMP		Time at which the work class was created.
ALTER_TIME	TIMESTAMP		Time at which the work class was last altered.
WORKTYPE	SMALLINT		The type of DB2 activity. <ul style="list-style-type: none"> • 1 = ALL • 2 = READ • 3 = WRITE • 4 = CALL • 5 = DML • 6 = DDL • 7 = LOAD
RANGEUNITS	CHAR (1)		The units to use for the bottom and top range. <ul style="list-style-type: none"> • C = Cardinality • T = Timerons • Blank = Not applicable
FROMVALUE	DOUBLE	Y	The low value of the range in the units specified by the RANGEUNITS. Null value when RANGEUNITS is blank.
TOVALUE	DOUBLE	Y	The high value of the range in the units specified by the RANGEUNITS. Null value when RANGEUNITS is blank. -1 value is used to indicate no upper bound.
ROUTINESHEMA	VARCHAR (128)	Y	Schema name of the procedures that are called from the CALL statement. Null value when WORKTYPE is not 4 (CALL) or 1 (ALL).
INITIALSQLDATAPRIORITY	CHAR (1)		Reserved for future use.
EVALUATIONORDER	SMALLINT		Uniquely identifies the evaluation order used for choosing a work class within a work class set.

SYSCAT.WORKCLASSETS

Each row represents a work class set.

Table 187. SYSCAT.WORKCLASSETS Catalog View

Column Name	Data Type	Nullable	Description
WORKCLASSETNAME	VARCHAR (128)		Name of the work class set.
WORKCLASSETID	INTEGER		Identifier for the work class set.

Table 187. SYSCAT.WORKCLASSETS Catalog View (continued)

Column Name	Data Type	Nullable	Description
CREATE_TIME	TIMESTAMP		Time at which the work class set was created.
ALTER_TIME	TIMESTAMP		Time at which the work class set was last altered.
REMARKS	VARCHAR (254)	Y	User-provided comments, or the null value.

SYSCAT.WORKLOADAUTH

Each row represents a user, group, or role that has been granted USAGE privilege on a workload.

Table 188. SYSCAT.WORKLOADAUTH Catalog View

Column Name	Data Type	Nullable	Description
WORKLOADID	INTEGER		Identifier for the workload.
WORKLOADNAME	VARCHAR (128)		Name of the workload.
GRANTOR	VARCHAR (128)		Grantor of the privilege.
GRANTORTYPE	CHAR (1)		<ul style="list-style-type: none"> • U = Grantee is an individual user
GRANTEE	VARCHAR (128)		Holder of the privilege.
GRANTEETYPE	CHAR (1)		<ul style="list-style-type: none"> • G = Grantee is a group • R = Grantee is a role • U = Grantee is an individual user
USAGEAUTH	CHAR (1)		Indicates whether grantee holds USAGE privilege on the workload. <ul style="list-style-type: none"> • N = Not held • Y = Held

SYSCAT.WORKLOADCONNATTR

Each row represents a connection attribute in the definition of a workload.

Table 189. SYSCAT.WORKLOADCONNATTR Catalog View

Column Name	Data Type	Nullable	Description
WORKLOADID	INTEGER		Identifier for the workload.
WORKLOADNAME	VARCHAR (128)		Name of the workload.

Table 189. SYSCAT.WORKLOADCONNATTR Catalog View (continued)

Column Name	Data Type	Nullable	Description
CONNATTRTYPE	VARCHAR (30)		Type of the connection attribute. <ul style="list-style-type: none"> • 1 = APPLNAME • 2 = SYSTEM_USER • 3 = SESSION_USER • 4 = SESSION_USER GROUP • 5 = SESSION_USER ROLE • 6 = CURRENT CLIENT_USERID • 7 = CURRENT CLIENT_APPLNAME • 8 = CURRENT CLIENT_WRKSTNNAME • 9 = CURRENT CLIENT_ACCTNG • 10 = ADDRESS
CONNATTRVALUE	VARCHAR (1000)		Value of the connection attribute.

SYSCAT.WORKLOADS

Each row represents a workload.

Table 190. SYSCAT.WORKLOADS Catalog View

Column Name	Data Type	Nullable	Description
WORKLOADID	INTEGER		Identifier for the workload.
WORKLOADNAME	VARCHAR (128)		Name of the workload.
EVALUATIONORDER	SMALLINT		Evaluation order used for choosing a workload.
CREATE_TIME	TIMESTAMP		Time at which the workload was created.
ALTER_TIME	TIMESTAMP		Time at which the workload was last altered.
ENABLED	CHAR (1)		<ul style="list-style-type: none"> • N = This workload is disabled. • Y = This workload is enabled.
ALLOWACCESS	CHAR (1)		<ul style="list-style-type: none"> • N = A UOW associated with this workload will be rejected. • Y = A unit of work (UOW) associated with this workload can access the database.
SERVICECLASSNAME	VARCHAR (128)		Name of the service subclass to which a unit of work (associated with this workload) is assigned.
PARENTSERVICECLASSNAME	VARCHAR (128)	Y	Name of the service superclass to which a unit of work (associated with this workload) is assigned.
COLLECTAGGACTDATA	CHAR (1)		<p>Specifies what aggregate activity data should be captured for the workload by the applicable event monitor.</p> <ul style="list-style-type: none"> • B = Collect base aggregate activity data • E = Collect extended aggregate activity data • N = None

Table 190. SYSCAT.WORKLOADS Catalog View (continued)

Column Name	Data Type	Nullable	Description
COLLECTACTDATA	CHAR (1)		<p>Specifies what activity data should be collected by the applicable event monitor.</p> <ul style="list-style-type: none"> • D = Activity data with details • N = None • S = Activity data with details and section environment • V = Activity data with details and values. Applies when the COLLECT column is set to 'C' • W = Activity data without details • X = Activity data with details, section environment, and values
COLLECTACTPARTITION	CHAR (1)		<p>Specifies where activity data is collected.</p> <ul style="list-style-type: none"> • C = Database partition of the coordinator of the activity • D = All database partitions
COLLECTDEADLOCK	CHAR (1)		<p>Specifies that deadlock events should be collect by the applicable event monitor.</p> <ul style="list-style-type: none"> • H = Collect deadlock data with past activities only • N = Do not not collect deadlock data • V = Collect deadlock data with past activities and values • W = Collect deadlock data without past activities and values
COLLECTLOCKTIMEOUT	CHAR (1)		<p>Specifies that lock timeout events should be collect by the applicable event monitor.</p> <ul style="list-style-type: none"> • H = Collect lock timeout data with past activities only • N = Do not not collect lock timeout data • V = Collect lock timeout data with past activities and values • W = Collect lock timeout data without past activities and values
COLLECTLOCKWAIT	CHAR (1)		<p>Specifies that lock wait events should be collect by the applicable event monitor.</p> <ul style="list-style-type: none"> • H = Collect lock wait data with past activities only • N = Do not not collect lock wait data • V = Collect lock wait data with past activities and values • W = Collect lock wait data without past activities and values
LOCKWAITVALUE	INTEGER		<p>Specifies the time in milliseconds a lock should wait before a lock event is collected by the applicable event monitor; 0 when COLLECTLOCKWAIT = 'N'</p>

Table 190. SYSCAT.WORKLOADS Catalog View (continued)

Column Name	Data Type	Nullable	Description
COLLECTACTMETRICS	CHAR (1)		Specifies the monitoring level for activities submitted by an occurrence of the workload. <ul style="list-style-type: none"> • B = Collect base activity metrics • E = Collect extended activity metrics • N = None
COLLECTUOWDATA	CHAR (1)		Specifies what unit of work data should be collected by the applicable event monitor. <ul style="list-style-type: none"> • B = Collect base unit of work data • N = None
EXTERNALNAME	VARCHAR (128)	Y	Reserved for future use.
REMARKS	VARCHAR (254)	Y	User-provided comments, or the null value.

Appendix A. Naming rules

Rules exist for the naming of all database objects, users, groups, files, and paths. Some of these rules are specific to the platform you are working on.

For example, there is a rule regarding the use of upper and lowercase letters in a name.

- On UNIX platforms, names must be in lowercase.
- On Windows platforms, names can be in upper, lower, and mixed-case.

Unless otherwise specified, all names can include the following characters:

- A through Z. When used in most names, characters A through Z are converted from lowercase to uppercase.
- 0 through 9.
- ! % () { } . - ^ ~ _ (underscore) @, #, \$, and space.
- \ (backslash).

Names cannot begin with a number or with the underscore character.

Do not use SQL reserved words to name tables, views, columns, indexes, or authorization IDs.

There are other special characters that might work separately depending on your operating system and where you are working with the DB2 database. However, while they might work, there is no guarantee that they will work. It is not recommended that you use these other special characters when naming objects in your database.

User and group names also must follow the rules forced on specific operation systems by the related systems. For example, on Linux and UNIX platforms, allowed characters for user names and primary group names must be lowercase a through z, 0 through 9, and _ (underscore) for names not starting with 0 through 9.

Lengths must be less than or equal to the lengths listed in “SQL and XML limits” in the *SQL Reference*.

You also must consider object naming rules, naming rules in an NLS environment, and naming rules in a Unicode environment.

Restrictions on the AUTHID identifier: Version 9.5, and later, of the DB2 database system allows you to have an 128-byte authorization ID, but when the authorization ID is interpreted as an operating system user ID or group name, the operating system naming restrictions apply (for example, Linux and UNIX operating systems have a limitation to 8 characters and Windows operating systems have a limitation of 30 characters for user IDs and group names). Therefore, while you can grant an 128-byte authorization ID, it is not possible to connect as a user that has that authorization ID. If you write your own security plugin, you should be able to take full advantage of the extended sizes for the authorization ID. For example, you can give your security plugin a 30-byte user ID and it can return an 128-byte authorization ID during authentication that you are able to connect with.

Appendix B. Roles

Roles simplify the administration and management of privileges by offering an equivalent capability as groups but without the same restrictions.

A role is a database object that groups together one or more privileges and can be assigned to users, groups, PUBLIC, or other roles by using a GRANT statement, or can be assigned to a trusted context by using a CREATE TRUSTED CONTEXT or ALTER TRUSTED CONTEXT statement. A role can be specified for the SESSION_USER ROLE connection attribute in a workload definition.

Roles provide several advantages that make it easier to manage privileges in a database system:

- Security administrators can control access to their databases in a way that mirrors the structure of their organizations (they can create roles in the database that map directly to the job functions in their organizations).
- Users are granted membership in the roles that reflect their job responsibilities. As their job responsibilities change, their membership in roles can be easily granted and revoked.
- The assignment of privileges is simplified. Instead of granting the same set of privileges to each individual user in a particular job function, the administrator can grant this set of privileges to a role representing that job function and then grant that role to each user in that job function.
- A role's privileges can be updated and all users who have been granted that role receive the update; the administrator does not need to update the privileges for every user on an individual basis.
- The privileges and authorities granted to roles are always used when you create views, triggers, materialized query tables (MQTs), static SQL and SQL routines, whereas privileges and authorities granted to groups (directly or indirectly) are not used.

This is because the DB2 database system cannot determine when membership in a group changes, as the group is managed by third-party software (for example, the operating system or an LDAP directory). Because roles are managed inside the database, the DB2 database system can determine when authorization changes and act accordingly. Roles granted to groups are not considered, due to the same reason groups are not considered.

- All the roles assigned to a user are enabled when that user establishes a connection, so all privileges and authorities granted to roles are taken into account when a user connects. Roles cannot be explicitly enabled or disabled.
- The security administrator can delegate management of a role to others.

All DB2 privileges and authorities that can be granted within a database can be granted to a role. For example, a role can be granted any of the following authorities and privileges:

- DBADM, SECADM, DATAACCESS, ACCESSCTRL, SQLADM, WLMADM, LOAD, and IMPLICIT_SCHEMA database authorities
- CONNECT, CREATETAB, CREATE_NOT_FENCED, BINDADD, CREATE_EXTERNAL_ROUTINE, or QUIESCE_CONNECT database authorities
- Any database object privilege (including CONTROL)

A user's roles are automatically enabled and considered for authorization when a user connects to a database; you do not need to activate a role by using the SET ROLE statement. For example, when you create a view, a materialized query table (MQT), a trigger, a package, or an SQL routine, the privileges that you gain through roles apply. However, privileges that you gain through roles granted to groups of which you are a member do not apply.

A role does not have an owner. The security administrator can use the WITH ADMIN OPTION clause of the GRANT statement to delegate management of the role to another user, so that the other user can control the role membership.

Restrictions

There are a few restrictions in the use of roles:

- A role cannot own database objects.
- Permissions and roles granted to groups are not considered when you create the following database objects:
 - Packages containing static SQL
 - Views
 - Materialized query tables (MQT)
 - Triggers
 - SQL Routines

Only roles granted to the user creating the object or to PUBLIC, directly or indirectly (such as through a role hierarchy), are considered when creating these objects.

Appendix C. Trusted contexts and trusted connections

A trusted context is a database object that defines a trust relationship for a connection between the database and an external entity such as an application server.

The trust relationship is based upon the following set of attributes:

- System authorization ID: Represents the user that establishes a database connection
- IP address (or domain name): Represents the host from which a database connection is established
- Data stream encryption: Represents the encryption setting (if any) for the data communication between the database server and the database client

When a user establishes a database connection, the DB2 database system checks whether the connection matches the definition of a trusted context object in the database. When a match occurs, the database connection is said to be trusted.

A trusted connection allows the initiator of this trusted connection to acquire additional capabilities that may not be available outside the scope of the trusted connection. The additional capabilities vary depending on whether the trusted connection is explicit or implicit.

The initiator of an explicit trusted connection has the ability to:

- Switch the current user ID on the connection to a different user ID with or without authentication
- Acquire additional privileges via the role inheritance feature of trusted contexts

An implicit trusted connection is a trusted connection that is not explicitly requested; the implicit trusted connection results from a normal connection request rather than an explicit trusted connection request. No application code changes are needed to obtain an implicit connection. Also, whether you obtain an implicit trusted connection or not has no effect on the connect return code (when you request an explicit trusted connection, the connect return code indicates whether the request succeeds or not). The initiator of an implicit trusted connection can only acquire additional privileges via the role inheritance feature of trusted contexts; they cannot switch the user ID.

How using trusted contexts enhances security

The three-tiered application model extends the standard two-tiered client and server model by placing a middle tier between the client application and the database server. It has gained great popularity in recent years particularly with the emergence of web-based technologies and the Java™ 2 Enterprise Edition (J2EE) platform. An example of a software product that supports the three-tier application model is IBM WebSphere Application Server (WAS).

In a three-tiered application model, the middle tier is responsible for authenticating the users running the client applications and for managing the interactions with the database server. Traditionally, all the interactions with the database server occur through a database connection established by the middle tier using a combination of a user ID and a credential that identify that middle tier to the database server. In other words, the database server uses the database privileges

associated with the middle tier's user ID for all authorization checking and auditing that must occur for any database access, including access performed by the middle tier on behalf of a user.

While the three-tiered application model has many benefits, having all interactions with the database server (for example, a user request) occur under the middle tier's authorization ID raises several security concerns, which can be summarized as follows:

- **Loss of user identity**
Some enterprises prefer to know the identity of the actual user accessing the database for access control purposes.
- **Diminished user accountability**
Accountability through auditing is a basic principle in database security. Not knowing the user's identity makes it difficult to distinguish the transactions performed by the middle tier for its own purpose from those performed by the middle tier on behalf of a user.
- **Over granting of privileges to the middle tier's authorization ID**
The middle tier's authorization ID must have all the privileges necessary to execute all the requests from all the users. This has the security issue of enabling users who do not need access to certain information to obtain access anyway.
- **Weakened security**
In addition to the privilege issue raised in the previous point, the current approach requires that the authorization ID used by the middle tier to connect must be granted privileges on all resources that might be accessed by user requests. If that middle-tier authorization ID is ever compromised, then all those resources will be exposed.
- **"Spill over" between users of the same connection**
Changes by a previous user can affect the current user.

Clearly, there is a need for a mechanism whereby the actual user's identity and database privileges are used for database requests performed by the middle tier on behalf of that user. The most straightforward approach of achieving this goal would be for the middle-tier to establish a new connection using the user's ID and password, and then direct the user's requests through that connection. Although simple, this approach suffers from several drawbacks which include the following:

- **Inapplicability for certain middle tiers.** Many middle-tier servers do not have the user authentication credentials needed to establish a connection.
- **Performance overhead.** There is an obvious performance overhead associated with creating a new physical connection and re-authenticating the user at the database server.
- **Maintenance overhead.** In situations where you are not using a centralized security set up or are not using single sign-on, there is maintenance overhead in having two user definitions (one on the middle tier and one at the server). This requires changing passwords at different places.

The trusted contexts capability addresses this problem. The security administrator can create a trusted context object in the database that defines a trust relationship between the database and the middle-tier. The middle-tier can then establish an explicit trusted connection to the database, which gives the middle tier the ability to switch the current user ID on the connection to a different user ID, with or without authentication. In addition to solving the end-user identity assertion problem, trusted contexts offer another advantage. This is the ability to control when a privilege is made available to a database user. The lack of control on when privileges are available to a user can weaken overall security. For example,

privileges may be used for purposes other than they were originally intended. The security administrator can assign one or more privileges to a role and assign that role to a trusted context object. Only trusted database connections (explicit or implicit) that match the definition of that trusted context can take advantage of the privileges associated with that role.

Enhancing performance

When you use trusted connections, you can maximize performance because of the following advantages:

- No new connection is established when the current user ID of the connection is switched.
- If the trusted context definition does not require authentication of the user ID to switch to, then the overhead associated with authenticating a new user at the database server is not incurred.

Example of creating a trusted context

Suppose that the security administrator creates the following trusted context object:

```
CREATE TRUSTED CONTEXT CTX1
  BASED UPON CONNECTION USING SYSTEM AUTHID USER2
  ATTRIBUTES (ADDRESS '192.0.2.1')
  DEFAULT ROLE managerRole
  ENABLE
```

If user *user1* requests a trusted connection from IP address 192.0.2.1, the DB2 database system returns a warning (SQLSTATE 01679, SQLCODE +20360) to indicate that a trusted connection could not be established, and that user *user1* simply got a non-trusted connection. However, if user *user2* requests a trusted connection from IP address 192.0.2.1, the request is honored because the connection attributes are satisfied by the trusted context CTX1. Now that user *user2* has established a trusted connection, he or she can now acquire all the privileges and authorities associated with the trusted context role *managerRole*. These privileges and authorities may not be available to user *user2* outside the scope of this trusted connection.

Appendix D. Overview of the DB2 technical information

DB2 technical information is available through the following tools and methods:

- DB2 Information Center
 - Topics (Task, concept and reference topics)
 - Help for DB2 tools
 - Sample programs
 - Tutorials
- DB2 books
 - PDF files (downloadable)
 - PDF files (from the DB2 PDF DVD)
 - printed books
- Command line help
 - Command help
 - Message help

Note: The DB2 Information Center topics are updated more frequently than either the PDF or the hardcopy books. To get the most current information, install the documentation updates as they become available, or refer to the DB2 Information Center at ibm.com.

You can access additional DB2 technical information such as technotes, white papers, and IBM Redbooks® publications online at ibm.com. Access the DB2 Information Management software library site at <http://www.ibm.com/software/data/sw-library/>.

Documentation feedback

We value your feedback on the DB2 documentation. If you have suggestions for how to improve the DB2 documentation, send an e-mail to db2docs@ca.ibm.com. The DB2 documentation team reads all of your feedback, but cannot respond to you directly. Provide specific examples wherever possible so that we can better understand your concerns. If you are providing feedback on a specific topic or help file, include the topic title and URL.

Do not use this e-mail address to contact DB2 Customer Support. If you have a DB2 technical issue that the documentation does not resolve, contact your local IBM service center for assistance.

DB2 technical library in hardcopy or PDF format

The following tables describe the DB2 library available from the IBM Publications Center at www.ibm.com/shop/publications/order. English and translated DB2 Version 9.7 manuals in PDF format can be downloaded from www.ibm.com/support/docview.wss?rs=71&uid=swg2700947.

Although the tables identify books available in print, the books might not be available in your country or region.

The form number increases each time a manual is updated. Ensure that you are reading the most recent version of the manuals, as listed below.

Note: The *DB2 Information Center* is updated more frequently than either the PDF or the hard-copy books.

Table 191. DB2 technical information

Name	Form Number	Available in print	Last updated
<i>Administrative API Reference</i>	SC27-2435-00	Yes	August, 2009
<i>Administrative Routines and Views</i>	SC27-2436-00	No	August, 2009
<i>Call Level Interface Guide and Reference, Volume 1</i>	SC27-2437-00	Yes	August, 2009
<i>Call Level Interface Guide and Reference, Volume 2</i>	SC27-2438-00	Yes	August, 2009
<i>Command Reference</i>	SC27-2439-00	Yes	August, 2009
<i>Data Movement Utilities Guide and Reference</i>	SC27-2440-00	Yes	August, 2009
<i>Data Recovery and High Availability Guide and Reference</i>	SC27-2441-00	Yes	August, 2009
<i>Database Administration Concepts and Configuration Reference</i>	SC27-2442-00	Yes	August, 2009
<i>Database Monitoring Guide and Reference</i>	SC27-2458-00	Yes	August, 2009
<i>Database Security Guide</i>	SC27-2443-00	Yes	August, 2009
<i>DB2 Text Search Guide</i>	SC27-2459-00	Yes	August, 2009
<i>Developing ADO.NET and OLE DB Applications</i>	SC27-2444-00	Yes	August, 2009
<i>Developing Embedded SQL Applications</i>	SC27-2445-00	Yes	August, 2009
<i>Developing Java Applications</i>	SC27-2446-00	Yes	August, 2009
<i>Developing Perl, PHP, Python, and Ruby on Rails Applications</i>	SC27-2447-00	No	August, 2009
<i>Developing User-defined Routines (SQL and External)</i>	SC27-2448-00	Yes	August, 2009
<i>Getting Started with Database Application Development</i>	GI11-9410-00	Yes	August, 2009
<i>Getting Started with DB2 Installation and Administration on Linux and Windows</i>	GI11-9411-00	Yes	August, 2009

Table 191. DB2 technical information (continued)

Name	Form Number	Available in print	Last updated
<i>Globalization Guide</i>	SC27-2449-00	Yes	August, 2009
<i>Installing DB2 Servers</i>	GC27-2455-00	Yes	August, 2009
<i>Installing IBM Data Server Clients</i>	GC27-2454-00	No	August, 2009
<i>Message Reference Volume 1</i>	SC27-2450-00	No	August, 2009
<i>Message Reference Volume 2</i>	SC27-2451-00	No	August, 2009
<i>Net Search Extender Administration and User's Guide</i>	SC27-2469-00	No	August, 2009
<i>Partitioning and Clustering Guide</i>	SC27-2453-00	Yes	August, 2009
<i>pureXML Guide</i>	SC27-2465-00	Yes	August, 2009
<i>Query Patroller Administration and User's Guide</i>	SC27-2467-00	No	August, 2009
<i>Spatial Extender and Geodetic Data Management Feature User's Guide and Reference</i>	SC27-2468-00	No	August, 2009
<i>SQL Procedural Languages: Application Enablement and Support</i>	SC27-2470-00	Yes	August, 2009
<i>SQL Reference, Volume 1</i>	SC27-2456-00	Yes	August, 2009
<i>SQL Reference, Volume 2</i>	SC27-2457-00	Yes	August, 2009
<i>Troubleshooting and Tuning Database Performance</i>	SC27-2461-00	Yes	August, 2009
<i>Upgrading to DB2 Version 9.7</i>	SC27-2452-00	Yes	August, 2009
<i>Visual Explain Tutorial</i>	SC27-2462-00	No	August, 2009
<i>What's New for DB2 Version 9.7</i>	SC27-2463-00	Yes	August, 2009
<i>Workload Manager Guide and Reference</i>	SC27-2464-00	Yes	August, 2009
<i>XQuery Reference</i>	SC27-2466-00	No	August, 2009

Table 192. DB2 Connect-specific technical information

Name	Form Number	Available in print	Last updated
<i>Installing and Configuring DB2 Connect Personal Edition</i>	SC27-2432-00	Yes	August, 2009
<i>Installing and Configuring DB2 Connect Servers</i>	SC27-2433-00	Yes	August, 2009

Table 192. DB2 Connect-specific technical information (continued)

Name	Form Number	Available in print	Last updated
DB2 Connect User's Guide	SC27-2434-00	Yes	August, 2009

Table 193. Information Integration technical information

Name	Form Number	Available in print	Last updated
Information Integration: Administration Guide for Federated Systems	SC19-1020-02	Yes	August, 2009
Information Integration: ASNCLP Program Reference for Replication and Event Publishing	SC19-1018-04	Yes	August, 2009
Information Integration: Configuration Guide for Federated Data Sources	SC19-1034-02	No	August, 2009
Information Integration: SQL Replication Guide and Reference	SC19-1030-02	Yes	August, 2009
Information Integration: Introduction to Replication and Event Publishing	GC19-1028-02	Yes	August, 2009

Ordering printed DB2 books

If you require printed DB2 books, you can buy them online in many but not all countries or regions. You can always order printed DB2 books from your local IBM representative. Keep in mind that some softcopy books on the *DB2 PDF Documentation DVD* are unavailable in print. For example, neither volume of the *DB2 Message Reference* is available as a printed book.

Printed versions of many of the DB2 books available on the *DB2 PDF Documentation DVD* can be ordered for a fee from IBM. Depending on where you are placing your order from, you may be able to order books online, from the IBM Publications Center. If online ordering is not available in your country or region, you can always order printed DB2 books from your local IBM representative. Note that not all books on the *DB2 PDF Documentation DVD* are available in print.

Note: The most up-to-date and complete DB2 documentation is maintained in the DB2 Information Center at <http://publib.boulder.ibm.com/infocenter/db2luw/v9r7>.

To order printed DB2 books:

- To find out whether you can order printed DB2 books online in your country or region, check the IBM Publications Center at <http://www.ibm.com/shop/publications/order>. You must select a country, region, or language to access publication ordering information and then follow the ordering instructions for your location.
- To order printed DB2 books from your local IBM representative:

1. Locate the contact information for your local representative from one of the following Web sites:
 - The IBM directory of world wide contacts at www.ibm.com/planetwide
 - The IBM Publications Web site at <http://www.ibm.com/shop/publications/order>. You will need to select your country, region, or language to the access appropriate publications home page for your location. From this page, follow the "About this site" link.
2. When you call, specify that you want to order a DB2 publication.
3. Provide your representative with the titles and form numbers of the books that you want to order. For titles and form numbers, see "DB2 technical library in hardcopy or PDF format" on page 387.

Displaying SQL state help from the command line processor

DB2 products return an SQLSTATE value for conditions that can be the result of an SQL statement. SQLSTATE help explains the meanings of SQL states and SQL state class codes.

To start SQL state help, open the command line processor and enter:

```
? sqlstate or ? class code
```

where *sqlstate* represents a valid five-digit SQL state and *class code* represents the first two digits of the SQL state.

For example, ? 08003 displays help for the 08003 SQL state, and ? 08 displays help for the 08 class code.

Accessing different versions of the DB2 Information Center

For DB2 Version 9.7 topics, the DB2 Information Center URL is <http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/>

For DB2 Version 9.5 topics, the DB2 Information Center URL is <http://publib.boulder.ibm.com/infocenter/db2luw/v9r5/>

For DB2 Version 9 topics, the DB2 Information Center URL is <http://publib.boulder.ibm.com/infocenter/db2luw/v9/>

For DB2 Version 8 topics, go to the Version 8 Information Center URL at: <http://publib.boulder.ibm.com/infocenter/db2luw/v8/>

Displaying topics in your preferred language in the DB2 Information Center

The DB2 Information Center attempts to display topics in the language specified in your browser preferences. If a topic has not been translated into your preferred language, the DB2 Information Center displays the topic in English.

- To display topics in your preferred language in the Internet Explorer browser:
 1. In Internet Explorer, click the **Tools** —> **Internet Options** —> **Languages...** button. The Language Preferences window opens.
 2. Ensure your preferred language is specified as the first entry in the list of languages.
 - To add a new language to the list, click the **Add...** button.

Note: Adding a language does not guarantee that the computer has the fonts required to display the topics in the preferred language.

- To move a language to the top of the list, select the language and click the **Move Up** button until the language is first in the list of languages.
- 3. Clear the browser cache and then refresh the page to display the DB2 Information Center in your preferred language.
- To display topics in your preferred language in a Firefox or Mozilla browser:
 1. Select the button in the **Languages** section of the **Tools** —> **Options** —> **Advanced** dialog. The Languages panel is displayed in the Preferences window.
 2. Ensure your preferred language is specified as the first entry in the list of languages.
 - To add a new language to the list, click the **Add...** button to select a language from the Add Languages window.
 - To move a language to the top of the list, select the language and click the **Move Up** button until the language is first in the list of languages.
 3. Clear the browser cache and then refresh the page to display the DB2 Information Center in your preferred language.

On some browser and operating system combinations, you must also change the regional settings of your operating system to the locale and language of your choice.

Updating the DB2 Information Center installed on your computer or intranet server

A locally installed DB2 Information Center must be updated periodically.

Before you begin

A DB2 Version 9.7 Information Center must already be installed. For details, see the “Installing the DB2 Information Center using the DB2 Setup wizard” topic in *Installing DB2 Servers*. All prerequisites and restrictions that applied to installing the Information Center also apply to updating the Information Center.

About this task

An existing DB2 Information Center can be updated automatically or manually:

- Automatic updates - updates existing Information Center features and languages. An additional benefit of automatic updates is that the Information Center is unavailable for a minimal period of time during the update. In addition, automatic updates can be set to run as part of other batch jobs that run periodically.
- Manual updates - should be used when you want to add features or languages during the update process. For example, a local Information Center was originally installed with both English and French languages, and now you want to also install the German language; a manual update will install German, as well as, update the existing Information Center features and languages. However, a manual update requires you to manually stop, update, and restart the Information Center. The Information Center is unavailable during the entire update process.

Procedure

This topic details the process for automatic updates. For manual update instructions, see the “Manually updating the DB2 Information Center installed on your computer or intranet server” topic.

To automatically update the DB2 Information Center installed on your computer or intranet server:

1. On Linux operating systems,
 - a. Navigate to the path where the Information Center is installed. By default, the DB2 Information Center is installed in the /opt/ibm/db2ic/V9.7 directory.
 - b. Navigate from the installation directory to the doc/bin directory.
 - c. Run the ic-update script:

```
ic-update
```
2. On Windows operating systems,
 - a. Open a command window.
 - b. Navigate to the path where the Information Center is installed. By default, the DB2 Information Center is installed in the <Program Files>\IBM\DB2 Information Center\Version 9.7 directory, where <Program Files> represents the location of the Program Files directory.
 - c. Navigate from the installation directory to the doc\bin directory.
 - d. Run the ic-update.bat file:

```
ic-update.bat
```

Results

The DB2 Information Center restarts automatically. If updates were available, the Information Center displays the new and updated topics. If Information Center updates were not available, a message is added to the log. The log file is located in doc\eclipse\configuration directory. The log file name is a randomly generated number. For example, 1239053440785.log.

Manually updating the DB2 Information Center installed on your computer or intranet server

If you have installed the DB2 Information Center locally, you can obtain and install documentation updates from IBM.

Updating your locally-installed DB2 Information Center manually requires that you:

1. Stop the DB2 Information Center on your computer, and restart the Information Center in stand-alone mode. Running the Information Center in stand-alone mode prevents other users on your network from accessing the Information Center, and allows you to apply updates. The Workstation version of the DB2 Information Center always runs in stand-alone mode. .
2. Use the Update feature to see what updates are available. If there are updates that you must install, you can use the Update feature to obtain and install them

Note: If your environment requires installing the DB2 Information Center updates on a machine that is not connected to the internet, mirror the update site to a local file system using a machine that is connected to the internet and has the DB2 Information Center installed. If many users on your network will be installing the documentation updates, you can reduce the time required for

individuals to perform the updates by also mirroring the update site locally and creating a proxy for the update site.

If update packages are available, use the Update feature to get the packages. However, the Update feature is only available in stand-alone mode.

3. Stop the stand-alone Information Center, and restart the DB2 Information Center on your computer.

Note: On Windows 2008, Windows Vista (and higher), the commands listed later in this section must be run as an administrator. To open a command prompt or graphical tool with full administrator privileges, right-click the shortcut and then select **Run as administrator**.

To update the DB2 Information Center installed on your computer or intranet server:

1. Stop the DB2 Information Center.
 - On Windows, click **Start** → **Control Panel** → **Administrative Tools** → **Services**. Then right-click **DB2 Information Center** service and select **Stop**.
 - On Linux, enter the following command:

```
/etc/init.d/db2icdv97 stop
```


2. Start the Information Center in stand-alone mode.

- On Windows:
 - a. Open a command window.
 - b. Navigate to the path where the Information Center is installed. By default, the DB2 Information Center is installed in the <Program Files>\IBM\DB2 Information Center\Version 9.7 directory, where <Program Files> represents the location of the Program Files directory.
 - c. Navigate from the installation directory to the doc\bin directory.
 - d. Run the help_start.bat file:

```
help_start.bat
```
- On Linux:
 - a. Navigate to the path where the Information Center is installed. By default, the DB2 Information Center is installed in the /opt/ibm/db2ic/V9.7 directory.
 - b. Navigate from the installation directory to the doc/bin directory.
 - c. Run the help_start script:

```
help_start
```

The systems default Web browser opens to display the stand-alone Information Center.

3. Click the **Update** button (). (JavaScript™ must be enabled in your browser.) On the right panel of the Information Center, click **Find Updates**. A list of updates for existing documentation displays.
4. To initiate the installation process, check the selections you want to install, then click **Install Updates**.
5. After the installation process has completed, click **Finish**.
6. Stop the stand-alone Information Center:
 - On Windows, navigate to the installation directory's doc\bin directory, and run the help_end.bat file:

```
help_end.bat
```


Note: The help_end batch file contains the commands required to safely stop the processes that were started with the help_start batch file. Do not use Ctrl-C or any other method to stop help_start.bat.

- On Linux, navigate to the installation directory's doc/bin directory, and run the help_end script:
help_end

Note: The help_end script contains the commands required to safely stop the processes that were started with the help_start script. Do not use any other method to stop the help_start script.

7. Restart the DB2 Information Center.

- On Windows, click **Start** → **Control Panel** → **Administrative Tools** → **Services**. Then right-click **DB2 Information Center** service and select **Start**.
- On Linux, enter the following command:
/etc/init.d/db2icdv97 start

The updated DB2 Information Center displays the new and updated topics.

DB2 tutorials

The DB2 tutorials help you learn about various aspects of DB2 products. Lessons provide step-by-step instructions.

Before you begin

You can view the XHTML version of the tutorial from the Information Center at <http://publib.boulder.ibm.com/infocenter/db2help/>.

Some lessons use sample data or code. See the tutorial for a description of any prerequisites for its specific tasks.

DB2 tutorials

To view the tutorial, click the title.

“pureXML®” in *pureXML Guide*

Set up a DB2 database to store XML data and to perform basic operations with the native XML data store.

“Visual Explain” in *Visual Explain Tutorial*

Analyze, optimize, and tune SQL statements for better performance using Visual Explain.

DB2 troubleshooting information

A wide variety of troubleshooting and problem determination information is available to assist you in using DB2 database products.

DB2 documentation

Troubleshooting information can be found in the *DB2 Troubleshooting Guide* or the Database fundamentals section of the *DB2 Information Center*. There you will find information about how to isolate and identify problems using DB2 diagnostic tools and utilities, solutions to some of the most common problems, and other advice on how to solve problems you might encounter with your DB2 database products.

DB2 Technical Support Web site

Refer to the DB2 Technical Support Web site if you are experiencing problems and want help finding possible causes and solutions. The Technical Support site has links to the latest DB2 publications, TechNotes, Authorized Program Analysis Reports (APARs or bug fixes), fix packs, and other resources. You can search through this knowledge base to find possible solutions to your problems.

Access the DB2 Technical Support Web site at http://www.ibm.com/software/data/db2/support/db2_9/

Terms and Conditions

Permissions for the use of these publications is granted subject to the following terms and conditions.

Personal use: You may reproduce these Publications for your personal, non commercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these Publications, or any portion thereof, without the express consent of IBM.

Commercial use: You may reproduce, distribute and display these Publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these Publications, or reproduce, distribute or display these Publications or any portion thereof outside your enterprise, without the express consent of IBM.

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the Publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the Publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Appendix E. Notices

This information was developed for products and services offered in the U.S.A. Information about non-IBM products is based on information available at the time of first publication of this document and is subject to change.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
3-2-12, Roppongi, Minato-ku, Tokyo 106-8711 Japan

The following paragraph does not apply to the United Kingdom or any other country/region where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web

sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information that has been exchanged, should contact:

IBM Canada Limited
Office of the Lab Director
8200 Warden Avenue
Markham, Ontario
L6G 1C7
CANADA

Such information may be available, subject to appropriate terms and conditions, including, in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems, and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements, or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information 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.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application

programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (*your company name*) (*year*). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. *_enter the year or years_*. All rights reserved.

Trademarks

IBM, the IBM logo, and `ibm.com`[®] are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

The following terms are trademarks or registered trademarks of other companies

- Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
- Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
- UNIX is a registered trademark of The Open Group in the United States and other countries.
- Intel[®], Intel logo, Intel Inside[®], Intel Inside logo, Intel[®] Centrino[®], Intel Centrino logo, Celeron[®], Intel[®] Xeon[®], Intel SpeedStep[®], Itanium[®], and Pentium[®] are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
- Microsoft[®], Windows, Windows NT[®], and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

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

Index

A

- activation time
 - last_wlm_reset monitor element 344
- activities
 - analysis example 57
 - application of work actions to 133
 - assignment to work classes 51
 - business intelligence reports 120, 122
 - cancelling 197
 - cancelling example 267
 - capturing for later analysis example 199
 - collecting data about 195
 - identifying long-running activities
 - overview 267
 - importing information into the Design Advisor 197
 - low estimated cost and high runtime 276
 - mapping to service classes 68
 - monitor elements
 - act_total 328
 - activity_collected 324
 - activity_id 324
 - activity_secondary_id 325
 - activity_type 325
 - coord_act_aborted_total 331
 - coord_act_completed_total 331
 - coord_act_rejected_total 336
 - parent_activity_id 345
 - overviewnested 13
 - queuing 102
 - rogue 198
 - states in service class 74
 - thresholds 95
- activity event monitor 160
- ACTIVITYTOTALTIME activity threshold 96
- agents
 - investigating usage by service class 88
 - priority 72
- aggregate thresholds
 - definition 102
- aggregating data 159
- AGGSQLTEMPSPACE aggregate threshold
 - description 103
- AIX Workload Manager
 - integrating DB2 workload manager 203
 - processor priority 72
- APIs
 - sqleseti
 - workload assignment 35
- assignment of connections to workloads 21
- AUTHID identifier
 - restrictions 379

B

- bins
 - purpose 186
- books
 - printed
 - ordering 390

- buffer pool priority
 - service classes 73

C

- CALL statement
 - classification by schema 48
- catalog views
 - HISTOGRAMTEMPLATEBINS 366
 - HISTOGRAMTEMPLATES 366
 - HISTOGRAMTEMPLATEUSE 366
 - SERVICECLASSES 367
 - THRESHOLDS 369
 - WORKACTIONS 371
 - WORKACTIONSETS 373
 - WORKCLASSES 373
 - WORKCLASSESSETS 374
 - WORKLOADAUTH 375
 - WORKLOADCONNATTR 375
 - WORKLOADS 376
- commands
 - SET WORKLOAD 26, 364
- CONCURRENTDBCOORDACTIVITIES aggregate
 - threshold 103
- CONCURRENTWORKLOADACTIVITIES aggregate
 - threshold 105
- CONCURRENTWORKLOADOCCURRENCES aggregate
 - threshold 106
- configuration parameters
 - wlm_collect_int configuration parameter 365
- CONNECTIONIDLETIME connection threshold 94
- connections
 - assigning to default administration workload 26
 - assignment
 - workload connection attributes with multiple values 44
 - assignment to workload 21
 - mapping to workloads
 - example 39
 - states in service class 74
 - transient 108
- CPUTIME activity threshold
 - description 96
- CPUTIMEINSC activity threshold
 - description 97
- creating
 - service classes 76
 - threshold 109
 - workload 28

D

- data
 - aggregation 159
- Data Definition Language (DDL)
 - statements
 - DB2 workload manager 16
- database objects
 - DB2 workload manager 16
 - roles 381

- DB2 Information Center
 - languages 391
 - updating 392, 393
 - versions 391
 - viewing in different languages 391
- DB2 workload manager
 - activities 95
 - analyzing workloads by activity type (example) 57
 - assignment of work actions 133
 - assignment to work classes 51
 - canceling 197
 - capturing information for analysis (example) 199
 - collecting data 195
 - controlling business intelligence reports (scenario) 120, 122
 - identifying activities with low estimated costs and high runtimes (example) 276
 - identifying long-running activities (example) 267
 - importing information into Design Advisor 197
 - mapping to service classes 68
 - overview 13
 - rogue 198
 - activity queues 102
 - agents
 - investigating usage by service class (example) 88
 - priority 72
 - AIX Workload Manager integration 203
 - concepts 1
 - connections
 - assignment to workloads 21
 - states in service classes 74
 - controlling work 63
 - DDL statements 16
 - event monitors
 - overview 160
 - frequently asked questions 4
 - histogram templates
 - altering 185
 - creating 185
 - dropping 186
 - histograms
 - computing averages and standard deviation (example) 186
 - overview 180
 - historical analysis tool 188
 - identification of work 13
 - introduction
 - tutorial 215
 - Linux workload management integration 208
 - management stage 63, 88, 125
 - metrics 192
 - monitoring
 - data 166
 - event monitors 160
 - monitoring system behavior at different levels (example) 155
 - overview 149
 - real-time 149
 - service class tiers 178
 - monitoring work 149
 - object ownership 16
 - performance modeling 199
 - priority aging 113
 - sample scripts 117
 - Query Patroller 4
 - remapping activities
 - sample scripts 117
- DB2 workload manager (*continued*)
 - sample application 265
 - service class tiering 113
 - sample scripts 117
 - service classes 63
 - altering 78
 - analyzing system slowdown (example) 86
 - creating 76
 - dropping 80
 - entities not tracked 75
 - example 81
 - obtaining point-in-time statistics (example) 158
 - prefetch priority 72, 73
 - service subclasses 63
 - stages 1
 - statistics 169
 - collecting using statistics event monitor 188
 - resetting 190
 - stored procedures
 - WLM_CANCEL_ACTIVITY 168
 - WLM_CAPTURE_ACTIVITY_IN_PROGRESS 168
 - WLM_COLLECT_STATS 168
 - WLM_SET_CLIENT_INFO 168
 - table functions
 - aggregating data (example) 159
 - operational information 149
 - understanding what is running on data server (example) 154
 - using with snapshot monitor table functions 193
 - thresholds 88, 95
 - ACTIVITYTOTALTIME 96
 - aggregate 102
 - AGGSQLTEMPSPACE 103
 - altering 110
 - CONCURRENTDBCOORDACTIVITIES 103
 - CONCURRENTWORKLOADACTIVITIES 105
 - CONCURRENTWORKLOADOCCURRENCES 106
 - connection 94
 - CONNECTIONIDLETIME 94
 - CPUTIME 96
 - CPUTIMEINSC 97
 - creating 109
 - dropping 110
 - ESTIMATEDSQLCOST 98
 - evaluation order 92
 - managing database resources across departments (example) 111
 - monitoring violations 194
 - overview 88
 - remapping activities 113, 123
 - scope resolution 95
 - SQLROWSREAD 99
 - SQLROWSREADINSC 100
 - SQLROWSRETURNED 101
 - SQLTEMPSPACE 102
 - summary 91
 - TOTALDBPARTITIONCONNECTIONS 107
 - TOTALSCPARTITIONCONNECTIONS 108
 - using work action set and database threshold (example) 145
 - tuning
 - with capacity planning data (example) 269
 - without capacity planning data (example) 271
 - tutorial 215
 - basic monitoring 215
 - canceling activities 249
 - capturing activity information 253

- DB2 workload manager (*continued*)
 - tutorial (*continued*)
 - differentiating activities 230
 - discovering activities 250
 - generating historical data and reports 255
 - investigating delay 246
 - isolating activities 221
 - using extended aggregates 258
 - using histograms 237
 - using thresholds 227
 - unit of work
 - workload assignment when multiple workloads exist (example) 41
- USAGE privilege on workloads
 - granting 33
 - revoking 34
- WLMADM authority
 - overview 3
- work action sets 125
 - altering 138
 - creating 137
 - determining types of work being run (example) 147
 - disabling 139
 - dropping 140
 - overview 128
 - using with database threshold (example) 145
- work actions
 - altering 143
 - assignment to database activities 133
 - creating 140
 - defining for work action set 130
 - disabling 145
 - dropping 145
 - thresholds 133
- work class sets
 - altering 56
 - creating 55
 - dropping 56
 - managing DML activities (example) 58
 - overview 48
- work classes
 - altering 55
 - classification of CALL statement by schema 48
 - creating 52
 - defined with the ALL keyword (example) 59
 - dropping 55
 - evaluation order 50
 - overview 45
- workloads
 - altering 29
 - analyzing system slowdown (example) 266
 - assignment (example) 35
 - assignment when multiple workloads exist (example) 41
 - assignment when workload attributes have multiple values (example) 44
 - connection assignment to the default administration workload 26
 - creating 28
 - default 23
 - disabling 33
 - dropping 34
 - enabling 32
 - overview 17
 - permitting database access 31
 - preventing database access 31

- default maintenance service superclass
 - overview 66
- default system service superclass
 - overview 66
- default user service superclass
 - overview 66
- default workloads 23
- Design Advisor
 - importing activity information 197
- documentation
 - overview 387
 - PDF 387
 - printed 387
 - terms and conditions of use 396
- dropping
 - histogram templates 186
 - service classes 80
 - thresholds 110
 - work action sets 140
 - work class sets 56
 - workloads 34

E

- enforcement scopes
 - thresholds 91
- ESTIMATEDSQLCOST activity threshold
 - description 98
- evaluation order
 - DB2 workload manager thresholds 92
 - workloads 21
- event monitors
 - activity data 195
 - DB2 workload manager
 - statistics collection 188
 - threshold violations 194
 - types 160
- examples
 - mapping connections to workloads 39
 - service classes 81
 - work action set and threshold 145
 - work class defined with ALL keyword 59
 - work class set management of activities 58

F

- functions
 - table
 - WLM_GET_QUEUE_STATS 288
 - WLM_GET_SERVICE_CLASS_AGENTS_V97 291
 - WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97 298
 - WLM_GET_SERVICE_SUBCLASS_STATS_V97 302
 - WLM_GET_SERVICE_SUPERCLASS_STATS 310
 - WLM_GET_WORK_ACTION_SET_STATS 311
 - WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97 313
 - WLM_GET_WORKLOAD_STATS_V97 318
 - table functions
 - WLM_GET_ACTIVITY_DETAILS 281

G

- granting
 - USAGE privilege on workload 33

H

- help
 - configuring language 391
 - SQL statements 391
- histogram templates
 - altering 185, 190
 - creating 185
 - dropping 186
- histograms
 - example 186
 - monitor elements
 - histogram_type 342
 - number_in_bin 345
 - top 359
 - overview 180

I

- identifiers
 - monitor elements
 - arm_correlator 329
 - bin_id 329
 - db_work_action_set_id 341
 - db_work_class_id 342
 - sc_work_action_set_id 351
 - sc_work_class_id 351
 - service_class_id 352
 - work_action_set_id 361
 - work_class_id 361
- in-service-class thresholds
 - definition 113
- integration of DB2 and operating system workload managers 203

L

- Linux workload management
 - integrating DB2 workload manager 208

M

- metrics
 - DB2 workload manager objects 192
- monitor elements
 - act_remapped_in 327
 - act_remapped_out 327
 - activation time
 - last_wlm_reset 344
 - activities
 - act_total 328
 - activity_collected 324
 - activity_id 324
 - activity_secondary_id 325
 - activity_type 325
 - coord_act_aborted_total 331
 - coord_act_completed_total 331
 - coord_act_rejected_total 336
 - parent_activity_id 345
 - agg_temp_tablespace_top 328
 - coord_act_est_cost_avg 332, 340
 - coord_act_exec_time_avg 332, 339
 - coord_act_interarrival_time_avg 333, 341
 - coord_act_lifetime_avg 334, 337
 - coord_act_queue_time_avg 335, 338
 - destination_service_class_id 342

- monitor elements (*continued*)
 - executing
 - act_exec_time 326
 - histograms
 - histogram_type 342
 - number_in_bin 345
 - top 359
 - identifiers
 - arm_correlator 329
 - bin_id 329
 - db_work_action_set_id 341
 - db_work_class_id 342
 - sc_work_action_set_id 351
 - sc_work_class_id 351
 - service_class_id 352
 - work_action_set_id 361
 - work_class_id 361
 - names
 - service_subclass_name 353
 - service_superclass_name 354
 - work_action_set_name 361
 - work_class_name 362
 - num_remaps 344
 - partitions
 - coord_partition_num 336
 - queries
 - queue_assignments_total 346
 - queue_size_top 347
 - queue_time_total 347
 - ranges
 - bottom 329
 - request_exec_time_avg 339, 347
 - rows
 - rows_fetched 348
 - rows_modified 348
 - rows_returned 349
 - sections
 - section_env 352
 - source_service_class_id 354
 - thresholds
 - num_threshold_violations 344
 - threshold_action 356
 - threshold_domain 356
 - threshold_maxvalue 357
 - threshold_name 357
 - threshold_predicate 357
 - threshold_queuesize 357
 - thresholdid 358
 - time
 - prep_time 346
 - time_completed 358
 - time_created 358
 - time_of_violation 359
 - time_started 359
 - timestamps
 - activate_timestamp 323
 - statistics_timestamp 355
 - units of work (UOW)
 - parent_uow_id 345
 - uow_id 360
 - watermarks
 - act_cpu_time_top 326
 - act_rows_read_top 327
 - concurrent_act_top 329
 - concurrent_connection_top 330
 - concurrent_wlo_act_top 330
 - concurrent_wlo_top 330

- monitor elements (*continued*)
 - watermarks (*continued*)
 - coord_act_lifetime_top 334
 - cost_estimate_top 337
 - rows_returned_top 350
 - temp_tablespace_top 355
 - uow_total_time_top 360
 - Workload management
 - overview 323
 - workloads
 - wlo_completed_total 360
 - workload_id 362
 - workload_name 363
 - workload_occurrence_id 364
- monitoring
 - entities not tracked by service class 75
 - historical 160
 - overview 149
 - priority aging 178
 - real-time 149
 - service class tiers 178
- monitoring data, overview 166

N

- names
 - monitor elements
 - service_subclass_name 353
 - service_superclass_name 354
 - work_action_set_name 361
 - work_class_name 362
- naming rules
 - restrictions 379
- notices 397

O

- operating systems
 - integrating DB2 workload manager 203
- ordering DB2 books 390
- ownership
 - DB2 workload manager objects 16

P

- partitioned database environments
 - monitor elements
 - coord_partition_num 336
- performance
 - DB2 workload manager
 - examples 269, 271
 - performance modeling 199
- permitting workloads to access database 31
- prefetch priority
 - service classes 72
- privileges
 - roles 381
- problem determination
 - information available 395
 - tutorials 395
- procedures
 - WLM_CANCEL_ACTIVITY 277
 - WLM_CAPTURE_ACTIVITY_IN_PROGRESS 278
 - WLM_COLLECT_STATS 280
 - WLM_SET_CLIENT_INFO 321

Q

- queries
 - monitor elements
 - queue_assignments_total 346
 - queue_size_top 347
 - queue_time_total 347
- queues
 - prefetch 72

R

- ranges
 - monitor elements
 - bottom 329
- recognized activities 13
- REMAP ACTIVITY
 - definition 113, 117
- remapping
 - activities 123
- remapping activities
 - sample scripts 117
- restrictions
 - naming rules 379
- revoking
 - USAGE privilege on workload 34
- roles 381
- routines
 - WLM_CANCEL_ACTIVITY
 - example of cancelling activity 267
- rows
 - monitor elements
 - rows_fetched 348
 - rows_modified 348
 - rows_returned 349
 - rows_returned_top 350

S

- schemas
 - classification of CALL statement 48
- sections
 - monitor elements
 - section_env 352
- security
 - using trusted contexts 383
- service classes
 - agent priority 72
 - altering 78
 - changes occur at statistics reset 190
 - analyzing system slowdown, example 86
 - buffer pool priority 73
 - connection and activity states 74
 - creating 76
 - dropping 80
 - entities not tracked by 75
 - example 81
 - mapping of activities 68
 - point-in-time statistics, obtaining 158
 - prefetch priority 72
 - service subclasses
 - default 66
 - service superclasses
 - default 66
- service subclasses
 - altering 78
 - creating 76

- service subclasses (*continued*)
 - dropping 80
 - monitoring data, overview 166
- service superclasses
 - altering 78
 - creating 76
 - dropping 80
 - monitoring data, overview 166
- SET WORKLOAD command 26, 364
- snapshot monitor
 - using to supplement table functions 193
- SQL statements
 - displaying help 391
- sqleseti API
 - workload assignment 35
- SQLROWSREAD activity threshold
 - description 99
- SQLROWSREADINSC activity threshold
 - description 100
- SQLROWSRETURNED activity threshold
 - description 101
- SQLTEMPSPACE activity threshold 102
- statistics
 - collection for workload management 188
 - DB2 workload manager objects 169
- statistics event monitor 160
- stored procedures
 - WLM_CANCEL_ACTIVITY 168
 - WLM_CAPTURE_ACTIVITY_IN_PROGRESS 168
 - WLM_COLLECT_STATS 168
 - WLM_SET_CLIENT_INFO 168
- SYSDEFAULTMAINTENANCECLASS (default maintenance service superclass)
 - overview 66
- SYSDEFAULTSYSTEMCLASS (default system service superclass)
 - overview 66
- SYSDEFAULTUSERCLASS (default user service superclass)
 - overview 66

T

- table functions
 - aggregating data 159
 - example of using 154
 - monitoring at different levels 155
 - snapshot monitor 193
 - WLM_COLLECT_STATS 190
 - WLM_GET_SERVICE_CLASS_AGENTS_V97 88
 - WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97
 - long-running activity, identifying 267
 - WLM_GET_SERVICE_SUBCLASS_STATS_V97 158
 - analyzing system slowdown, example 86, 266
 - WLM_GET_WORK_ACTION_SET_STATS
 - analysis of activities, example 57
 - obtaining count of activities, example 57
- table spaces
 - SQLTEMPSPACE threshold 102
- terms and conditions
 - use of publications 396
- threshold violations event monitor 160
- thresholds
 - action 88
 - ACTIVITYTOTALTIME 96
 - aggregate 88, 102
 - AGGSQLTEMPSPACE 103
 - altering 110

- thresholds (*continued*)
 - CONCURRENTDBCOORDACTIVITIES 103
 - CONCURRENTWORKLOADACTIVITIES 105
 - CONCURRENTWORKLOADOCCURRENCES 106
 - CONNECTIONIDLETIME 94
 - CPUTIME 96
 - CPUTIMEINSC 97
 - creating 109
 - DB2 workload manager
 - activities 95
 - connection 94
 - evaluation order 92
 - queuing 102
 - DB2 workload manager connection 94
 - definition 102
 - domain 91
 - dropping 110
 - enforcement scope 91
 - ESTIMATEDSQLCOST 98
 - example 145
 - monitor elements
 - num_threshold_violations 344
 - threshold_action 356
 - threshold_domain 356
 - threshold_maxvalue 357
 - threshold_name 357
 - threshold_predicate 357
 - threshold_queuesize 357
 - thresholdid 358
 - monitoring violations 194
 - purpose 88
 - remapping activities 113, 123
 - scope resolution of activity 95
 - SQLROWSREAD 99
 - SQLROWSREADINSC 100
 - SQLROWSRETURNED 101
 - SQLTEMPSPACE 102
 - supported work classifications 51
 - TOTALDBPARTITIONCONNECTIONS 107
 - TOTALSCPARTITIONCONNECTIONS 108
 - usage example 111
 - work actions 133
- time
 - monitor elements
 - prep_time 346
 - time_completed 358
 - time_created 358
 - time_of_violation 359
 - time_started 359
- timestamps
 - monitor elements
 - activate_timestamp 323
 - statistics_timestamp 355
- TOTALDBPARTITIONCONNECTIONS connection
 - threshold 107
- TOTALSCPARTITIONCONNECTIONS connection
 - threshold 108
- troubleshooting
 - online information 395
 - tutorials 395
- trusted connections 383
- trusted contexts 383
- tutorials
 - problem determination 395
 - troubleshooting 395
 - Visual Explain 395

U

- units of work (UOW)
 - assignment to default workload 23
 - matching to workload
 - example 41
 - monitor elements
 - parent_uow_id 345
 - uow_id 360
- unrecognized activities 13
- updates
 - DB2 Information Center 392, 393

V

- Visual Explain
 - tutorial 395

W

- watermark monitor elements
 - act_cpu_time_top 326
 - act_rows_read_top 327
 - concurrent_connection_top 330
 - concurrent_wlo_act_top 330
 - concurrent_wlo_top 330
 - coord_act_lifetime_top 334
 - cost_estimate_top 337
 - uow_total_time_top 360
- watermarks
 - monitor elements
 - concurrent_act_top 329
 - rows_returned_top 350
 - temp_tablespace_top 355
- WLM_CANCEL_ACTIVITY procedure 277
- WLM_CAPTURE_ACTIVITY_IN_PROGRESS procedure 278
- wlm_collect_int database configuration parameter
 - description 365
- WLM_COLLECT_STATS procedure
 - description 280
 - resetting in-memory statistics 190
- WLM_GET_ACTIVITY_DETAILS table function
 - description 281
- WLM_GET_QUEUE_STATS table function 288
- WLM_GET_SERVICE_CLASS_AGENTS_V97 table function
 - description 291
 - investigating agent usage by service class (scenario) 88
- WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97 table function
 - description 298
- WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97 table function
 - examples
 - aggregating data 159
 - identifying long-running activities 267
- WLM_GET_SERVICE_SUBCLASS_STATS_V97 table function
 - analyzing system slowdown (example) 86
 - description 302
 - examples
 - aggregating data 159
 - analyzing system slowdown 266
- WLM_GET_SERVICE_SUPERCLASS_STATS table function 310
- WLM_GET_WORK_ACTION_SET_STATS table function
 - analyzing workload by activity type (example) 57
 - description 311

- WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97 table function
 - description 313
- WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97 table function
 - aggregating data (example) 159
- WLM_GET_WORKLOAD_STATS_V97 table function
 - description 318
- WLM_SET_CLIENT_INFO procedure 321
- WLMADM authority
 - overview 3
- work action sets
 - altering 138
 - creating 137
 - disabling 139
 - domain and permitted work actions 130
 - dropping 140
 - examples
 - association with other objects 126
 - determining types of work being run 147
 - using work action set and database threshold 145
 - overview 128
 - work actions specifying thresholds 133
- work actions
 - altering 143
 - assigning to database activities 133
 - association with other objects (example) 126
 - creating 140
 - disabling 145
 - dropping 145
 - thresholds 133
 - work action sets 130
- work class sets
 - altering 56
 - association with other objects (example) 126
 - creating 55
 - dropping 56
 - managing DML activities (example) 58
 - overview 48
 - work class evaluation order 50
- work classes
 - altering 55
 - assigning activities 51
 - creating 52
 - dropping 55
 - evaluation order 50
 - examples
 - association with other objects 126
 - defined with ALL keyword 59
 - overview 45
 - supported thresholds 51
- workload management
 - monitor elements 323
 - SET WORKLOAD command 364
- workload manager administrator authority
 - overview 3
- workloads
 - altering 29
 - assignment 21
 - assignment examples 35
 - connection assignment to the default administration workload 26
 - creating 28
 - default 23
 - disabling 33
 - dropping 34
 - enabling 32

- workloads (*continued*)
 - evaluation order 21
 - examples
 - analyzing system slowdown 266
 - assignment when multiple workloads exist 41
 - assignment when workload attributes have multiple values 44
 - assignment when workload attributes have single values 39
 - monitor elements
 - wlo_completed_total 360
 - workload_id 362
 - workload_name 363
 - workload_occurrence_id 364
 - monitoring data 166
 - overview 17
 - permitting database access 31
 - position in workload list 21
 - preventing database access 31
 - USAGE privilege
 - granting 33
 - revoking 34
 - work action set comparison 135



Printed in USA

SC27-2464-00



Spine information:

IBM DB2 9.7 for Linux, UNIX, and Windows

Workload Manager Guide and Reference

