


Query Patroller Administration and User's Guide



Query Patroller Administration and User's Guide

Note

Before using this information and the product it supports, read the general information under Appendix B, "Notices," on page 271.

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 1998, 2009.

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

Contents

Part 1. DB2 Query Patroller overview 1

Chapter 1. Scenarios 3

- Scenario: Managing query submitter needs by configuring submitter profiles 3
- Scenario: Handling very large queries 4
- Scenario: Running large emergency queries 4
- Scenario: Managing queries of different sizes using query classes 5
- Scenario: Using historical analysis to improve performance 6

Chapter 2. Query Patroller components 7

Chapter 3. Query Patroller limitations and restrictions 11

Part 2. Upgrading Query Patroller 19

Chapter 4. Upgrading Query Patroller servers 21

- Enabling Query Patroller after upgrade 21

Chapter 5. Upgrading Query Patroller client tools 23

Part 3. DB2 Query Patroller installation tasks overview 25

Chapter 6. Query Patroller installation environment overview 27

Chapter 7. Installing Query Patroller with the DB2 Setup Wizard (Linux and UNIX) 31

- Installing Query Patroller server using the DB2 Setup wizard (Linux and UNIX) 31
- Verifying the installation of Query Patroller server 32
- Installing Query Patroller client tools using the DB2 Setup wizard (Linux and UNIX) 34

Chapter 8. Installing Query Patroller with the DB2 Setup Wizard (Windows) . 37

- Installing Query Patroller server using the DB2 Setup wizard (Windows) 37
- Installing the Query Patroller client tools using the DB2 Setup Wizard (Windows) 38

Chapter 9. Configuring the client and server connection 41

- Configuring a client to Query Patroller server connection using the command line processor 41

Part 4. Configuring Query Patroller 43

Chapter 10. Query Patroller configuration roadmap 45

- Setting up Query Patroller server manually. 46
- Query Patroller administration tasks overview. . . . 47
 - Starting Query Patroller 48
 - Enabling Query Patroller to intercept queries . . . 48
 - Stopping Query Patroller 49
 - Query processing by Query Patroller 49

Chapter 11. Query Patroller system configuration 53

- Setting query thresholds for the Query Patroller system 55
- Updating the list of databases in Query Patroller . . 56
- Enabling e-mail notification of Query Patroller submitters 56
- Enabling collection of historical data 57

Chapter 12. Managing users 59

- Administering operators 59
 - Query Patroller operators. 59
 - Query Patroller operator profiles 59
 - Creating operator profiles for users and groups . 60
 - Suspending or restoring operator privileges for users and groups 61
- Submitter profile configuration 62
 - Query Patroller submitters 65
 - Query Patroller submitter profiles 65
 - Configuring submitter profiles 68
 - Creating submitter profiles for users and groups . 69
 - Setting submitter resource limits 71
 - Suspending or restoring submitter privileges for users and groups 72
- Query Patroller query submission preferences . . . 72
 - Setting individual query submission preferences . 73
 - Setting query submission preferences for another submitter 74

Chapter 13. Query interception and management in Query Patroller 77

- Managed query status 78
- Changing the status of queries using Query Patroller 80
- Query Patroller variables 81
- Viewing managed query details using Query Patroller 81
- Viewing the SQL of managed queries using Query Patroller 83
- Viewing the query submitter profiles 83

Filtering managed queries using Query Patroller . . .	84
Result tables and result sets in Query Patroller . . .	85
Viewing result tables using Query Patroller. . . .	86
Running held queries at a scheduled time	87
Scheduling the start time for running held queries	88
Cost estimation in Query Patroller.	89
Definition of your query management policy . . .	89

Chapter 14. Managing space 93

Setting Query Patroller maintenance schedules for queries and result tables	93
Managing historical queries	94
Scheduling purges of managed queries and result tables	95
Scheduling purges of historical queries	96
Dropping result tables manually using Query Patroller	97
Removing managed queries manually	98
Removing historical queries manually	98
Removing orphaned result table aliases	99

Chapter 15. Query class configuration 101

Configuring query classes	103
Creating query classes for Query Patroller. . . .	104
Removing query classes for Query Patroller . . .	105
Query Patroller query classes	106

Chapter 16. Query Patroller thresholds 109

Part 5. About Query Patroller historical analysis 111

Chapter 17. Getting started with the Query Patroller historical analysis interface 113

Chapter 18. Enabling collection of historical data 117

Chapter 19. Uses for historical analysis reports 119

Chapter 20. Collecting historical data with Query Patroller 121

Chapter 21. Generating historical data using Query Patroller 123

Chapter 22. Determining when historical data was last generated . . 125

Chapter 23. Viewing historical query details using Query Patroller. 127

Chapter 24. Filtering tables for historical analysis using Query Patroller 129

Chapter 25. Filtering queries for historical analysis using Query Patroller 131

Chapter 26. Viewing index details using Query Patroller 133

Part 6. Using Query Patroller with other DB2 components 135

Part 7. Tuning Query Patroller . . . 139

Part 8. Query Patroller commands 141

Chapter 27. Query Patroller command line support 143

Chapter 28. ADD OPERATOR_PROFILE. 145

Chapter 29. ADD QUERY_CLASS . . . 149

Chapter 30. ADD SUBMISSION_PREFERENCES 151

Chapter 31. ADD SUBMITTER_PROFILE 153

Chapter 32. CANCEL QUERY. 157

Chapter 33. GENERATE HISTORICAL_DATA 159

Chapter 34. FILE RESULT 161

Chapter 35. GET OPERATOR_PROFILE. 163

Chapter 36. GET QP_SYSTEM 165

Chapter 37. GET QUERY 167

Chapter 38. GET QUERY_CLASS . . . 169

Chapter 39. GET SUBMISSION_PREFERENCES 171

Chapter 40. GET SUBMITTER_PROFILE	173	Chapter 59. RUN IN BACKGROUND QUERY	215
Chapter 41. LIST OPERATOR_PROFILES	175	Chapter 60. SHOW RESULT	217
Chapter 42. LIST QUERIES	177	Chapter 61. UPDATE OPERATOR_PROFILE.	219
Chapter 43. LIST QUERY_CLASSES	181	Chapter 62. UPDATE QUERY_CLASS	223
Chapter 44. LIST SUBMISSION_PREFERENCES	183	Chapter 63. UPDATE SUBMISSION_PREFERENCES	225
Chapter 45. LIST SUBMITTER_PROFILES	185	Chapter 64. UPDATE SUBMITTER_PROFILE	229
Chapter 46. qpcenter - Start Query Patroller Center	187	Chapter 65. UPDATE QP_SYSTEM	233
Chapter 47. qpsetup - Set up Query Patroller server	189	Query Patroller system threshold settings	234
Chapter 48. qpstart - Start Query Patroller	193	Held query handling settings	235
Chapter 49. qpstop - Stop Query Patroller	195	Query interception settings	236
Chapter 50. REMOVE OPERATOR_PROFILE.	197	System maintenance settings	237
Chapter 51. REMOVE QUERY_CLASS	199	Historical data collection settings	239
Chapter 52. REMOVE QUERY_INFO	201	E-mail notification settings	240
Chapter 53. REMOVE QUERY_INFO_HISTORY	203	<hr/>	
Chapter 54. REMOVE RESULT	205	Part 9. Reference tables.	243
Chapter 55. REMOVE RESULT_TABLE_ALIASES Command .	207	Chapter 66. Control tables	245
Chapter 56. REMOVE SUBMISSION_PREFERENCES	209	<hr/>	
Chapter 57. REMOVE SUBMITTER_PROFILE	211	Part 10. Appendixes	259
Chapter 58. RUN HELD_QUERY	213	Appendix A. Overview of the DB2 technical information	261
		DB2 technical library in hardcopy or PDF format	261
		Ordering printed DB2 books	264
		Displaying SQL state help from the command line processor	265
		Accessing different versions of the DB2 Information Center	265
		Displaying topics in your preferred language in the DB2 Information Center	265
		Updating the DB2 Information Center installed on your computer or intranet server	266
		Manually updating the DB2 Information Center installed on your computer or intranet server	267
		DB2 tutorials	269
		DB2 troubleshooting information	269
		Terms and Conditions	270
		Appendix B. Notices	271
		Index	275

Part 1. DB2 Query Patroller overview

DB2® Query Patroller is a powerful query management system that you can use to proactively and dynamically control the flow of queries against your DB2 database in the following key ways:

- Define separate query classes for queries of different sizes to better share system resources among queries and to prevent smaller queries from getting stuck behind larger ones
- Give queries submitted by certain users high priority so that these queries run sooner
- Automatically put large queries on hold so that they can be canceled or scheduled to run during off-peak hours
- Track and cancel runaway queries

Important: With the new workload management features introduced in DB2 Version 9.5, Query Patroller and its components have been deprecated in Version 9.7 and might be removed in a future release. For more information, see the “DB2 Governor and Query Patroller have been deprecated” topic in the *What's New for DB2 Version 9.7* book.

The features of Query Patroller allow you to regulate your database's query workload so that small queries and high-priority queries can run promptly and your system resources are used efficiently. In addition, information about completed queries can be collected and analyzed to determine trends across queries, heavy users, and frequently used tables and indexes.

Administrators can use Query Patroller to:

- Set resource usage policies at the system level and at the user level
- Actively monitor and manage system usage by canceling or rescheduling queries that could impact database performance
- Generate reports that assist in identifying trends in database usage such as which objects are being accessed, and which individuals or groups of users are the biggest contributors to the workload

Query submitters can use Query Patroller to:

- Monitor the queries they have submitted
- Store query results for future retrieval and reuse, effectively eliminating the need for repetitive query submission
- Set a variety of preferences to customize their query submissions, such as whether to receive e-mail notification when a query completes

Chapter 1. Scenarios

Scenario: Managing query submitter needs by configuring submitter profiles

The Shopmart data warehouse has three main types of users: sales representatives, sales managers, and business analysts. Each of these three groups has different requirements with respect to the types of queries that they submit and the response times that they require.

When the sales representatives submit queries, they are often seeking an immediate answer to a customer inquiry. At times, receiving query results quickly can make the difference between winning or losing an important account. The sales representatives use a query application to submit their queries.

The sales managers use the data warehouse primarily to prepare weekly reports. These queries are written with a customized query-building application and are therefore reasonably consistent in terms of the size of the queries and the efficiency of the SQL. Sales managers submit their queries under a time constraint: they must have up-to-date figures compiled in time to present at a weekly meeting.

The business analysts write complex ad hoc queries. When written by an analyst who has limited experience with SQL, these queries sometimes become very large and unwieldy, which can severely impact the performance of the data warehouse. The response time for these queries is not particularly critical, although the results should return within a matter of hours.

To reflect the different characteristics of these groups and their queries, the administrator, Mel, has created three Query Patroller submitter profiles based on the following existing DB2 user groups: Managers, Salesreps, and Analysts.

Because of the rapid response requirements of the sales representatives, the administrator, Mel, has configured the Salesreps submitter profile so that queries from this group of submitters will not be intercepted by Query Patroller. This avoids the performance cost associated with having Query Patroller intercept and manage these queries.

To ensure that the report queries for sales managers execute relatively quickly, Mel has assigned a queue priority of 999 to the Managers submitter profile. Having a priority of 999 means that when a manager's query is placed in a queue, it is given priority over queries from other submitters. Because the size of the sales manager's queries is consistent, Mel has left the query cost threshold for this submitter profile as the default.

The business analysts may need to submit queries that are larger than the default cost allowed. However, to prevent the business analysts from submitting queries that are too costly, Mel has set the maximum query cost threshold for the Analysts submitter profile to 100 million. If an analyst submits a query whose cost exceeds this threshold, the query will be held by Query Patroller.

These Shopmart submitter profile settings are summarized in the following table.

Table 1. Shopmart submitter profile settings

Submitter profile	Intercepted by Query Patroller?	Maximum query cost (in timerons)	Queue priority
Salesreps	N	n/a	n/a
Managers	Y	Default (10 000 000)	999
Analysts	Y	100 000 000	Default (500)

Scenario: Handling very large queries

Jane is a business analyst examining customer buying patterns in Shopmart stores. She needs to make recommendations regarding the suite of services to be offered at a new Shopmart that is currently under plan.

She suspects that there is an important relationship between the type of fresh produce a customer purchases and the type of in-store services such a customer uses, such as dry cleaning or photofinishing. To explore such a relationship, Jane submits a very large query, involving dozens of tables from different departments.

In order to ensure that an unusually large query from a business analyst does not tie up system resources and prevent smaller queries from running, the system administrator, Mel, has set the Query Patroller thresholds so that extremely large queries from the business analysts will be held. When Query Patroller holds a query that exceeds the size allowed for the business analysts, Query Patroller alerts the submitter that the query is being held. A decision can then be made about whether the query should be run immediately, scheduled for later execution, or canceled. Mel has set up the Shopmart system so that any queries that are still held at the end of the day will be run overnight.

The size of Jane's query exceeds the maximum size specified in the Query Patroller group profile for business analysts. Query Patroller holds the query, and Jane gets an error message from Query Patroller saying that her query is being held. So she calls Quentin from the data warehouse support team.

Quentin examines the query to make sure that the query is not unusually large due to an error in the SQL statement. He decides that the query is legitimately large and therefore does not warrant being canceled. Quentin now has two options for dealing with Jane's held query: he can release the query to run as soon as resources are available, or he can let the query run at the scheduled time for held queries.

Quentin tells Jane that he will allow her query to run even though it exceeds her normal cost threshold. However, to minimize the impact on other users, he explains that he has decided to set the query to run overnight.

Scenario: Running large emergency queries

There has been an emergency recall of a line of Shopmart healthcare products sold between May and October of last year. The legal department has asked Arun, the pharmacy sales manager, for a detailed report on all sales of these products during this period and on any remaining stock.

Although this means executing a very large query at a time of day when the system is very busy, it is urgent that the query run immediately. Arun notifies Mel, the system administrator, that this query cannot be scheduled for a later time but must run right away.

To safeguard against a single large query crippling the entire system, Mel has set the Query Patroller thresholds in query submitter profiles so that extremely large queries from any group of users, including the sales managers, are held.

Knowing how important Arun's query is, Mel decides to temporarily raise the queue priority value in Arun's submitter profile and asks Arun to wait until she has done this before he submits his query. Raising Arun's queue priority means that if the query is queued, it will become the first query to run as soon as some other query completes.

When submitted, the query is intercepted and held by the system because it exceeds the maximum query cost for Arun's submitter profile. Mel releases the query from a held state to allow it to run as soon as sufficient resources are available.

After the query executes, Mel resets the queue priority value in Arun's submitter profile to its original value.

Scenario: Managing queries of different sizes using query classes

Head office requires each product line manager to review the corporation-wide sales figures from the previous week to prepare for the Monday afternoon sales meeting. Since the sales figures are not available until after close of business on Saturday, the managers usually run their weekly sales queries between 8:30 and 11:00 a.m. Monday morning.

Even though the system could potentially run all of the weekly sales queries simultaneously, these queries would monopolize system resources and prevent other users from running small queries, such as customer account queries, at the same time. To address this problem, the system administrator, Mel, has set up query classes to allocate system resources to both smaller queries and larger queries.

Query classes are Query Patroller mechanisms for grouping and running queries according to size. For each query class, you can specify the size of the queries in the class and the number of queries that can run concurrently. The queries in each query class are queued separately, so you can control the flow of queries against a database.

Small queries take only seconds to complete, so in setting up the query class for small queries, Mel has not set a limit on the number of queries from this class that can run simultaneously. To limit the amount of resources that the large weekly sales queries can tie up, she has set the maximum number of queries for the query class for large queries to 10. This means that if 15 weekly sales queries are submitted at once, 5 of these will be placed in a queue until some of the running queries in this query class complete.

When Bill, the dairy sales manager, submits his weekly sales query on Monday at 9:55 a.m., there are already 12 other managers trying to run similar queries. Meanwhile, Alphonso, a credit representative, is submitting a query to verify some customer account information.

Because of the query classes that Mel has set up, Alphonso's small query runs immediately. Bill's query will be placed in the queue and will run in the large query class once some of the other sales queries complete.

Scenario: Using historical analysis to improve performance

The Shopmart data warehouse has grown very large in the past few years, and the database administrator, Mel, is looking for ways to restructure the database and improve performance.

Query Patroller has been collecting query activity data on the data warehouse for the past six months. The historical analysis function uses this data to generate reports on which database objects have been used, by which submitters, and when.

Mel uses the historical analysis Tables Not Hit report to identify tables that have not been accessed by any queries in the past six months. Several of these tables are very large and are located on the most powerful processors in the data warehouse. Because the historical data suggests that these tables are infrequently or never accessed, Mel decides to move them to a slower access device, freeing up space for more heavily-used tables.

Mel also uses another report, Tables Hit, to look at the most frequently accessed tables in the database. Once he has identified the tables that appear to be the most critical to the data warehouse users, he drills down to look more closely at the columns that are accessed to find suitable candidates for indexes. He also looks at the Indexes not hit report to see which indexes are not being used and should be removed.

After Mel has implemented these changes to the Shopmart data warehouse, he can monitor the effect of these changes on query performance. To do this, he examines the historical analysis graphs and reports on query execution time and compares the current average execution time of queries with the average execution time prior to the database changes.

Chapter 2. Query Patroller components

DB2 Query Patroller is a client and server solution consisting of the following components:

- Query Patroller server
- Query Patroller Center (available only for Windows® on x86_32, Linux® on x86_32, and Linux on x64)
- Query Patroller command line support

DB2 Query Patroller can be deployed on a system running DB2 Enterprise Server Edition.

Important: With the new workload management features introduced in DB2 Version 9.5, Query Patroller and its components have been deprecated in Version 9.7 and might be removed in a future release. For more information, see the “DB2 Governor and Query Patroller have been deprecated” topic in the *What’s New for DB2 Version 9.7* book.

Query Patroller server

When you install Query Patroller server, the following software elements are deployed to the target computer:

Query Patroller stored procedures

Query Patroller stored procedures are called by other Query Patroller components to perform the necessary database tasks.

Control tables

When Query Patroller is set up to manage queries issued against a database, the DB2QP schema, control tables, triggers, functions, and procedures are created within that database. The control tables store all of the information that Query Patroller requires to manage queries. This information includes the following:

- Query Patroller system properties settings
- Query class information
- Submitter information, including query submission preferences
- Operator information
- Managed query properties information
- Historical query properties information
- Query result information
- Historical analysis data
- Scheduled purge job details

For example, the SUBMITTER_PROFILE table contains information such as the submitter’s ID, authority level, and the maximum number of queries that the user can have running simultaneously. When the user submits a query, Query Patroller references the SUBMITTER_PROFILE table for these parameters.

Log files

Diagnostic information about errors is recorded in these Query Patroller log files:

qpsetup.log

Query Patroller writes to qpsetup.log during installation and when the qpsetup command is issued. On UNIX® operating systems, qpsetup.log resides in the INSTANCE/db2dump directory, where INSTANCE is the directory where you installed DB2. On Windows, qpsetup.log resides in the directory specified in the *diagpath* database configuration parameter. The qpsetup.log file is intended for use by Query Patroller administrators.

qpuser.log

Query Patroller starts writing to the qpuser.log file when the system becomes active. The information written to the qpuser.log file is used for problem determination and is intended for use by Query Patroller administrators. On UNIX operating systems, qpuser.log resides in the INSTANCE/db2dump directory, where INSTANCE is the directory where you installed DB2. On Windows, qpuser.log resides in the directory specified in the *diagpath* database configuration parameter. If you choose to locate qpuser.log in a place other than the default log path, ensure that the permissions in the new directory allow write access for the fenced user ID. Query Patroller commands call a fenced stored procedure which must have write access to this file for diagnostic information to be logged.

qpdiag.log

Query Patroller starts writing to the qpdiag.log file when the system becomes active. The information written to the qpdiag.log file is used for problem determination and is intended for use by DB2 technical support. On UNIX operating systems, qpdiag.log resides in the INSTANCE/db2dump directory, where INSTANCE is the directory where you installed DB2. On Windows, qpdiag.log resides in the directory specified in the *diagpath* database configuration parameter. If you choose to locate qpdiag.log in a place other than the default log path, ensure that the permissions in the new directory allow write access for the fenced user ID. Query Patroller commands call a fenced stored procedure which must have write access to this file for diagnostic information to be logged.

Use a text editor to view the log files on the server where Query Patroller is installed. The most recent events are recorded at the end of the file.

Generally, each entry contains the following parts:

- A time stamp
- Instance, database, and database partition name details
- Process ID (PID) and name, or Thread ID (TID) and name
- The component reporting the error
- A diagnostic message (usually beginning with "DQP") explaining the error

The log files grow continuously. Occasionally, they should be backed up then erased. A new log file is generated automatically the next time it is required by Query Patroller.

Query Patroller directory access

All userids that require access to Query Patroller must have read access to the directory <instance directory>\ctrl\qp on Windows and INSTHOME/sql1lib/ctrl/qp on Linux and UNIX. If any userids that are running the Query Patroller tools do not have access to this directory, they might be unable to connect to the Query Patroller server, resulting in error SQL29007. A file located in this directory contains configuration information that allows DB2 and the Query Patroller tools to communicate with Query Patroller server. You must not modify or delete the files in this directory while the Query Patroller server is running.

Query Patroller Center

The Query Patroller Center is a graphical user interface that allows administrators to manage Query Patroller system properties, users, and queries, and to view historical analysis reports. The Query Patroller Center also allows query submitters to manage their queries, save query results, and customize their query submission preferences.

The look and functionality of the Query Patroller Center varies depending on different factors, such as the authority of the user and whether the DB2 administration tools are also installed.

An administrator has access to the Query Patroller Center's full functionality. The following list shows some of the tasks that administrators can do with Query Patroller Center:

- Manage the Query Patroller system parameters
- Create, update, or delete profiles for Query Patroller submitters and operators
- Create, update, or delete submission preferences for Query Patroller submitters
- Create, update, or delete query classes
- Monitor and manage queries that have been intercepted by the Query Patroller system
- Generate and analyze reports that display database usage history

A submitter has access to a subset of the Query Patroller Center's functionality. The following list shows some of the tasks that submitters can do with Query Patroller Center:

- Monitor and manage queries that they have submitted through the Query Patroller system
- Store results of the queries that they have submitted for future retrieval
- Show or file results of the queries that they have submitted
- Create, update, or delete their own query submission preferences

Query Patroller command line support

Command line support enables Query Patroller administrators and submitters to perform most Query Patroller tasks from the DB2 CLP or from the operating system's command line prompt. Query Patroller commands can also be combined with shell scripts or languages such as Perl, awk, and REXX™.

Chapter 3. Query Patroller limitations and restrictions

Incompatible SQL statements and Query Patroller functions

There are several types of SQL queries which are incompatible with particular Query Patroller functions. These query types are listed in the following table.

Table 2. Query Patroller limitations by SQL statement type

Statement type	Intercept	Manage	Historical analysis	Schedule	Hold	Queue
Non-SELECT statements (UPDATE, INSERT, DELETE)	yes	yes	yes ¹	no	no	yes
Static SQL containing host variables	yes	yes	yes ¹	no	no	yes
SQL queries from stored procedures	yes	yes	yes ¹	no	no	no
User-defined functions containing non-zero nested queries	yes	yes	yes ¹	no	no	no ²
SQL queries containing parameter markers or special registers	yes	yes	yes ¹	no	no	yes
SQL queries containing declared temporary tables or created temporary tables	yes	yes	no	no	no	yes
SQL queries containing an identity value function (IDENTITY_VAL_LOCAL()) or a sequence value function (seqno())	yes	yes	yes ¹	no	no	yes
SQL queries containing an encrypt or decrypt function	yes	yes	yes ¹	no	no	yes
Xquery queries, and SQL queries that return xml data type	yes	yes	no	no	no	yes
SQL queries containing global variables	yes	yes	no	no	no	yes
SQL queries that return DECFLOAT data types	yes	yes	no	no	no	yes

Note:

1. Historical analysis is only performed on queries with a completion status of "D" (done). If a query of this type exceeds a threshold, then it will be assigned a completion status of "R" (rejected) rendering the query ineligible for historical analysis.
 2. Nested queries cannot be queued. Instead, a nested query will run immediately if it exceeds a threshold that would normally cause it to be queued.
-

Quiesce mode

When DB2 is in quiesce mode, Query Patroller will be bypassed by all queries.

Potential inconsistency between db2qp.result_info table with the database

A situation may arise in which there is an inconsistency between the db2qp.result_info table and the result tables in the database. This inconsistency may occur if the table space containing the control tables undergoes a roll forward recovery to a point in time prior to the last update of the result tables which reside on a different table space. For example, if you perform a purge of your result tables, and subsequently restore and perform a roll-forward of the control table space to a point in time prior to the purge, the db2qp.result_info table will contain information that says the result tables still exist when they do not.

UDF provided to select SQL statement text from db2dqp.track_query_info table

Since the data type for the statement field in the db2dqp.track_query_info table is BLOB, you will not be able to perform a simple SELECT statement to retrieve the values in the table. In order to query this field, you need to use the db2qp.convertToString function to convert the values for retrieval. For example, to select all the values from the statement field in this table, you will enter the following:

```
select db2qp.convertToString(statement) from db2qp.track_query_info
```

Timing considerations for held, queued, and analyzed queries

If the execution of a query is delayed for any reason, such as when the query is held or queued, and if any tables accessed by the query have been changed or removed between the time that the query was submitted and when the query executes, there may be unexpected results or an error may occur.

If a query is being analyzed and the tables involved in the query have been changed or removed between the time the query was run and when the query is being analyzed, there may be unexpected results and an error may occur.

Command line support for non-English languages

In order for non-English characters to appear properly during command line usage, Query Patroller commands must be entered in a DB2 CLP window for languages other than English.

Restriction on the generation of historical data with positioned UPDATE or DELETE statements

Historical data cannot be generated for queries that include positioned UPDATE or positioned DELETE statements.

Query recovery

On rare occasions, when the status of a queued or running query is changed, Query Patroller may be unable to record the new status immediately. This usually happens in an abnormal termination situation such as one of the following:

- The DB2 server terminates unexpectedly or is shut down with the db2stop force command while there are queries that are queued, running, or both.

- The Query Patroller server terminates unexpectedly or is shut down with the `qptstop dbname force` command while there are queries that are queued, running, or both.
- Both the DB2 server and the Query Patroller server terminate unexpectedly while there are queries that are queued, running, or both. This might occur in the event of a power outage or hardware failure.

Query Patroller server performs an automatic recovery at startup time and at periodic intervals. It checks for queries with a current status of queued or running and checks to see if the status is still accurate. If the current status is accurate, the query is handled normally and, if Query Patroller server was terminated and restarted, internal Query Patroller server data structure is restored. If, however, a query with a status of queued or running is found to no longer exist in DB2 because the DB2 server was terminated and restarted, or the Query Patroller was down and could not update the status of the query, then recovery will be performed on the query. The recovery action taken depends on whether the query was to return the results to a client application or to a DB2 result table:

- A running query that was to return results to a client application will have its status changed to unknown.
- A queued query that was to return results to a client application will have its status changed to aborted.
- A running query that was to return results to a DB2 result table will be automatically re-run.
- A queued query that was to return results to a DB2 result table will be automatically re-queued.

DBCLOB objects not available in Show result dialog

Due to a JDBC limitation, DBCLOB objects can not be displayed in the Query Patroller **Show result** dialog window. Instead, an empty string will appear in place of DBCLOB objects in the dialog window. This limitation applies only to the Query Patroller Center, and not to the Query Patroller command line.

Apply filters when viewing a large number of queries

The response time of Query Patroller Center might slow down considerably if you are viewing several hundred managed or historical queries. To alleviate this problem, it is recommended that you apply a filter to the views to reduce the number of queries displayed. For information on how to apply filters in the Query Patroller Center, see the *Query Patroller Guide: Installation, Administration, and Usage* or the Query Patroller information in the DB2 Information Center.

Queries run in the background issue loads

When a query is run in the background, the results of the query are stored in a result table. Any query that will generate a result table is run by a process called *qprunquery*. This process creates a result table and issues a load from cursor to fill the table with the results of the query. This means that queries that produce result tables are subject to all of the same restrictions as any other load from cursor. For a complete description of these restrictions, see the documentation of the LOAD command in the *DB2 Command Reference*.

During each load performed by *qprunquery*, entries are placed in the db2diag log files. On UNIX operating systems, one or more messages will be created in a subdirectory under the *INSTANCE/db2dump* directory, where *INSTANCE* is the

directory where you installed DB2. On Windows, one or more messages will be created in a subdirectory under the directory specified in the `diagpath` database configuration parameter. The name of the message file subdirectory is generated based on the details of the load operation. For example, the following is the name of a generated message file subdirectory:

```
qpTbLoad_SAMPLE_349_2003-05-21-16.51.32
```

where:

- `qpTbLoad` specifies that this is a message file from the `LOAD` command run by `qprunquery`
- `SAMPLE` is the name of the database
- `349` is the query id for which the load was initiated
- `2003-05-21-16.51.32` is the timestamp just before the load is initiated

The message file name contained in this subdirectory would be as follows.

```
qpTbLoad_SAMPLE_349_2003-05-21-16.51.32.MSG.*
```

Note: For multipartioned databases, the `LOAD` command will create more than one message file and append different file extensions to each file name.

The message files are deleted once the load completes successfully. To aid in problem determination, the message files are not deleted if the load fails.

There is a limit to the number of simultaneous loads that can be run in parallel. Exceeding this limit results in an aborted query and an error `SQL6555` recorded in the `qpdiag.log` file. If this error is encountered, you can remedy the situation by changing the range specified by the `DB2ATLD_PORTS` registry variable, which determines the number of parallel loads permitted at once. To calculate the approximate number of ports required in your system, decide on the maximum number of loads that need to be run at one time, including those issued by `qprunquery` and other load operations. Multiply this number by the number of logical partitions per physical partition in your environment. Add 25% to this amount.

To set the `DB2ATLD_PORTS` registry variable, issue the following command:

```
db2set DB2ATLD_PORTS=num1:num2
```

where `num1 < num2`

Query Patroller uses a default of 6000 ports in the range 50000–56000; setting `DB2ATLD_PORTS` will override this value.

Out of memory errors with Query Controller or Query Patroller Center

When Query Patroller is managing a large number of queries, and the Query Controller or the Query Patroller Center is running, you may receive an out of memory error even when there is sufficient memory available on the machine. In order to be able to use more of the available memory, you can increase the java heap environment variable settings from their default levels.

The environment variables to update are `QP_INIT_JAVA_HEAP_SIZE` and `QP_MAX_JAVA_HEAP_SIZE`. If these variables are not set, the default is 32 mb and 512 mb, respectively. The Query Patroller server or Query Patroller Center must be restarted in order for changes to `QP_INIT_JAVA_HEAP_SIZE` and

QP_MAX_JAVA_HEAP_SIZE to take effect.

Query Patroller limitations when DYN_QUERY_MGMT is disabled

Query Patroller cannot perform the following actions if the database configuration parameter DYN_QUERY_MGMT is disabled:

- Release queries from a held state
- Make a running or queued query run in the background when the query is in the foreground

If you attempt to release a query from held state, or change a foreground query to a background query when DYN_QUERY_MGMT is set to DISABLE, an error message will be displayed and the state of the query will not change. If held queries are scheduled to run and DYN_QUERY_MGMT is disabled at the time they start running, an error message is written to the qpdiag.log file and the queries are left in held state.

Resolution limitation when using the Terminal Services Client

When using the Terminal Services Client at resolution 640x480 to connect to a remote desktop that is running the Query Patroller Center, the Submission Preferences window might appear blank. For the Submission Preferences window to display properly, you must use a resolution higher than 640x480.

Query Patroller schedule limitations

When working with schedules in the Query Patroller Center, you can use the Schedule window to save schedules to a file and import them later. If you have a schedule that you saved using Version 8.1 FixPak 6 or earlier, you cannot import the schedule using Version 8.2 or later. This limitation is due to the change in serialization between Java™ Software Developer's Kit levels introduced with DB2 UDB Version 8.2.

Abnormal shutdown of the historical data generator

If you run the historical data generator and shut it down in an abnormal way, you will receive an error the next time you attempt to run the historical data generator. Examples of abnormal shutdown include:

- DB2 database stops unexpectedly
- Issuing a db2stop force command
- Issuing a killdb2 command

When the historical data generator shuts down abnormally, you must issue the following command before attempting to rerun the historical data generator:

```
qp -d database generate historical_data stop
```

where *database* identifies the database that the command is being run against.

Binding Query Patroller packages

If the Query Patroller packages are not bound after applying a fixpak, a user without DBADM authority or proper Query Patroller privileges can encounter the following error when using the Query Patroller Center or Query Patroller command line:

```
SQL0001N - Binding or precompilation did not complete successfully.
```

If you are using the Query Patroller Center, the SQL0001N error is logged in qpdiaq.log file. If you are using the Query Patroller command line, the SQL0001N is returned to the console

Auto-bind code exists to initiate automatic binding. However, the automatic binding fails when the connecting user does not have the necessary privileges to execute all statements in the Query Patroller packages. A symptom of this problem is missing folders in the Query Patroller Center.

To avoid this problem, the qpserver.lst packages should be bound manually by a user with DBADM authority or necessary privileges after applying a fixpak.

Binding or rebinding using the REOPT ALWAYS option

If you bind or rebind queries and specify the REOPT ALWAYS option, Query Patroller will not queue those queries.

Scheduling feature cannot save schedules to a file

The Query Patroller scheduling feature cannot save schedules to a file. The problem is due to an incompatibility with the IBM® Software Developer's Kit for Java(TM) level used in this version of the DB2 product

Group support for query submission

If you are logged in, but do not have either DBADM authority or Edit privilege for Query Patroller user administration, you can only add or update a submission preference for yourself. In this case, the **Submitter Profile to Use** drop-down list contains existing submitter profiles of the DB2 groups that you belong to, instead of just the operating system groups that you belong to.

If you are logged in, and have either DBADM authority or Edit privilege for Query Patroller user administration, you can add or update submission preferences for other users. In this case, the **Submitter Profile to Use** drop-down list contains all existing group submitter profiles.

Definition of new database partition group in Query Patroller

When you define a new table space during Query Patroller installation in a partitioned environment, you can select an existing database partition group. Alternatively, you can set up a new database partition group for the new table space. If you choose to set up a new database partition group, you can select only one database partition from the list on the install panel for that new database partition group. If you want to add partitions into the new database partition group, you must add them manually after the installation is complete.

Unavailability of ports might cause an SQL error (Windows)

Submitted queries in Query Patroller might receive SQL code -29007 when there are no available ports on Windows XP or Windows 2003. The likelihood of this error increases as the number of clients accessing Query Patroller increases.

The solution to this problem is to set the following Windows registry variables to the following values:

```
MaxUserPort=65534  
TcpTimedWaitDelay=30
```

and restart your system for the changes to take effect. You can find details about setting Windows registry variables on the Microsoft® Help and Support Web site at <http://support.microsoft.com/>

Scheduling a statement that contains global variables

Query Patroller cannot schedule a statement that contains global variables.

DB2 roles profile

Query Patroller does not support roles profile created in DB2. Submitter profiles can be created for user and group profiles only.

DB2 non-root install feature

Query Patroller does not support DB2 non-root install feature.

Part 2. Upgrading Query Patroller

Chapter 4. Upgrading Query Patroller servers

Upgrading a Query Patroller server involves upgrading the DB2 server to Version 9.7, installing a Query Patroller Version 9.7, and enabling Query Patroller. No additional Query Patroller-specific command is required to upgrade a Query Patroller server.

You must have installed Query Patroller Version 9.5, Version 9.1, or Version 8.

To upgrade Query Patroller server to Query Patroller server Version 9.7:

1. Upgrade your DB2 server where Query Patroller is installed using any of the following tasks: In the upgrade task, you must install Query Patroller server Version 9.7 after installing DB2 Version 9.7 to successfully upgrade your instances.
2. Enable Query Patroller.
3. “Verifying the installation of Query Patroller server” on page 32

If you have a system running Query Patroller client tools, you should upgrade it to Version 9.7. Refer to Chapter 5, “Upgrading Query Patroller client tools,” on page 23 for details.

Many of the successful management concepts available within Query Patroller have been used to help design a number of new Workload Management features introduced since of DB2 9.5. To learn more about these features (workloads, service classes, and thresholds), please refer to discussions on workload management in the DB2 Information Center Version 9.7.

Enabling Query Patroller after upgrade

After upgrading your DB2 server and installing Query Patroller server Version 9.7, you need to enable Query Patroller to capture query information and intercept queries in those databases.

Upgrading your DB2 server requires that you upgrade your instances and all databases within the instance. When you upgrade the Query Patroller databases, the **dyn_query_mgmt** database configuration parameter is set to DISABLE. After upgrading the Query Patroller databases and installing Query Patroller Version 9.7, you need to set the **dyn_query_mgmt** database configuration parameter to ENABLE.

To enable Query Patroller to capture query information and intercept queries using the Command Line Processor interface:

1. Connect to the Query Patroller database, as follows:
`CONNECT TO database-name`
2. Use the UPDATE DATABASE CONFIGURATION command to set the **dyn_query_mgmt** database configuration parameter to ENABLE, as follows:
`UPDATE DB CFG FOR database-name USING DYN_QUERY_MGMT ENABLE`
3. Verify that the **dyn_query_mgmt** database configuration parameter is updated and that the new value has become effective, as follows:
`GET DB CFG FOR database-name SHOW DETAIL`

Chapter 5. Upgrading Query Patroller client tools

Upgrading a Query Patroller client tools involves upgrading the DB2 product to Version 9.7, installing a Query Patroller client tools Version 9.7, and enabling Query Patroller. No additional Query Patroller-specific command is required to upgrade a Query Patroller server.

You must have installed Query Patroller client tools Version 9.5, Version 9.1, or Version 8.

To upgrade Query Patroller client tools to Query Patroller client tools Version 9.7:

1. Upgrade your DB2 product on the system where Query Patroller client tools is installed .
2. Install Query Patroller client tools Version 9.7 on your system using any of the following tasks:
 - “Installing Query Patroller client tools using the DB2 Setup wizard (Linux and UNIX)” on page 34
 - “Installing the Query Patroller client tools using the DB2 Setup Wizard (Windows)” on page 38
3. Verify the installation of Query Patroller client tools is successful.

Part 3. DB2 Query Patroller installation tasks overview

The following chapters describe the steps to take for installing Query Patroller, Version 9 and setting up your Query Patroller server.

Important: With the new workload management features introduced in DB2 Version 9.5, Query Patroller and its components have been deprecated in Version 9.7 and might be removed in a future release. For more information, see the “DB2 Governor and Query Patroller have been deprecated” topic in the *What’s New for DB2 Version 9.7* book.

First steps




Figure 1. Query Patroller installation tasks overview

Chapter 6. Query Patroller installation environment overview

This topic provides an overview of two typical Query Patroller installation environments. After reading this topic you should understand what components make up Query Patroller, and understand the nonpartitioned and partitioned environments in which they can operate.

Important: With the new workload management features introduced in DB2 Version 9.5, Query Patroller and its components have been deprecated in Version 9.7 and might be removed in a future release. For more information, see the “DB2 Governor and Query Patroller have been deprecated” topic in the *What's New for DB2 Version 9.7* book.

A typical Query Patroller installation environment consists of the following computers:

- Query Patroller server installed on all database partitions in your DB2 Enterprise Server Edition environment.
- One or more system administration workstations with Query Patroller client tools installed on them.
- One or more DB2 clients with Query Patroller client tools installed on them.
- One or more DB2 clients without the Query Patroller client tools installed on them.
- One or more Business Intelligence clients without the Query Patroller client tools installed on them.

Query Patroller server and client tools

Query Patroller server

The Query Patroller server is installed on all of the computers that make up your DB2 database environment. It is installed in either a nonpartitioned or a partitioned database environment. The Query Patroller server accepts, analyzes, prioritizes, and schedules queries that run against your database. Query Patroller server can also notify users when their queries are completed.

Query Patroller client tools (system administration workstation)

You can install the Query Patroller client tools on the DB2 client that will function as the system administration workstation in your environment. One or more of your DB2 clients can function as a system administration workstation. You should install these tools on any machine that will perform remote administration of the data warehouse.

The Query Patroller client tools that get installed on your system administration client are the Query Patroller Center and the Query Patroller command line support. These tools enable you to configure and manage your Query Patroller server, to create and delete user profiles, to manage queries and result destinations, and to monitor the usage history of a database.

Query Patroller client tools (DB2 clients)

You can install the Query Patroller client tools on your DB2 clients. You can install either the Query Patroller Center or the Query Patroller command line support (or both). In order to install the Query Patroller Center on your clients they must all have either a DB2 Version 9 client product installed on them, any DB2® Connect™ Version 9 product, or any DB2 Version 9 server product. You may also have your DB2 clients that do not have the Query Patroller Center or the Query Patroller command line support installed on them, submit queries to your Query Patroller server.

Note: You may also have a number of Business Intelligence clients with third-party GUI tools (such as Business Objects) that can submit queries. These clients first submit their queries to a dedicated Business Intelligence server which, in turn, directs the queries to the Query Patroller server for processing.

Query Patroller typical environments

Nonpartitioned DB2 database environment

You can install Query Patroller in a nonpartitioned database environment. For example, a typical nonpartitioned environment consists of the following computers:

- One Query Patroller server.
- One system administration workstation with Query Patroller client tools installed on it.
- Two DB2 Run-Time clients (UNIX and Windows).

Figure 1 illustrates how these computers are typically set up in a nonpartitioned database environment.




Figure 2. Query Patroller nonpartitioned database installation environment

In a typical nonpartitioned database environment, the Query Patroller server is installed on the computer that is running a DB2 server product.

Partitioned DB2 database environment

You can install Query Patroller in a partitioned database environment. For example, a typical partitioned database environment consists of the following computers:

- Query Patroller server installed on all database partitions.
- One Business Intelligence server for receiving and directing queries from Business Intelligence clients.
- One system administration workstation with Query Patroller client tools installed on it.
- Two DB2 clients with Query Patroller client tools installed on them.
- Two DB2 clients without Query Patroller client tools installed on them.

Figure 2 illustrates how these computers are typically set up in a partitioned database environment.




Figure 3. Query Patroller partitioned database installation environment

DB2 Clients which have Query Patroller Client Tools installed must connect to a database partition that has the Query Patroller server installed. In a typical partitioned database environment, the Query Patroller server is installed on all database partitions so that the user can choose any database partition to be the coordinator partition and still be able to use the Query Patroller client tools.

Chapter 7. Installing Query Patroller with the DB2 Setup Wizard (Linux and UNIX)

If you know which Query Patroller components and tools you want to install, and the kind of environment you will be installing on, you can follow the steps below to install Query Patroller on your Linux or UNIX server and clients.

Important: With the new workload management features introduced in DB2 Version 9.5, Query Patroller and its components have been deprecated in Version 9.7 and might be removed in a future release. For more information, see the “DB2 Governor and Query Patroller have been deprecated” topic in the *What's New for DB2 Version 9.7* book.

To install Query Patroller:

1. Open the Query Patroller System Properties notebook:
2. Ensure that the *dyn_query_mgmt* database configuration parameter is set to DISABLE. This parameter must be set to ENABLE for Query Patroller to intercept and manage queries, but it must be set to DISABLE prior to installation to prevent the interception of any internal queries run by the installer. You must connect to the database to update *dyn_query_mgmt* .
If you upgraded your DB2 server to DB2 Version 9.7, the *dyn_query_mgmt* configuration parameter would have automatically been disabled during database upgrade. The owner of the instance containing the database that you want to run Query Patroller on must also have SETSESSIONUSER privilege on PUBLIC.
3. Install Query Patroller server with the DB2 Setup wizard (Linux and UNIX).
4. Verify the installation of Query Patroller server.
5. Install the Query Patroller client tools with the DB2 Setup wizard (Linux and UNIX).

After Query Patroller is installed, you must set the *dyn_query_mgmt* configuration parameter to ENABLE so that Query Patroller can intercept and manage queries.

Installing Query Patroller server using the DB2 Setup wizard (Linux and UNIX)

This task outlines the steps for installing Query Patroller server on a Linux or UNIX operating system using the DB2 Setup wizard.

- To install Query Patroller Version 9 server, a DB2 Version 9 server must already be installed. This list of qualified servers include:
 - DB2 Enterprise Server Edition
 - DB2 Workgroup Server Edition
 - DB2 Express™ Edition
- The instance running on your DB2 server must be stopped. To stop an instance, log on to the system as a user with SYSADM authority and issue the `db2stop` command at a Linux or UNIX shell prompt.
- All of the languages that the clients will use when submitting queries to Query Patroller server must be installed on the server in order to eliminate any server to client dependencies.

This task is part of the larger task of Installing Query Patroller (Linux and UNIX).

To install Query Patroller server (Linux and UNIX):

1. Log on to the system as a user with root authority.
2. Stop the DB2 administration server (DAS) and any DB2 instances related to the DB2 installation path where Query Patroller will be installed so that the DB2 Setup wizard can update files as required.
3. Mount the Query Patroller product CD-ROM.
4. Change to the directory where the CD-ROM is mounted by entering the following command:

```
cd /cdrom
```

Where */cdrom* represents the mount point of the CD-ROM:

5. Enter the `./db2setup` command to start the DB2 Setup wizard. The IBM DB2 Setup Launchpad window opens.
6. Choose **Install a Product** from the DB2 Launchpad.
7. The DB2 Setup Wizard opens. Click on **Work with existing** and choose the DB2 installation path where you want to install Query Patroller.
8. Follow the instructions on the DB2 Setup Wizard.

Note: If you select **Computer usage based** from the **Select the installation type** window using the DB2 Setup wizard and you select **Server** from the **Select how this computer will be used** window, the Query Patroller Center is not installed. If you want to install the Query Patroller Center, select either the **Complete** or the **Custom installation** types from the **Select the installation type** window. You do not have to use the **Computer usage based** installation type to install the Query Patroller server. You can use either the **Complete** or the **Custom** installation types to install the Query Patroller server.

Online help is available to guide you through the installation. To invoke the online help, click the **Help** button on the bottom right hand corner of each installation window, or press **F1**. You can click the **Cancel** button at any time to end the installation.

When you have completed your installation, Query Patroller server will be installed in the directory you specified in the DB2 Setup Wizard.

You have completed the task of installing Query Patroller server on a Linux or UNIX operating system using the DB2 Setup wizard. You are now ready to verify the installation of Query Patroller server. On Windows and Linux clients, you can also install the Query Patroller client tools using the DB2 Setup wizard.

Verifying the installation of Query Patroller server

After you have installed Query Patroller server, you can verify your installation by submitting a query and checking its status.

- You must set up Query Patroller server by issuing the **qpsetup** command. This must be issued prior to verifying the installation of Query Patroller server. You have the choice of either having the **qpsetup** command be issued automatically during the installation, or issuing it manually after the installation at a Linux or UNIX shell prompt or a Windows command prompt.
- Make sure the DB2 instance that the database belongs to is running on the server.

To verify your Query Patroller server installation

1. Log on to the database that is managed by Query Patroller as a user with DBADM authority.
2. Start Query Patroller by issuing the `qpstart` command:

```
qpstart dbname
```

where *dbname* is the name of the database that is managed by Query Patroller. You should receive output that is similar to the following:

```
IBM DB2 Query Patroller V9, (c) Copyright IBM Corp. 1998-2006.  
All rights reserved.  
Initializing.  
Query Patroller started.
```

Alternatively, on Linux operating systems you can start Query Patroller from **Main Menu** by selecting **IBM DB2** → **Query Patroller**.

Note: On a Linux or UNIX operating system the `qpstart` command will release control back to the UNIX shell prompt. You do not need to open up a second shell prompt to continue with the rest of this procedure. On Windows the control is released back to the Windows command prompt immediately.

3. Enable the `dyn_query_mgmt` database configuration parameter by issuing the following command:

```
db2 update db cfg for dbname using DYN_QUERY_MGMT enable
```

where *dbname* is the name of the database that is managed by Query Patroller.

4. Submit a query that will be intercepted by Query Patroller server. This query must have an estimated cost of 15000 timerons or more. The estimated cost must be 15000 timerons or more because queries with costs lower than this are not managed by Query Patroller. This is determined by the default Query Patroller parameters. For example, submit the following query::

```
db2 select count(*) from syscat.tables,syscat.tables,syscat.tables
```

5. Wait for the query to return. If you find that the cost of this query is not high enough (it is less than 15000 timerons), then append enough tables to the command until the cost is greater than the minimal cost of 15000 timerons. For example:

```
db2 select count(*) from syscat.tables,syscat.tables,syscat.tables,...
```

The above query is just an example. You can choose to submit a query that is more appropriate to your particular system. However, ensure that the query has a cost that is greater than 15000 timerons.

6. Check the status of the query that you submitted by entering the following command:

```
qp -d dbname list queries
```

where *dbname* is the name of the database that is managed by Query Patroller. The above query is just an example. You can choose to submit a query that is more appropriate to your particular system. However, ensure that the query has a cost that is greater than 15000 timerons.

7. If the query has completed successfully, you should receive output that is similar to the following:

ID	Status	Created	Completed
1	Done	2003-10-30 18:36:37.615000	2003-10-30 18:36:37.615000

8. If the query was aborted, you should receive output that is similar to the following:

```
   ID      Status      Created                      Completed
-----
   1      Aborted    2003-10-30 18:36:37.615000 2003-10-30 18:36:37.615000
```

9. You can find out why the query was aborted by issuing the following:

```
qp -d dbname get query 1
```

Under "Message Description" you will see the error that caused the query to be aborted.

10. Stop the Query Patroller service by entering the `qpstop` command. You should receive output that is similar to the following:

```
IBM DB2 Query Patroller V9, (c) Copyright IBM Corp. 1998-2006.
All rights reserved.
Stopping Query Patroller.
Query Patroller stopped.
```

11. Investigate the error and resubmit the query. Investigate the error and resubmit the query.

Installing Query Patroller client tools using the DB2 Setup wizard (Linux and UNIX)

This task outlines the steps for installing the Query Patroller client tools on a Linux or UNIX DB2 client using the DB2 Setup wizard.

Before you begin

A DB2 Version 9 product must be installed computer that you will install the DB2 Query Patroller client tools on. The following products are appropriate prerequisites:

- Any DB2 Version 9 client product
- Any DB2 Version 9 Connect product
- Any DB2 Version 9 server product

About this task

This task is part of the larger task of Chapter 7, "Installing Query Patroller with the DB2 Setup Wizard (Linux and UNIX)," on page 31. Note that you do not need to install Query Patroller Client Tools on a DB2 client if you only want to submit queries from that client to the Query Patroller server. You only need to install Query Patroller Client tools on a DB2 client if you need to run the Query Patroller client tools from that client machine.

Procedure

To install the Query Patroller client tools on a Linux or UNIX DB2 client:

- Open the Query Patroller System Properties notebook.
- Log in as a user with root authority.
- Shut down any other programs so that the DB2 Setup wizard can update files as required.
- Check the Query Patroller product CD-ROM label to ensure that you are using the CD-ROM with the appropriate language.
- Mount the Query Patroller product CD-ROM.

- Change to the directory where the CD-ROM is mounted by entering the following command:

```
cd /cdrom
```

Where */cdrom* represents the mount point of the CD-ROM.

- Enter the `./db2setup` command to start the DB2 Setup wizard. The IBM DB2 Setup Launchpad window opens.
- Choose **Install a Product** from the DB2 Launchpad.
The DB2 Setup wizard opens.
- Click on **Work with existing** and choose the DB2 installation path where you want to install Query Patroller.
- Follow the instructions in the DB2 Setup wizard.

Online help is available to guide you through the installation. To invoke the online help, click the **Help** button on the bottom right hand corner of each installation window, or press **F1**. You can click the **Cancel** button at any time to end the installation.

Note: If you select **Computer usage based** from the Select the installation type window using the DB2 Setup wizard and you select **Server** from the Select how this computer will be used window, the Query Patroller Center is not installed. If you want to install the Query Patroller Center, select either the **Complete** or the **Custom installation** types from the Select the installation type window. You do not have to use the **Computer usage based** installation type to install the Query Patroller server. You can use either the **Complete** or the **Custom** installation types to install the Query Patroller server.

Results

When you have completed your installation, the Query Patroller client tools will be installed in the directory you specified in the DB2 Setup wizard.

What to do next

You have completed the task of installing Query Patroller server on a Linux or UNIX operating system using the DB2 Setup wizard. You are now ready to verify the installation of Query Patroller server. On Windows and Linux clients, you can also install the Query Patroller client tools using the DB2 Setup wizard.

Chapter 8. Installing Query Patroller with the DB2 Setup Wizard (Windows)

If you know which Query Patroller components and tools you want to install, and the kind of environment you will be installing on, you can follow the steps below to install Query Patroller on your Windows server and your Windows clients.

Important: With the new workload management features introduced in DB2 Version 9.5, Query Patroller and its components have been deprecated in Version 9.7 and might be removed in a future release. For more information, see the “DB2 Governor and Query Patroller have been deprecated” topic in the *What's New for DB2 Version 9.7* book.

To install Query Patroller (Windows):

1. Ensure that the *dyn_query_mgmt* database configuration parameter is set to DISABLE. This parameter must be set to ENABLE for Query Patroller to intercept and manage queries, but it must be set to DISABLE prior to installation to prevent the interception of any internal queries run by the installer. You must connect to the database to update *dyn_query_mgmt*.

If you upgraded your DB2 server to DB2 Version 9.7, the *dyn_query_mgmt* configuration parameter would have automatically been disabled during database upgrade. The owner of the instance containing the database that you want to run Query Patroller on must also have SETSESSIONUSER privilege on PUBLIC.

2. Install Query Patroller server with the DB2 Setup wizard (Windows).
3. Verify the installation of Query Patroller server.
4. Install the Query Patroller client tools with the DB2 Setup wizard (Windows).

After Query Patroller is installed, you must set the *dyn_query_mgmt* configuration parameter to ENABLE so that Query Patroller can intercept and manage queries.

Installing Query Patroller server using the DB2 Setup wizard (Windows)

This task outlines the steps for installing Query Patroller server on Windows using the DB2 Setup wizard.

- To install Query Patroller Version 9 server, a DB2 Version 9 server must already be installed. This list of qualified servers include:
 - DB2 Enterprise Server Edition
 - DB2 Workgroup Server Edition
 - DB2 Express Edition
- The instance running on your DB2 server must be stopped. To stop an instance, log on to the system as a user with SYSADM authority and issue the db2stop command at a Windows command prompt.
- All of the languages that the clients will use when submitting queries to Query Patroller server must be installed on the server in order to eliminate any server to client dependencies.

Online help is available to guide you through the installation. To invoke the online help, click the **Help** button on the bottom right hand corner of each installation window, or press **F1**. You can click the **Cancel** button at any time to end the installation.

To install Query Patroller server on Windows:

1. Log on to the system with the Administrator account you will use to install Query Patroller server.
2. Close all programs so the installation program can update files as required.
3. Insert the CD-ROM into the drive. If enabled, the auto-run feature automatically starts the DB2 Setup Launchpad. From this window you can proceed directly to the installation. If the auto-run does not work, use Windows Explorer to browse the DB2 product CD and double-click the setup icon.
4. Choose **Install a Product** from the DB2 Launchpad.
5. The DB2 Setup Wizard opens. Click on **Work with existing** and choose the DB2 installation path where you want to install Query Patroller.
6. Follow the instructions on the DB2 Setup Wizard.

Note: If you select **Computer usage based** from the **Select the installation type** window using the DB2 Setup wizard and you select **Server** from the **Select how this computer will be used** window, the Query Patroller Center is not installed. If you want to install the Query Patroller Center, select either the **Complete** or the **Custom installation** types from the **Select the installation type** window. You do not have to use the **Computer usage based** installation type to install the Query Patroller server. You can use either the **Complete** or the **Custom** installation types to install the Query Patroller server.

For information on errors encountered during installation, see the installation log file. This file stores general information and error messages resulting from install and uninstall activities. The installation log file is located in the My Documents\DB2LOG\ directory. The location of the My Documents directory depends on your computer settings. By default, the name of the installation log file follows this format:

```
DB2-<prod>-<Mon/Tue/...> <date>
```

For example:

```
DB2-QP_MON May 22 10-20-46 2006
```

You have completed the task of installing Query Patroller server on Windows using the DB2 Setup wizard.

Installing the Query Patroller client tools using the DB2 Setup Wizard (Windows)

This task outlines the steps for installing the Query Patroller client tools on a DB2 client for Windows using the DB2 Setup wizard.

A DB2 Version 9 product must be installed on the computer that you will install the Query Patroller client tools on. The following products are appropriate prerequisites:

- Any DB2 Version 9 client product
- Any DB2 Version 9 Connect product
- Any DB2 Version 9 server product

This task is part of the larger task of Installing Query Patroller (Windows).

Note that you do not need to install Query Patroller Client Tools on a DB2 client if you only want to submit queries from that client to the Query Patroller server. You only need to install Query Patroller Client tools on a DB2 client if you need to run the Query Patroller client tools from that client machine.

To install the Query Patroller client tools on Windows:

1. Open the Query Patroller System Properties notebook:
2. Log on to the system with the Administrator account you will use to install the Query Patroller client tools.
3. Stop the DB2 administration server (DAS) and any DB2 instances related to the DB2 installation path where Query Patroller will be installed so that the DB2 Setup wizard can update files as required.
4. Insert the CD-ROM into the drive. If enabled, the auto-run feature automatically starts the DB2 Setup Launchpad. From this window you can proceed directly to the installation.
5. The DB2 Setup wizard determines the system language, and launches the setup program for that language. If you want to run the setup program in a different language, or if the setup program fails to auto-start, you can start the DB2 Setup wizard manually. To start the DB2 Setup wizard manually:

- a. Click **Start** -> **Run**.

- b. In the **Open** field, enter the following command:

```
x:\setup /i language
```

where:

- *x*: represents your CD-ROM drive
- *language* represents the territory identifier for your language (for example, EN for English).

If the /i flag is not specified, the installation program runs in the default language of the operating system.

- c. Click **OK**.

6. Choose **Install Products** once the DB2 Launchpad opens.
7. Click **Next** on the **Select the product you would like to install** window.
8. On the **Welcome to the DB2 Setup wizard** window, click **Next**.
9. On the **License Agreement** window, select **I accept the terms in the license agreement** if you accept the terms. Click **Next**. If you do not accept the terms, then select **I do not accept the terms in the license agreement**. Click **Cancel** to end the installation.
10. On the **Select the installation type** window select **Computer usage based** and click **Next**.
11. On the **Select how this computer will be used** window select **Administrator or end user with self-service capability** and click **Next**.
12. Proceed by following the setup program's prompts. Online help is available to guide you through the remaining steps. To invoke the online help, click the **Help** button on the bottom right hand corner of each installation window, or press **F1**. You can click the **Cancel** button at any time to end the installation.
- 13.

You are now ready to configure your Query Patroller tool-enabled DB2 client to access a remote Query Patroller server.

Note: Query Patroller Client Tools must connect to a database partition on which the Query Patroller server is installed.

Chapter 9. Configuring the client and server connection

Configuring a client to Query Patroller server connection using the command line processor

This task describes how to configure a connection from a client to a Query Patroller server using the command line processor (CLP).

Before you configure a client to Query Patroller server connection:

- Communications must be configured on the client computer. Communications must be TCP/IP.
- Communications must be configured on the Query Patroller server. Communications must be TCP/IP.

To configure a client to Query Patroller server connection using the command line processor:

1. Catalog the TCP/IP node on the DB2 client.
2. Catalog the database on the DB2 client.
3. Test the client to server connection.

Part 4. Configuring Query Patroller

Chapter 10. Query Patroller configuration roadmap

You can use Query Patroller to help you address specific business goals, such as preventing runaway queries from monopolizing your system. Table 3 indicates the specific configuration settings that apply to different business goals.

Important: With the new workload management features introduced in DB2 Version 9.5, Query Patroller and its components have been deprecated in Version 9.7 and might be removed in a future release. For more information, see the “DB2 Governor and Query Patroller have been deprecated” topic in the *What’s New for DB2 Version 9.7* book.

Table 3. Road map to configuration information

Goal	Parameter to set	Topics to read
To prevent runaway queries or limit the size of queries from particular submitters	Maximum cost of a query allowed (MAX_COST_ALLOWED) in submitter profiles	Submitter profile configuration
To limit the number of queries from a particular submitter or group	Maximum number of queries allowed (MAX_QUERIES_ALLOWED) in submitter profiles	Submitter profile configuration
To set the optimum number of concurrent queries in the system	Maximum total queries in system (MAX_TOTAL_QUERIES) in system settings	Query Patroller thresholds
To limit the amount of system resources used for all queries or for intercepted queries	Maximum system workload cost (MAX_TOTAL_COST) in system settings	Query Patroller thresholds
To allow the queries from some submitters or groups to bypass Query Patroller	Select do not intercept option (INTERCEPT set to 'N') in submitter profiles or use the DB2_QP_BYPASS_USERS registry variable	Submitter profile configuration or Query Patroller variables
To allow the queries from some applications to bypass interception by Query Patroller	Specify which applications to not intercept (INTERCEPT_APPLICATION and EXCLUDE_APPLICATIONS) in system settings or use the DB2_QP_BYPASS_APPLICATIONS registry variable	Setting query thresholds for the Query Patroller system or Query Patroller variables
To specify particular applications for Query Patroller to intercept queries from	Specify which applications to intercept (INTERCEPT_APPLICATION and INCLUDE_APPLICATIONS) in system settings	Setting query thresholds for the Query Patroller system
To run small queries without Query Patroller managing them	Minimum cost to manage (MIN_COST_TO_MANAGE) in submitter profiles or use the DB2_QP_BYPASS_COST registry variable	Submitter profile configuration or Query Patroller variables
To prioritize queries of particular users when these queries are in the queue	Queue priority (PRIORITY) in submitter profiles	Submitter profile configuration

Table 3. Road map to configuration information (continued)

Goal	Parameter to set	Topics to read
To prevent small queries from getting queued behind large queries	Set up query classes based on grouping queries by size	Query class configuration
To limit the number of concurrent large queries	Set up a query class for large queries	Query class configuration
To manage disk space used by result tables and historical query information	Set up purge jobs for removing result tables, managed query information and historical query information	Setting Query Patroller maintenance schedules for queries and result tables

Setting up Query Patroller server manually

This task outlines how to manually set up Query Patroller server on a specified DB2 database.

- Query Patroller server must be installed on your computer.
- A DB2 instance must be running on your computer. This is the instance that contains the database that Query Patroller server will monitor.
- You must have SYSADM authority to create new table spaces using the `qpsetup` command.
- You require DBADM authority to use existing table spaces using the `qpsetup` command.
- Only one table space container can be specified. This is to avoid an overly-complicated installation process and command line input. If you require more than one container for the table space, then the table space must be created before running the `qpsetup` command.
- You cannot specify the buffer pool. The table space will be created using the default buffer pool, `IBMDEFAULTBP`. If you require the table space to use a different buffer pool, then the buffer pool and the table space must be created before running the `qpsetup` command.
- Since the table space will be created using the default buffer pool, `IBMDEFAULTBP`, which has a 4K page size, the page size for the table space will be 4K as well. If you require table space with page sizes other than 4K, the table space and its buffer pool must be created before running the `qpsetup` command.
- If you specify a database partition group that spans more than one logical database partition in the `CONTROL_DBPARTITIONGROUP` and `RESULT_DBPARTITIONGROUP` parameter of the `qpsetup` command, then the following restriction applies:
 - You can specify a database partition expression for container string syntax when creating either SMS or DMS containers. You would typically specify the database partition expression if you were using multiple logical database partitions in the partitioned database system. This ensures that container names are unique across database partition servers. When you specify the expression, the database partition number is part of the container name or, if you specify additional arguments, the result of the argument is part of the container name.

The DB2 Setup wizard will automatically run the `qpsetup` command during the Query Patroller server installation (UNIX or Windows). However, if you decide not to automatically run this command during the installation you can manually run it

after the installation by issuing the **qpsetup** command from a UNIX shell prompt or from a Windows command prompt. You can use qpsetup in the following two situations:

1. You can use qpsetup if you install the Query Patroller server files without setting up Query Patroller server.
2. You can use qpsetup if you decide to manage a different database after you install Query Patroller server.

The qpsetup command performs the following:

- Creates either an SMS or a DMS table space for the Query Patroller control tables, if the specified table space does not exist.
- Creates either an SMS or a DMS table space for the Query Patroller result tables, if the specified table space does not exist.
- Creates Query Patroller control tables on the specified table space, if they do not exist.
- Binds Query Patroller packages to the database.

You need to run the qpsetup command for each database that you want to use Query Patroller with.

To set up Query Patroller server on a specified DB2 database:

1. Check the appropriate command syntax and command parameters before issuing the qpsetup command. For more information see the qpsetup command topic in the related links.
2. Issue the qpsetup command at a UNIX shell prompt or a Windows command prompt, depending on your operating system.

Query Patroller administration tasks overview

Once you have completed your installation, set up, and upgrade, and you have enabled Query Patroller to intercept queries, you can begin to administer and use your Query Patroller system. The following diagram presents an overview of the different Query Patroller post-installation tasks.

Important: With the new workload management features introduced in DB2 Version 9.5, Query Patroller and its components have been deprecated in Version 9.7 and might be removed in a future release. For more information, see the “DB2 Governor and Query Patroller have been deprecated” topic in the *What’s New for DB2 Version 9.7* book.




Figure 4. Query Patroller administration tasks overview

Starting Query Patroller

You might need to start or stop Query Patroller during normal business operations; for example, after creating, removing, or changing the maximum query cost of query classes, if you want the changes you have made to take effect immediately, rather than waiting for all queued and running queries to complete, you must stop and restart Query Patroller.

- You must have DBADM authority.
- DB2 must be started.

To start Query Patroller, issue the `qpstart` command.

Enabling Query Patroller to intercept queries

Before you can use the interception, management, and historical analysis functions of Query Patroller, you must enable Query Patroller to intercept queries by setting the `dyn_query_mgmt` configuration parameter.

You must be connected to the database.

To enable Query Patroller to intercept queries, use any of the following methods.

- DB2 Control Center method
- Command line method

Note: In a partitioned database environment, Query Patroller must be enabled for all database partitions.

- To enable Query Patroller to intercept queries using the Query Patroller Center method:
 1. Open the Database Configuration window.
 2. Click on the **Value** for DYN_QUERY_MGMT. A  button appears.
 3. Click . The Change Database Configuration Parameter window opens.
 4. Click the **Enable** radio button.
 5. Click **OK** to close the Change Database Configuration Parameter window.
 6. Click **OK** to close the Database Configuration window and save your changes.
- To enable Query Patroller to intercept queries using the Command line method:
 1. Issue the UPDATE DATABASE CONFIGURATION command and set *dyn_query_mgmt* to ENABLE for the database you want Query Patroller to intercept queries from.

Stopping Query Patroller

You might need to start or stop Query Patroller during normal business operations; for example, after creating, removing, or changing the maximum query cost of query classes, if you want the changes you have made to take effect immediately, rather than waiting for all queued and running queries to complete, you must stop and restart Query Patroller. After you have stopped Query Patroller, forced queries will be in an inconsistent state until you restart Query Patroller and query recovery is completed.

- You must have DBADM authority.
- Query Patroller must be started.

To stop Query Patroller, issue the `qptest` command.

Query processing by Query Patroller

Query Patroller interacts with DB2 to execute several processes when a query is submitted. This topic covers the processing that takes place from the time a query is submitted until the time results are generated.

Query submission

Query Patroller can intercept queries submitted through a variety of methods, including:

- A query application
- A middle tier business intelligence tool
- The DB2 graphical user interface (dynamically)
- A command line interface

When a query is submitted, DB2 checks the *dyn_query_mgmt* database configuration parameter. If the parameter is set to ENABLE, then Query Patroller captures information about the query such as the SQL statement and the submitter's ID. Query Patroller also evaluates other information at this point, such

as the query submitter's minimum cost to manage value for a query (MIN_COST_TO_MANAGE), and the submitter's maximum cost of a query value (MAX_COST_ALLOWED).

Cost analysis

After a query is submitted, the DB2 query optimizer performs a cost estimation on the query. The query optimizer provides Query Patroller with an estimated cost to run the query, measured in timerons.

Query interception

Query Patroller determines whether to intercept the query or allow it to run directly against the database.

You can set properties at the system level or at the submitter level to specify which queries Query Patroller intercepts. At the system level, query interception is determined based on the application name. At the submitter level, query interception is determined based on a submitter profile property.

If the query is not intercepted by Query Patroller, the query runs directly against the database without further Query Patroller involvement.

Query management

If the query's estimated cost is less than the submitter's minimum cost to manage value for a query (MIN_COST_TO_MANAGE), Query Patroller allows the query to run directly against the database.

If the query's estimated cost is greater than the submitter's minimum cost to manage value for a query, then Query Patroller assesses, prioritizes, and holds or queues the query.

Query Patroller holds the query under either of the following conditions:

- The estimated cost of the query exceeds the submitter's maximum cost of a query value (MAX_COST_ALLOWED).
- The estimated cost of the query exceeds the maximum workload cost value for the system (MAX_TOTAL_COST).

Queries that are in a held state will not run unless an administrator or operator manually releases the query, or a scheduled release job runs.

Query Patroller queues the query under any of the following conditions:

- The maximum number of queries value for the submitter (MAX_QUERIES_ALLOWED) has already been reached.
- The maximum number of queries value for the system (MAX_TOTAL_QUERIES) has already been reached.
- The estimated cost of the query plus the current workload cost exceeds the maximum workload cost value for the system (MAX_TOTAL_COST).
- The maximum number of queries value for the query class in which the query runs (MAX_QUERIES) value has already been reached.

Query Patroller routinely checks the query queue to identify queries that are eligible to execute. For example, if a submitter is allowed to submit five concurrent

queries, then the sixth query is put in the queued state. However, as soon as one of the other five queries completes, Query Patroller executes the sixth query, provided that the query falls within other system and submitter thresholds.

If query classes have not been defined, the query executes within the default query class. If query classes have been defined, the query executes within the appropriate query class.

Query execution

DB2 executes the query, and Query Patroller directs the result set to one of two destinations, based on settings in the submitter's submission preferences, illustrated in Table 4.

Table 4. Result set destinations

Destination	Query Patroller Center method	Command line method
The original DB2 agent process executes the query and returns the result set back to the application that submitted the query.	Wait until the results are returned is selected in the Query Submission Preferences window.	UPDATE SUBMISSION_PREFERENCES or ADD SUBMISSION_PREFERENCES command is issued with a value of 'A' for the RESULT_DESTINATION parameter.
A new DB2 agent process is created to execute the query and store the result set in a result table.	Release the application and retrieve the results from a result table is selected in the Query Submission Preferences window	UPDATE SUBMISSION_PREFERENCES or ADD SUBMISSION_PREFERENCES command is issued with a value of 'T' for the RESULT_DESTINATION parameter.

Query executions can be canceled or put in background any time from submission until they have been completed.

Notification

If the submitter indicated that the query's result set is to be directed to a result table and e-mail notification has been set up for the submitter, Query Patroller sends e-mail notification to the submitter when the result table is created.

The following diagram depicts the query processing described in this topic.




Figure 5. Query processing by Query Patroller

Chapter 11. Query Patroller system configuration

Query Patroller allows you to control the size of the query workload that runs concurrently against your database, either by setting a limit on the number of queries that run concurrently in the system, by setting a limit on the total cost of all queries that run concurrently, or both.

By default, both of these settings are set to unlimited. In most situations, these defaults will be sufficient because submitter-level resource limits and resource limits within query classes will effectively constrain the size of query workload that can run concurrently. It is strongly recommended that you use query classes and submitter-level resource limits to constrain database activity. However, if you wish to configure one or both of these system-level settings, you can use the following heuristics to guide you in determining the best settings for your system.

Setting total queries (MAX_TOTAL_QUERIES)

Your ability to effectively set a limit on the total number of managed queries depends on the nature of your database workload. Specifically, the degree of variation in the size of queries in your workload will determine whether or not it is beneficial to set this threshold.

If your workload is homogenous and includes queries that are similar in terms of estimated cost, you can determine the appropriate setting for this threshold by running test workloads with varying numbers of queries, and monitoring your system performance for how well it responds to the different size workloads.

If your workload is not homogenous, but is consistent in terms of composition, for example 75 percent small and 25 percent large queries, you can determine the appropriate setting for this threshold by running a set of test workload that mimic this composition. Monitor the performance of your system as you gradually increase the number of queries in the test workload to establish the optimal number of queries that should be allowed to run concurrently in your system.

Note:

1. It is strongly recommended that you run these test workloads prior to defining any query classes in your system.
2. Queries that are unintercepted or unmanaged by Query Patroller are not counted for the purposes of the MAX_TOTAL_QUERIES threshold. This means that, if you expect a significant number of queries to run unintercepted or unmanaged, you will need to take these queries into account when determining how many queries you can run concurrently in your system. For example, if you determine that you can comfortably run 1000 queries concurrently on your system, and you typically run 200 queries that are unintercepted or unmanaged, then you would set your maximum total number of queries to 800.

If your database workload is inconsistent in terms of the size of queries it includes, it may be extremely difficult to determine an appropriate setting for the maximum number of concurrent queries (MAX_TOTAL_QUERIES) threshold since system performance will vary based on both the size and the number of queries in the workload.

After setting the maximum number of concurrent queries (MAX_TOTAL_QUERIES) threshold, continue to monitor performance to ensure you have the correct settings. You might have to reevaluate your settings after significant database restructuring or changes to your typical database workload.

Setting total workload cost (MAX_TOTAL_COST)

Query Patroller calculates the total cost of your current database workload by adding up the estimated cost of all managed queries currently running in the system. The estimated cost of each query is calculated on the entire execution of the query, which might be spread out across a few minutes or several hours.

It is often difficult to set a useful limit on the total cost of all queries (MAX_TOTAL_COST) in the system because ten queries that cost 100,000 timerons have the same total cost as a single query costing 1,000,000 timerons, even though it would likely be more taxing on your system resources to run ten queries simultaneously in one hour than to run a single query over several hours.

To effectively ascertain how many timerons your system can handle at one point in time, you need to figure out how many timerons of work your system can handle in a particular period of time, or time slice. The following steps will guide you in doing this.

1. Choose a query or set of queries that are representative of your workload. You should be familiar with these queries and have a good sense of how long these queries typically take to execute.
2. Perform some test executions of the query or queries.
3. Use the historical analysis Queries report to see the execution time for the queries and the estimated cost of the queries in timerons.
4. To get a more refined measurement of cost, divide the query execution time into time slices. For example, divide a 50 minute query into ten five-minute time slices. Assuming that cost is relatively constant throughout the execution of the query, you can estimate that the cost of execution during each time slice is one tenth of the total cost of the query. For example, a query costing 100,000 timerons that takes 50 minutes might be divided into ten time slices costing 10,000 timerons each.
5. Determine how many of these representative queries you can run simultaneously before performance begins to deteriorate. For example, if you can run 20 such queries simultaneously before performance degradation, you might determine that your system can run 20,000,000 timerons concurrently.

Note: Note: Queries that are unintercepted or unmanaged by Query Patroller are not counted for the purposes of the MAX_TOTAL_COST parameter. This means that, if you expect a significant number of queries to run unintercepted or unmanaged, you will need to take these queries into account when determining the maximum cost of the workload you can run concurrently in your system

6. Set the maximum number of concurrent queries (MAX_TOTAL_COST) threshold, either through the Query Patroller Center interface or the command line.
7. Continue to monitor performance to ensure you have the correct settings. You might have to reevaluate your settings after significant database restructuring or changes to your typical database workload.

Setting query thresholds for the Query Patroller system

You can control the overall workload that executes against your database by setting system thresholds for cost and number of queries. By default, these query thresholds are set as unlimited. You may want to set more restrictive thresholds initially or you may want to wait until you've collected historical data that can guide you. First specify the applications you want to intercept, then set system-wide limits for managed queries.

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile has the CONFIGURATION privilege with edit authority

You can set query thresholds for the Query Patroller system using one of the following methods:

- Query Patroller Center method
- Command line method
- To set query thresholds using the Query Patroller Center method:
 1. Open the Query Patroller System Properties notebook:
 - a. From the Query Patroller Center, expand the object tree under the **Configuration** folder to find the **Query Patroller System** folder.
 - b. Click the **Query Patroller System** folder. A system property record is displayed in the pane on the right side of the window (the contents pane).
 - c. Right-click the system property record, and click **Properties** in the pop-up menu. The Query Patroller System Properties notebook opens.

You can also open the Query Patroller System Properties notebook by double-clicking the entry in the contents pane when the Query Patroller System folder is displayed.

2. Click the **Thresholds** tab.
3. To specify the applications that you want to intercept, select one of the following options:
 - a. **All applications**
 - b. **No applications or only the applications listed.** Type the executable file name for applications that are to be intercepted by Query Patroller. Separate multiple applications with commas.
 - c. **All applications except those listed.** Type the executable file name for applications that are not to be intercepted by Query Patroller. Separate multiple applications with commas.

Application names are case-sensitive.

4. In the **Maximum number of queries** field, type the number of queries that can be running simultaneously on the system. The value you enter will depend on the following factors:
 - The performance of your database
 - The performance of your database
 - The number of users who are submitting queries
 - The average cost of the queries being submitted at any given time

To set the Maximum number of queries value to unlimited, leave the field blank.

5. In the **Maximum workload cost** field, type a number to represent the maximum workload cost threshold, in timerons. To set the value to unlimited, leave the field blank.
 6. Click **OK** to accept your entries.
- To set query thresholds using the Command line method, issue the UPDATE QP_SYSTEM command using the following parameters:
 - INTERCEPT_APPLICATION
 - INCLUDE_APPLICATIONS
 - EXCLUDE_APPLICATIONS
 - MAX_TOTAL_QUERIES
 - MAX_TOTAL_COST

Updating the list of databases in Query Patroller

You can update the list of databases that you can select to work with in the Query Patroller Center.

To update the list of databases, use the following Query Patroller Center method:

1. In the Query Patroller Center, click the **Update Database List** push button. The Update Database List window opens.
2. Use the arrow buttons to move the databases that you want to add from the **Available databases** list to the **Selected databases** list.
3. Click **OK** to update the list of databases that you can work with in the Query Patroller Center.

Enabling e-mail notification of Query Patroller submitters

Enable e-mail notification if you want submitters to be notified when a query completes and a result table is created, or if an error occurs during the processing of a query that would have had its results sent to a result table.

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile has the CONFIGURATION privilege with edit authority

You must also have access to an SMTP mail server.

Note: A result table is created for submitters when:

- The query submitter's submission preferences specify that the application should be released after a query is submitted
- A query was held and then run
- A query was run in the background

You can enable e-mail notification using one of the following methods:

- Query Patroller Center method
- Command line method
- To enable e-mail notification using the Query Patroller Center method
 1. Open the Query Patroller System Properties notebook:
 - a. From the Query Patroller Center, expand the object tree under the **Configuration** folder to find the **Query Patroller System** folder.

- b. Click the **Query Patroller System** folder. A system property record is displayed in the pane on the right side of the window (the contents pane).
- c. Right-click the system property record, and click **Properties** in the popup menu. The Query Patroller System Properties notebook opens.

Note: You can also open the Query Patroller System Properties notebook by double-clicking the entry in the contents pane when the Query Patroller System folder is displayed.

2. Click the **E-mail** tab.
3. Select the **Enable e-mail notification** check box.
4. In the **E-mail server** field, enter the host name or IP address of the SMTP e-mail server you want to use.
5. For submitters who do not have an e-mail address specified in their query submission preferences, select either **Do not send e-mail** or **Send e-mail to a designated address**. If you are sending to a designated address, such as the administrator's address or a support desk, enter the address in the field.

Note: Submitters can specify the e-mail addresses they want to use in the Query Submission Preferences window.

6. Click **OK** to enable e-mail notification and close the notebook.
- To enable e-mail notification using the Command line method:
 1. Issue the UPDATE QP_SYSTEM command using the following parameters:
 - EMAIL_ENABL
 - EMAIL_SERVER
 - SEND_DESIGNATED
 - DESIGNATED_EMAIL_ADDRESS

Enabling collection of historical data

Query Patroller historical analysis can be a powerful tool for analyzing data warehouse usage. You can choose to collect historical data only for queries that were managed by Query Patroller, or for all queries intercepted by Query Patroller. Queries that are not intercepted cannot have historical data collected about them.

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile has the CONFIGURATION privilege with edit authority

To enable collection of historical data, use one of the following methods.

- Query Patroller Center method
- Command line method

The default behavior is that only managed queries will have historical data collected.

Note: If you are planning to collect data regarding query execution time, you need to ensure that the DB2 timestamp and statement monitor switches are set to 'ON'. If you are planning to collect data about how many rows are returned by queries, you need to ensure that the DB2 statement monitor switch is set to 'ON'.

- To enable collection of historical data using the Query Patroller Center method:

1. Open the Query Patroller System Properties notebook.
 2. Click the **Options** tab.
 3. Under **Historical analysis**, specify the queries that you want to save by selecting the **Only managed queries** or the **All intercepted queries** radio button.
 4. Click **OK** to save your changes.
- To enable collection of historical data using the Command line method:
 1. Issue the UPDATE QP_SYSTEM command using the following parameter:
 - QUERIES_TO_SAVE

Chapter 12. Managing users

Administering operators

Query Patroller operators

A Query Patroller operator is an ID used by Query Patroller to represent a user or a group of users who has a subset of administrator authorities and tasks defined in the operator's profile. Operator IDs are not defined by Query Patroller; they are created by selecting from existing DB2 user or group IDs.

Note: The ID that has DBADM authorities is automatically the Query Patroller administrator.

A Query Patroller operator performs some or all of the following administrative tasks.

- **Configuration**

Involves creating or deleting query classes and setting system-wide thresholds, as well as other configuration tasks such as setting up e-mail notification.

- **Monitoring**

Involves changing the status of queries, viewing the SQL of managed queries, and deleting result tables that are not needed.

- **User administration**

Involves creating, modifying, and deleting submitter profiles. Also involves removing managed queries that have completed and are no longer needed, and manually removing held queries.

- **Historical analysis**

Involves removing historical queries that are no longer needed.

Query Patroller operator profiles

The role of a Query Patroller operator is usually filled by users in a front-line support role, such as help-desk representatives.

The kinds of tasks that an operator can perform depend on the level of authority that the operator has: view authority or edit authority.

For instance, an operator whose profile has the MONITORING privilege with edit authority can change the status of queries, view the SQL of managed queries, and delete result tables that are no longer needed. However, an operator whose profile has the MONITORING privilege with view authority can only view the status of queries.

Active and suspended operator profiles

When an operator profile is suspended, users associated with the profile can still access the Query Patroller Center as submitters, provided that they are associated with an active submitter profile. However, they cannot access any other parts of the Query Patroller Center.

Which operator profile Query Patroller uses

When an operator's user ID has multiple group operator profiles associated with it, Query Patroller determines the operator's privileges by merging the privileges of all the operator profiles that the user is associated with. The privileges are merged so that the operator is granted the highest set of privileges from the operator profiles.

Creating operator profiles for users and groups

You can create a Query Patroller operator profile based on an existing operator profile or you can create an operator profile with new settings.

You must have DBADM authority.

You can create a Query Patroller operator profile using one of the following methods:

- Query Patroller Center method
- Command line method
- To create a Query Patroller operator profile using the Query Patroller Center method
 1. Open the Create Operator window:
 - a. If you are creating an operator profile based on an existing operator profile open the Create Operator window as follows:
 - 1) From the Query Patroller Center, expand the object tree under the **User Administration** folder to find the **Operators** folder.
 - 2) Click the **Operators** folder. Any existing operators are displayed in the pane on the right side of the window (the contents pane).
 - 3) In the contents pane, right-click the operator whose profile you want to use as the basis for the new operator you are creating, and click **Create Like** in the pop-up menu. The Create Operator window opens with certain fields prefilled.
 - b. If you are creating an operator profile with new settings, open the Create Operator window:
 - 1) From the Query Patroller Center, expand the object tree under the **User Administration** folder to find the **Operators** folder.
 - 2) Right-click the **Operators** folder, and click **Create** in the pop-up menu. The Create Operator window opens.
 2. In the **Operator** field, type a name for the operator you are creating. The name must exist as a DB2 authorization ID. This profile ID is case-sensitive. This means that if you create an operator profile for a user identified as "TESTUSER" there must also be an existing DB2 authorization ID called "TESTUSER". If you create an operator profile for a user identified as "testuser," this profile will not be associated with the DB2 authorization ID "TESTUSER" and will not be used by Query Patroller
 3. Use the **Profile type** field to select whether the operator profile applies to a user or a group.
 4. Optional: If you want to temporarily suspend the new operator's ability to access all parts of the Query Patroller Center, select the **Access suspended** check box.

Note: If this is a group operator profile, suspending access does not suspend all users belonging to that group, unless they have no other operator profiles.

Suspending this profile will, however, prevent this profile from giving these users any additional privileges beyond those assigned to the operator profiles of the other groups to which they belong.

5. Use the **Configuration** field to specify the level of access the operator will have to the Configuration tree element in the Query Patroller Center. The level of access you specify determines whether or not the operator can create or delete query classes or set system-wide thresholds.
 6. Use the **Monitoring** field to specify the level of access the operator will have to the Monitoring tree element in the Query Patroller Center. The level of access you specify determines whether or not the operator can modify the status of a query or delete a result table.
 7. Use the **User administration** field to specify the level of access the operator will have to the User Administration tree element in the Query Patroller Center. The level of access you specify determines whether or not the operator can create, modify, or delete submitter profiles.
 8. Use the **Historical analysis** field to specify the level of access the operator will have to the Historical Analysis tree element in the Query Patroller Center. The level of access you specify determines whether or not the operator can remove and view historical query data.
 9. Click **OK** to create the new operator profile.
- To create a Query Patroller operator profile using the Command line method:
 1. Issue the `ADD OPERATOR_PROFILE` command.

Note: Operator profiles do not need to be created for users with DBADM authority on a database. Such users already possess the maximum level of operator privileges, therefore adding operator profiles for them is redundant. It may also be misleading to create an operator profile for a user with DBADM authority since the user can automatically perform all Query Patroller tasks despite any restrictions on the operator privileges associated with the profile.

Suspending or restoring operator privileges for users and groups

You can temporarily suspend an operator's ability to access all parts of the Query Patroller Center, then restore these privileges at a later time.

You must have DBADM authority.

You can suspend privileges for an operator using one of the following methods:

- Query Patroller Center method
- Command line method

Note:

1. A user with a suspended operator profile can still access the Query Patroller Center as a submitter if he has an active submitter profile.
2. If you want to suspend or restore an operator's access to only some parts of the Query Patroller Center
 - To suspend privileges for an operator using the Query Patroller Center method
 1. Open the Operator Properties window:
 - a. From the Query Patroller Center, expand the object tree under the **User Administration** folder to find the **Operators** folder.

- b. Click the **Operators** folder. Any existing operators are displayed in the pane on the right side of the window (the contents pane).
- c. Right-click the operator profile that you want to change in the contents pane, and click **Properties** in the pop-up menu. The Operator Properties window opens.

Note: You can also open the Operator Properties window by double-clicking the operator profile that you want to change.

2. Select the **Access suspended** check box to suspend the operator's ability to access all parts of the Query Patroller Center, or clear it to restore the submitter's ability to access the indicated areas of the Query Patroller Center.

Note: If this is a group operator profile, suspending access does not suspend all users belonging to that group, unless they have no other operator profiles. Suspending this profile will, however, prevent this profile from giving these users any additional privileges beyond those assigned to the operator profiles of the other groups to which they belong.

3. Click **OK** to close the Operator Properties window.
- To suspend privileges for an operator using the Command line method:
 1. Issue the `UPDATE OPERATOR_PROFILE` command using the following parameters:
 - `SUSPENDED` parameter

Submitter profile configuration

Submitter profiles contain settings for submitter query interception, submitter resource limits, and submitter queue priority. Before creating your submitter profiles, you need to plan how you will configure these settings to best meet the needs of your users and your organization.

Step 1: Determine what type of submitter distinctions can be made in your system

Before you can plan for submitter profiles for your users and groups, you need to find out what submitter distinctions can be made based on the type of user information that is available to Query Patroller. A submitter may be a user, a group of users, an application, or a server that submits queries on behalf of a user.

In a two-tier setup where users are connecting directly to the database and submitting queries directly to Query Patroller, you can distinguish between every submitter who submits queries.

In a three-tier setup where users are submitting queries through a third-party submission application, you can only distinguish between submitters if the tool passes the appropriate user information on to Query Patroller. For example, if the submission application connects to the database with the DB2 ID of the user who submits the query, Query Patroller can use the submitter profile for the user to determine how to handle the query. If the submission application connects to the database using a different ID than the user ID submitting the query, Query Patroller uses the profile for the connecting ID to determine how to handle the query. In such cases, all queries submitted through the submission application will be identified as originating from a single submitter.

In situations where different user groups use different applications to submit their queries, and you can differentiate resource allocation, priority, or other user characteristics based on these groups.

Step 2: Determine the nature of the resource requirements that characterize each submitter or group of submitters

You may already be familiar with the warehouse usage patterns associated with the different departments or groups running queries against your database. However, if you do not know what kinds of queries are typically submitted by each group, you can use the Historical Analysis feature to collect historical data to get a representative sample of database activity and analyze the submitter activity data.

You can use the Submitters reports to view the following submitter activity information:

- the size of queries that are submitted by different submitters
- the number of queries submitted by a submitter or group of submitters within a given time frame
- which submitters generate unusually large queries

Step 3: Decide which individual or group submitters should be intercepted by Query Patroller

From the information about the query requirements of the different submitters in your system, you can get a sense of whether there are submitters in your organization whose queries are consistently so small that they are not worth intercepting. If such submitters can be identified, you may want to set up their submitter profiles such that queries from these submitters are not intercepted.

There may also be users in your organization whose response time requirements are so great that they cannot spare even the slightest performance cost that is incurred when a query is intercepted. These users should also be associated with submitter profiles whose queries will not be intercepted.

If you are deploying Query Patroller as a pilot or test project, you may want to intercept only the submitters who are participating in the project. The simplest way to do this is to set your test submitter profiles to be intercepted, and set the PUBLIC submitter profile to not be intercepted by Query Patroller.

Note: In a production environment, it is recommended that you intercept queries from any submitter who has the ability to submit ad hoc queries against the database.

Keep in mind that when a query is submitted under a submitter profile that is set to not be intercepted by Query Patroller (`INTERCEPT='N'`) the query cannot be tracked for the purposes of Historical Analysis. This means that once a submitter profile is set to not be intercepted by Query Patroller, any query activity submitted under this profile will not appear in any resource usage reports that you generate.

You can also use the `DB2_QP_BYPASS_USERS` registry variable to allow the queries from specific submitters to bypass Query Patroller.

Step 4: Decide which thresholds to set for submitters whose queries will be intercepted

There are several different resource thresholds that can be set within a submitter profile. Any or all of these thresholds can remain set to the default value, but in most cases where distinctions between submitter resource allocations need to be made, you will want to use one or more of these thresholds to control resource usage at the individual or group level.

Note: When setting resource thresholds for group profiles, keep in mind that the limit that you set applies to each individual in the group each time a query is submitted. This is not a resource limit for the entire group's submissions.

The thresholds are presented in the following table, along with their default values and their potential uses. For a more detailed description of each submitter profile parameter, see the description of the `ADD_SUBMITTER_PROFILE` command.

Table 5. Submitter threshold defaults and usage

Threshold	Parameter	Default value	Usage
Maximum query cost	<code>MAX_COST_ALLOWED</code>	10,000,000 timerons	Controls size of individual queries. Use to constrain any "problem" submitters or to prevent runaway queries.
Maximum number of queries	<code>MAX_QUERIES_ALLOWED</code>	100 queries	Controls maximum number of concurrent queries. Use to prevent submitters from monopolizing system resources by submitting too many queries at once.
Maximum result rows	<code>MAX_RESULT_ROWS</code>	1,000,000 rows	Controls the number of result rows that can be stored in a result table for a single query. Use to limit the amount of disk space taken up by large query results.
Minimum cost to manage	<code>MIN_COST_TO_MANAGE</code>	15,000 timerons	Determines whether a particular query will be managed based on size. Use to exclude small queries whose management would incur a noticeable performance hit relative to execution time.

From the information about the query activity by the different submitters in your system, you should be able to see the size of the queries that are typically submitted by each submitter or group of submitters. This will give you a sense of the value (in timerons) that would be appropriate for the maximum query cost (`MAX_COST_ALLOWED`) that you can assign to each submitter group. For some submitters, you may not want to change this value from the default. However, for submitters who occasionally submit queries that are excessively large, you can set a value in their submitter profile for the maximum query cost based on what you observe to be an acceptable size query, and allow Query Patroller to hold any queries that exceed this value. If an excessively large query is justified, a submitter will have to either notify the Query Patroller administrator to manually run the query despite its size, or the query can be run along with other held queries at a time when resource usage is lower.

Similarly, if you observe or know of a problem with certain individuals or groups submitting too many queries at once, you can set the maximum number of queries (`MAX_QUERIES_ALLOWED`) in their submitter profile to a value that seems reasonable given the resource requirements of the submitter.

Note: If your environment is one in which different submitters are not distinguished, you should set the maximum queries allowed

(MAX_QUERIES_ALLOWED) threshold to be unlimited (-1) because all intercepted queries will be identified as belonging to a single submitter.

If disk space limitations are a concern, you can set a limit on the number of result rows (MAX_RESULT_ROWS) that are stored in a result table for a query. You can also see from the submitter activity reports what the typical sizes of result sets are for different submitters. This will give you an indication of what values are appropriate for this submitter threshold for the different profiles.

The minimum cost to manage (MIN_COST_TO_MANAGE) threshold can be set to allow small queries to bypass Query Patroller management. You may have to lower this value if you feel that not enough of your workload is being managed by Query Patroller, or you may have to raise this value if you feel that too many queries are being managed by Query Patroller.

You can still specify that these queries are tracked for the purposes of Historical Analysis, but be aware that this tracking will have a small performance cost since it involves table updates for every query.

Step 5: Determine the queue priority level for each submitter profile.

The queue priority level for a submitter profile determines the order in which queued queries are run. The higher the queue priority level for a submitter, the sooner his or her queries are selected from the queue to run. You may want to adjust the queue priority level for a submitter or group of submitters if you have some submitters whose queries need to be managed by Query Patroller, but which have a higher degree of urgency than other submitter's queries.

Since only managed queries are queued, queue priority only affects queries that are intercepted and managed.

Query Patroller submitters

A submitter is an ID used by Query Patroller to represent any person, group, or application that submits queries. A submitter ID is not defined by Query Patroller; the submitter ID is the SQL authorization ID that DB2 assigns when the submitter connects to the instance (a case-sensitive ID, usually in uppercase). Query Patroller uses submitter IDs to differentiate the users or groups of users who submit queries through Query Patroller.

Query Patroller's ability to identify individual submitters varies depending on the environment in which you are using Query Patroller:

- In a two-tier environment in which users submit queries directly to Query Patroller, the DB2 ID associated with each submitter is passed along to Query Patroller.
- In a three-tier environment, Query Patroller treats all queries that are submitted through the application as if they come from the same submitter. In such an environment, the application server in the second tier uses the same ID to connect to the database regardless of which user is submitting a query.

Query Patroller submitter profiles

A Query Patroller submitter profile is a set of characteristics that define:

- Whether Query Patroller should intercept queries from a submitter
- If the submitter's queries are intercepted, what resource limits are applied to those queries

- What priority level the submitter's queries have in a queue
- The submitter's charge-back account code (to be used for cost tracking purposes)

You can create submitter profiles for individual users and for groups, but this is not required. During the Query Patroller installation process, a submitter profile called PUBLIC is created. By default, all submitters use this profile unless they also belong to a more restrictive profile. .

If you want to assign different characteristics to some submitters than you assign to others, the submitters must use different submitter profiles. For instance, if you want to track queries submitted by the Marketing department to determine the resources that they use, you will need to create a group submitter profile for submitters in that department. This means that you must use a DB2 group ID for the Marketing submitters.

Submitter resource limits

You can set submitter resource limits to ensure that no individual submitter or group of submitters uses too many system resources. You can set limits on the number of queries that a submitter can run simultaneously and on the maximum cost (in timerons) that any single query from a submitter can incur.

To limit the amount of disk space taken up by large query results, you can limit the number of result rows that can be stored in a result table for a single query from a submitter.

Interception and management of queries from a particular submitter profile

You can specify in the submitter profile if queries from the associated submitter should be allowed to bypass Query Patroller. Unintercepted queries are not managed, nor are they used for the collection of data for historical analysis.

You can also specify in the submitter profile that if a query from the submitter is below a certain size (in timerons) Query Patroller will not manage that query. Specify this using the minimum cost to manage value for the submitter (MIN_COST_TO_MANAGE).

Submitter queue priority

If you want to ensure that queries from certain users have priority over other queries when there is a queue, you can assign a higher queue priority to those users. The submitter queue priority is a numeric value between 0 and 999 that defines the priority of queries submitted by the associated submitter in a queue of queries. By default, all submitters have a queue priority of 500. You may, for example, want to assign a queue priority of 700 to submitters that perform queries needed for monetary transactions to ensure that these transactions are handled quickly.

For queue priority settings to work effectively, you should plan in advance the queue priority values you want to assign to different users.

Submitter charge-back accounts

A charge-back account is an alphanumeric account code to be used for cost tracking purposes. You can use the charge-back parameter to sort submitters into

logical groupings to track usage costs. To obtain historical analysis data for charge-back accounts, view the `TRACK_QUERY_INFO` table.

Active and suspended submitter profiles

Users must have an active submitter profile in order to submit queries. With DB2 Version 8, every DB2 user ID belongs to the `PUBLIC` group, so if the `PUBLIC` submitter profile is active then by default Query Patroller will allow every DB2 user to submit queries. To prevent an individual user from submitting queries without suspending the `PUBLIC` submitter profile, you can create an individual profile for the submitter and suspend the individual submitter profile. To create or suspend a submitter profile, you must be an administrator or an operator whose profile has the `USER ADMINISTRATION` privilege with edit authority.

Which submitter profile Query Patroller uses

The query submission preferences that can be set for a submitter include specifying the group submitter profile Query Patroller uses for that submitter. The default query submission preference for a submitter profile is `PUBLIC`. You can override the default query submission preference for a submitter profile by changing the query submission preferences of each submitter.

Query Patroller identifies which submitter profile to use when processing a query based on the following criteria, in order of importance:

1. Query Patroller chooses the submitter's own profile, if one exists. This submitter profile type is `USER`.
2. If the submitter does not have a `USER` submitter profile, but does have query submission preferences, Query Patroller chooses the active group profile that matches the group profile specified in the query submission preferences.
3. If there is no group submitter profile specified in the query submission preferences, or if the specified group submitter profile is not active, Query Patroller chooses the active group profile that is most restrictive by looking at the following characteristics in the following order of importance:
 - a. Whether Query Patroller does or does not intercept queries (`BYPASS Y/N`) from the group
 - b. The minimum cost to manage value for the group (`MIN_COST_TO_MANAGE`)
 - c. The maximum cost of a query value for the group (`MAX_COST_ALLOWED`)
 - d. The maximum number of queries value for the group (`MAX_QUERIES_ALLOWED`)
 - e. The queue priority of the group
 - f. The maximum number of return rows value for the group (`MAX_RESULT_ROWS`)
4. If more than one of the active group submitter profiles have the same degree of restrictiveness, Query Patroller chooses a profile arbitrarily.
5. If there is no active group submitter profile found, Query Patroller chooses the `PUBLIC` profile,
6. If the `PUBLIC` submitter profile is not active, Query Patroller returns an SQL error that states that a submitter profile was not found.

For example, consider a sales manager who has three different IDs that she uses when submitting queries:

- She uses the *sales_dept* group ID when accessing sales data.
- She uses the *sales_transactions* group ID when performing sales transactions.
- She uses the *managers* group ID when accessing the records of employees in her department.

Each of these group IDs has a group submitter profile. The settings for each of these group profiles are defined as shown in Table 6.

Table 6. Profile settings

Settings	<i>sales_dept</i>	<i>sales_transactions</i>	<i>managers</i>
Intercept	Y	Y	Y
Minimum cost to manage value (in timerons)	10 000	10 000	15 000
Maximum cost of a query allowed value (in timerons)	700 000	250 000	1 000 000
Maximum number of queries value	20	30	20
Queue priority	500	700	500
Maximum number of return rows value	1 000,000	400 [®] 000	1 200 000

The sales manager does not have her own submitter profile (of type USER), and her submission preferences do not specify which group profile she wants to use to submit queries. Query Patroller identifies which submitter profile to use when processing the sales manager's queries by determining which submitter profile is most restrictive:

- Queries submitted by all three groups can be intercepted, so they are all equally restrictive.
- The minimum cost to manage a query specified for the *managers* group profile is 15 000 timerons; whereas the minimum cost to manage specified for both the *sales_dept* group profile and the *sales_transactions* group profile is 10 000 timerons. Since the *managers* group profile is less restrictive, it is not used.
- The value specified for maximum cost of a query allowed for the *sales_dept* group profile is 700 000 timerons, whereas the value is 250 000 timerons for the *sales_transactions* group profile. The value specified for the *sales_transactions* group profile is more restrictive.
- Query Patroller identifies the *sales_transactions* group profile as the submitter profile it should use when processing queries from the sales manager.

If the sales manager needs to submit larger queries, she must either request that the administrator create her a submitter profile (of type USER) or she must change her submission preferences so that they specify which group profile she wants to use to submit queries.

Configuring submitter profiles

Before you create submitter profiles for your database, you need to make decisions about the type of submitters that will be using the database and what type of constraints you want to place on the resources used by each submitter or group of submitters.

You must read the Submitter profile configuration topic before completing this task.

To configure a submitter profile, use the following method.

1. Decide what type of submitter distinctions can be made in your system.
2. Optional: Collect historical data to get a representative sample of database activity.
3. Optional: Use the Historical Analysis submitters reports to view the following information:
 - the size of queries submitted by different submitters
 - the number of queries submitted by a submitter or group of submitters within a given time frame
 - which submitters generate unusually large queries
4. Decide which submitters or groups of submitters should be intercepted by Query Patroller.
5. Decide which thresholds to set for submitters whose queries will be intercepted.
6. Determine the priority level for each submitter profile.
7. Create submitter profiles for your system.
8. Evaluate your system performance and modify your submitter profiles if necessary.

Creating submitter profiles for users and groups

You can create a Query Patroller submitter profile based on an existing submitter profile or you can create a submitter profile with new settings. In the submitter profile, you can: define the submitter's charge-back account code (to be used for cost tracking purposes), specify whether queries from the submitter should be intercepted, and if the submitter's queries can be intercepted, set resource limits for the submitter

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile has the USER ADMINISTRATION privilege with edit authority

You create one or more submitter profiles for the Query Patroller system using one of the following methods:

- Query Patroller Center method
- Command line method
- To create one or more submitter profiles using the Query Patroller Center method:
 1. Open the Create Submitter notebook:
 - To create a submitter profile based on an existing submitter profile, open the Create Submitter notebook as follows:
 - a. From the Query Patroller Center, expand the object tree under the **User Administration** folder to find the **Submitters** folder.
 - b. Click the **Submitters** folder. Any existing submitters are displayed in the pane on the right side of the window (the contents pane).
 - c. In the contents pane, right-click the submitter whose profile you want to use as the basis for the new submitter you are creating, and click **Create Like** in the pop-up menu. The Create Submitter notebook opens with certain fields prefilled.
 - To create a submitter profile with new settings, open the Create Submitter notebook:

- a. From the Query Patroller Center, expand the object tree under the **User Administration** folder to find the **Submitters** folder.
 - b. Right-click the **Submitters** folder, and click **Create** in the pop-up menu. The Create Submitter notebook opens.
2. On the General page:
- a. In the **Submitter user ID** field, type the name that will be associated with this profile. The name is case sensitive and must match the SQL authorization ID in the Database Connection Information. The SQL authorization ID is usually same as the login ID, but in uppercase. You can type multiple values separated by commas, to create multiple submitter profiles.
 - b. Use the **Profile type** field to select whether the submitter profile applies to a user or a group.
 - c. Optional: In the **Charge-back account** field, type an alphanumeric account code to be used for cost tracking purposes. You can use the charge-back parameter to sort submitters into logical groupings to track usage costs.
Examples:
 - If three submitters belong to the legal department, then you could type LEGAL in the **Charge-back account** field.
 - If an expense code of MK001 exists for the Marketing department, then you could type MK001 in the **Charge-back account** field whenever you add a submitter from the Marketing department.
 - d. If you want to temporarily suspend the new submitter's ability to submit queries, select the **Access suspended** check box.
3. On the Resources page:
- a. If you do not want Query Patroller to intercept queries from the submitter you are creating, select the **Do not intercept queries from this submitter** check box. Query Patroller will neither perform cost evaluation nor create a query in the historical analysis view. Query Patroller will not manage queries submitted by submitters with this profile.
 - b. In the **Minimum cost to manage** field, type a number greater than or equal to 0 and less than the **Maximum cost of a query** value. Queries costing less than this value will not be managed by Query Patroller.
 - c. In the **Maximum number of queries** field, type the maximum number of queries that a submitter is permitted to run simultaneously. Additional queries will be queued. To allow an unlimited number of queries to run simultaneously, leave the field blank.
 - d. In the **Maximum cost of a query** field, type a number for the maximum query cost. If the submitter submits a query with an estimated cost that is higher than this value, the query is held. To set the value to unlimited, leave the field blank.
 - e. In the **Maximum size of a result table** field, type a value for the maximum number of result rows to be stored in a result table. If you leave the field blank, DB2 will always make this submitter's result table large enough to accommodate the entire result set.
 - f. In the **Queue priority** field, type a value from 0 to 999 to represent the priority assigned to a query when a query is submitted. Higher numbers indicate higher priority.
4. Click **OK** to create the new submitter.

- To create one or more submitter profiles using the Command line method, issue the ADD SUBMITTER_PROFILE command.

Setting submitter resource limits

You can specify whether or not you want a submitter's queries to be intercepted by Query Patroller. If you specify that they will be intercepted, you can also limit the resources that are used by a submitter, so that the submitter cannot use too many of the system's resources.

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile has the USER ADMINISTRATION privilege with edit authority

You can limit submitter resources using one of the following methods:

- Query Patroller Center method
- Command line method
- To limit submitter resources using the Query Patroller Center method
 1. Open the Submitter Properties notebook:
 - a. From the Query Patroller Center, expand the object tree under the **User Administration** folder to find the **Submitters** folder.
 - b. Click the **Submitters** folder. Any existing submitters are displayed in the pane on the right side of the window (the contents pane).
 - c. Right-click the submitter that you want to modify in the contents pane, and click **Properties** in the pop-up menu. The Submitter Properties notebook opens.

Note: You can also open the Submitter Properties window by double-clicking the submitter profile that you want to change.
 2. Click the **Resources** tab.
 3. If you do not want Query Patroller to intercept queries from the submitter you are creating, select the **Do not intercept queries from this submitter** check box. Query Patroller will neither perform cost evaluation nor create a query in the historical analysis view. Query Patroller will not manage queries submitted by submitters with this profile.
 4. In the **Minimum cost to manage** field, type a number greater than or equal to 0 and less than the **Maximum cost of a query** value. Queries costing less than this value will not be managed by Query Patroller.
 5. In the **Maximum number of queries** field, type the maximum number of queries that this submitter is permitted to run simultaneously. Additional queries will be queued. To allow an unlimited number of queries, to run simultaneously, leave the field blank.
 6. In the **Maximum cost of a query** field, type a number for the maximum query cost. If the submitter submits a query with an estimated cost that is higher than this value, the query is held. To set the value to unlimited, leave the field blank.
 7. In the **Maximum size of a result table** field, type a value for the maximum number of result rows to be stored in a result table. If you leave the field blank, DB2 will always make this submitter's result table large enough to accommodate the entire result set.

8. In the **Queue priority** field, type a value from 0 to 999 to represent the priority assigned to a query when a query is submitted. Higher numbers indicate higher priority.
 9. Click **OK** to update the submitter profile.
- To limit submitter resources using the Command line method:
 1. Issue the `UPDATE SUBMITTER_PROFILE` command using the following parameters:

Suspending or restoring submitter privileges for users and groups

You can temporarily suspend a submitter's ability to submit queries, then restore these privileges at a later time.

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile has the `USER ADMINISTRATION` privilege with edit authority

You can suspend or restore a submitter's privileges using one of the following methods:

- Query Patroller Center method
- Command line method
- To suspend or restore a submitter's privileges using the Query Patroller Center method
 1. Open the Submitter Properties notebook:
 - a. From the Query Patroller Center, expand the object tree under the **User Administration** folder to find the **Submitters** folder.
 - b. Click the **Submitters** folder. Any existing submitters are displayed in the pane on the right side of the window (the contents pane).
 - c. Right-click the submitter that you want to modify in the contents pane, and click **Properties** in the pop-up menu. The Submitter Properties notebook opens.

Note: You can also open the Submitter Properties window by double-clicking the submitter profile that you want to change.

2. On the General page, select the **Access suspended** check box to suspend the submitter's ability to submit queries, or clear it to restore the submitter's ability to submit queries.
 3. Click **OK** to close the Submitter Properties notebook.
- To suspend or restore a submitter's privileges using the Command line method:
 1. Issue the `UPDATE SUBMITTER_PROFILE` command using the following parameters:

Query Patroller query submission preferences

Query submission preferences are used to specify

- The submitter profile the submitter uses when submitting queries if the submitter does not have a user submitter profile, and has two or more group submitter profiles
- The location where Query Patroller should send the results of the submitter's queries

- Who can view the submitter's result tables
- What Query Patroller should do if the submitter's result tables are too large
- The e-mail address to be used for sending notifications to the submitter

The default query submission preferences are assigned to a submitter profile called PUBLIC. If an individual submitter requires submission preferences that are different from the default submission preferences, you need to create new submission preferences for that submitter.

Setting individual query submission preferences

The default query submission preferences are called PUBLIC. If you do not have your own submission preferences, by default you will use the settings specified in the PUBLIC submission preferences. If you do not want to use the values specified in the PUBLIC submission preferences, you can set different submission preferences for yourself.

To set your query submission preferences, use one of the following methods.

- From the Query Patroller Center:
 1. Open the Query Submission Preferences window.
 2. If you belong to multiple group profiles, use the **Submitter profile to use** field to select which group profile you want to submit your queries under. Select the group submitter profile that you want to submit your queries with. The **Select automatically** entry applies the profile with the lowest resource limits. If you lack a submitter profile whose type is user, yet belong to multiple group submitter profiles, the group submitter profiles are not listed in this field, but you can type the group submitter profile name.
 3. Specify what will happen after a query is submitted:
 - To specify that the application that submitted the query will wait for the result set to return while Query Patroller manages the query, select the **Wait until the results are returned** radio button. This option is the default. When this option is selected, the application that submitted the query can become unresponsive until the result set is returned.
 - To specify that the result set will be stored in a DB2 table and the application that submitted the query will become free for further processing, select the **Release the application and retrieve the results from a result table** radio button.
 4. Specify the access levels you want to use for your result tables:
 - To specify that only the submitter can access the result tables, select the **Restrict access to the submitter** radio button.
 - To make the result table accessible to specific DB2 users select the **Grant access to other users or groups** radio button. Type the names of specific users or groups in the associated field, separated by commas. The DB2 users listed must have access to the database where the query was submitted.
 5. Specify how to handle result sets that are longer than the maximum allowed:
 - To specify that the incomplete results should not be stored in a result table, select the **Do not return any results** radio button.
 - To specify that a truncated set of results should be stored in a result table, select the **Return a truncated set of results** radio button.
 6. In the **E-mail address** field, type the e-mail address to be used for sending notifications to the submitter. Notifications can be sent to the address or

addresses specified when a query completes or if a query encounters an error, but only in cases where a result table is created.

Note: A result table is created for submitters when:

- The **Release the application** option is specified in the Query Submission Preferences window.
 - The **Wait until results are returned** option is specified in the Query Submission Preferences window, yet the query was placed in a held state before completing.
 - The status of a query is changed to **Run query in the background**.
7. Click **OK** to set preferences for query submission.
- From the Query Patroller command line:
 1. To create new query submission preferences, issue the `ADD SUBMISSION_PREFERENCES` command.
 2. To change existing query submission preferences, issue the `UPDATE SUBMISSION_PREFERENCES` command.

Setting query submission preferences for another submitter

The default query submission preferences are called PUBLIC. Submitters who do not have their own submission preferences, use the settings specified in the PUBLIC submission preferences. If you do not want a submitter to use the values specified in the PUBLIC submission preferences, you can set different submission preferences for the submitter.

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile has the USER ADMINISTRATION privilege with edit authority

You can update preferences for query submission using one of the following methods:

- Query Patroller Center method
- Command line method
- To update preferences for query submission using the Query Patroller Center method:
 1. Open the Query Submission Preferences window.
 - To create new query submission preferences based on existing query submission preferences:
 - a. From the Query Patroller Center, expand the object tree under the **User Administration** folder to find the **Query Submission Preferences** folder.
 - b. Click the **Query Submission Preferences** folder. Any submitters with defined submission preferences are displayed in the pane on the right side of the window (the contents pane).
 - c. Right-click the submission preferences you want to use as the basis for the new submission preferences you are creating, and click **Create Like** in the pop-up menu. The Query Submission Preferences window opens, with certain fields pre-filled.
 - To create new query submission preferences with new settings:

- a. From the Query Patroller Center, expand the object tree under the **User Administration** folder to find the **Query Submission Preferences** folder.
- b. Right-click the **Query Submission Preferences** folder and click **Create** in the pop-up menu. The Query Submission Preferences window opens.
- To modify existing query submission preferences:
 - a. From the Query Patroller Center, expand the object tree under the **User Administration** folder to find the **Query Submission Preferences** folder.
 - b. Click the **Query Submission Preferences** folder. Any existing query submission preferences are displayed in the pane on the right side of the window (the contents pane).
 - c. Right-click the query submission preferences that you want to change in the contents pane, and click **Properties** in the pop-up menu. The Query Submission Preferences window opens.

Note: You can also open the Query Submission Preferences window by double-clicking the query submission preferences that you want to view or change.

2. If you are creating new query submission preferences, in the **Submitter** field enter the name of the submitter or the user ID that was used to start the application that submits queries.
3. If the submitter or end user belongs to multiple group profiles, use the **Submitter profile to use** field to select which group profile the submitter will use when submitting queries. If the submitter or end user has a submitter profile whose type is user, this field shows the submitter's user ID, and you cannot change this value.

If you lack a submitter profile whose type is user, yet belong to multiple group submitter profiles, this field lists the group submitter profiles. Select the group submitter profile that you want to submit your queries with. The Select automatically entry applies the profile with the lowest resource limits.

4. Specify what will happen after a query is submitted:
 - To specify that the application that submitted the query will wait for the result set to return while Query Patroller manages the query, select the **Wait until the results are returned** radio button. This option is the default.
When this option is selected, the application that submitted the query can become unresponsive until the result set is returned.
 - To specify that the result set will be stored in a DB2 table and the application that submitted the query will become free for further processing, select the **Release the application and retrieve the results from a result table** radio button.
5. Specify the access levels you want to use for your result tables:
 - To specify that only the submitter can access the result tables, select the **Restrict access to the submitter** radio button.
 - To make the result table accessible to specific DB2 users select the **Grant access to other users or groups** radio button. Type the names of specific users or groups in the associated field, separated by commas. The DB2 users listed must have access to the database where the query was submitted.
6. Specify how to handle result sets that are longer than the maximum allowed:

- To specify that a truncated set of results should be stored in a result table, select the **Return a truncated set of results** radio button.
 - To specify that the incomplete results should not be stored in a result table, select the **Do not return any results** radio button.
7. In the **E-mail address** field, type the e-mail address to be used for sending notifications to the submitter. You can type multiple values separated by commas, to send e-mail to more than one address. Notifications are sent to the address or addresses specified when a query completes or if a query encounters an error, but only in cases where a result table is created.

Note: A result table is created for submitters when:

- The **Release the application** option is specified in the Query Submission Preferences window.
 - The **Wait until results are returned** option is specified in the Query Submission Preferences window, yet the query was placed in a held state before completing.
 - The status of a query is changed to **Run query in the background**.
8. Click **OK** to set preferences for query submission.
- To create new query submission preferences using the Command line method:
 1. Issue the `ADD SUBMISSION_PREFERENCES` command.
 - To change existing query submission preferences using the Command line method:
 1. Issue the `UPDATE SUBMISSION_PREFERENCES` command..

Chapter 13. Query interception and management in Query Patroller

After queries are submitted, Query Patroller performs two stages of evaluation on them to determine the level of interaction Query Patroller will have with each query.

1. Query Patroller first evaluates the queries to determine if they meet specific criteria that allow them to bypass Query Patroller. Unintercepted queries are not managed, nor are they used for the collection of data for historical analysis.
2. Query Patroller then evaluates the intercepted queries to determine how each query will be handled. Depending on Query Patroller's evaluation of a query, one of the following occurs:
 - The query is both managed and historical data is collected for it.
 - The query is not managed, but historical data is collected for it.
 - The query is neither managed, nor is historical data collected for it.

Note: In order for Query Patroller to perform these evaluations and subsequent query interception and management, the *dyn_query_mgmt* database configuration parameter must be set to ENABLE.

Query interception

Query Patroller intercepts a query if the query meets the following criteria:

- The query is from an application whose queries you have specified you want to intercept in the Query Patroller system properties.
- The query is from a submitter whose submitter profile indicates that Query Patroller should intercept queries from the submitter.

If Query Patroller intercepts the query, it then evaluates the query to see if the query should be managed. Query Patroller can be configured to collect data for historical analysis on intercepted queries even if they do not meet the criteria for queries that should be managed.

Query management

Query Patroller manages an intercepted query based on the properties that have been set for the submitter of the query and on the Query Patroller system settings. Query Patroller also uses the query's estimated cost, which is assessed by the DB2 query optimizer. The management functions that Query Patroller performs on a query include prioritizing it (if applicable), assigning it to a query class (if applicable), and then either running, queuing, holding, or rejecting the query. If the cost of the query is smaller than the value specified in the submitter's profile for the minimum cost to manage a query (MIN_COST_TO_MANAGE), Query Patroller will not manage the query. Table 7 shows the differences in how Query Patroller treats managed and unmanaged queries.

Table 7. Query Patroller treatment of managed versus unmanaged queries

Query Patroller action	Managed queries	Unmanaged queries
Saves the query in the Query Patroller Center's Managed Queries folder	Yes	No

Table 7. Query Patroller treatment of managed versus unmanaged queries (continued)

Query Patroller action	Managed queries	Unmanaged queries
Saves the query in the Query Patroller Center's Historical Analysis folder	Yes	Optional
Prioritizes the query	Yes	No
Assigns the query to the appropriate query class	Yes	No
Runs, queues, holds, or rejects the query based on various thresholds	Yes	No
Returns the result set to the submitter's application, or creates a result table	Yes	No
Sends e-mail notification to submitter if a result table is created	Optional	No

Unintercepted queries

Query Patroller can be configured not to intercept queries from specific applications, or from certain submitters. Unintercepted queries are not managed by Query Patroller, so you cannot collect historical analysis data on them.

Bypassing query interception using Query Patroller variables

If you use Query Patroller submitter profiles to bypass queries, each bypassed query still incurs a Query Patroller overhead, because the Query Patroller server still needs to examine each query before deciding if the query can bypass Query Patroller management. When a large number of small queries are involved (as compared to a small number of large queries), the cumulative overhead can become significant, and affect overall database performance materially.

Alternatively, you can use one of three new Query Patroller registry variables that can bypass queries without any Query Patroller server involvement. Because the Query Patroller does not need to assess queries bypassed using the registry variables, no database performance problems occur when a lot of small queries pass through the system. The new variables are DB2_QP_BYPASS_APPLICATIONS, DB2_QP_BYPASS_USERS, and DB2_QP_BYPASS_COST.

Managed query status

While a query is being managed by Query Patroller, a query can be in one of several different states. Query status information can be obtained in several ways: from the Query Patroller Center Managed Queries folder, from the **Query status** field of the Managed Query Properties notebook, and by issuing the GET QUERY command.

If a query is managed by Query Patroller, it will pass through the following states during normal processing:

Initial The query has been intercepted by the Query Patroller server. (This status is never displayed, because the query moves quickly to a subsequent status.)

Running

The query is in progress, and has been passed to DB2 for execution.

Done The query completed successfully.

Note: Although the query itself completed without error, the application may receive an error if the completion was caused by an external event, such as a DB2 force application.

In some cases, queries pass through other states during processing:

Held The cost of the query exceeds the submitter's threshold. A held query can be released manually or it can be released automatically by a scheduled job. Releasing a held query will place it in *released* state.

Released

The query was held, but has been released by an administrator or an operator whose profile has the MONITORING privilege with edit authority, or it was released automatically by a scheduled job. A released query is processed by the Query Controller and placed in *running* state or *queued* state, depending on the current system workload.

Queued

The query is waiting to run. A query can be queued if one or more of the following thresholds is exceeded:

- Maximum number of queries value for the system (MAX_TOTAL_QUERIES)
- Maximum number of queries value for the submitter (MAX_QUERIES_ALLOWED)
- Maximum workload cost value for the system (MAX_TOTAL_COST)
- Maximum number of queries value for the query class in which the query runs (MAX_QUERIES)

The query will run when the situation that caused it to be queued has changed. For example, a query that was queued because the maximum number of queries for that query class has been exceeded will run when the number of queries belonging to that query class falls below the maximum for the query class. However, if a query caused more than one threshold to be exceeded, then even after the first threshold is no longer exceeded, the second threshold may still be exceeded. The query will run when it no longer causes any of these thresholds to be exceeded.

Canceled

The query was canceled, through either the Query Patroller Center or the Query Patroller command line, by the administrator, submitter, or an operator whose profile has the MONITORING privilege with edit authority. Only *running*, *held*, *released*, and *queued* queries can be *canceled*.

Aborted

The query was terminated by DB2 because of an error.

Rejected

The query was prevented from running.

Unknown

The status of the query cannot be determined. A query with status *unknown* is no longer *running*, but Query Patroller cannot determine whether the query completed or failed.

Changing the status of queries using Query Patroller

Depending on your authority level, you may change the status of queries managed by Query Patroller in several ways: cancel a query, release a held query, or run a query in the background

- To cancel a query, you must meet one of the following requirements:
 - Have DBADM authority
 - Be an operator whose profile has the MONITORING privilege with edit authority
 - Be the submitter of the query
- To release a held query, you must meet one of the following requirements:
 - Have DBADM authority
 - Be an operator whose profile has the MONITORING privilege with edit authority
- To run a query in the background, you must:
 - Be the submitter of the query

Canceling queries

Cancel a query if you realize after it has been submitted that it contains an error, or that its cost is too high. For example, you might receive notification that your query is being held because its cost exceeds the maximum amount of system resource that is allowed for any one of your queries. Canceling a query places it in the canceled state.

Releasing queries from the held state

Release a held query if you decide that a particular query should run, even though the query exceeds the submitter's maximum query cost. Releasing a held query places the query in the running state or the queued state, depending on the current system workload.

Running queries in the background

Run a query in the background if your query submission preferences specify that you will wait until the results of a query are returned, but you want to use your client application while that particular query runs. Running a query in the background places the query in the running state or the queued state, depending on the current system workload.

You can change the status of a query for the Query Patroller system using one of the following methods:

- Query Patroller Center method
- Command line method
- To change the status of a query using the Query Patroller Center method:
 1. Open the Change Query Status window.
 - a. From the Query Patroller Center, expand the object tree under the **Monitoring** folder to find the **Managed Queries** folder.
 - b. Click the **Managed Queries** folder. The managed queries are displayed in the pane on the right side of the window (the contents pane).
 - c. In the contents pane, right-click the query whose status you want to change, and click **Change Status** in the pop-up menu. The Change Query Status window opens.
 2. To cancel the query, click **Cancel query**.
 3. To run the query, click **Release query from held state**.

4. To regain control of the submitting application, click **Run query in the background**. Query Patroller stops the execution of the query and resubmits it. The results of the query will be returned to a result table.
 5. Click **OK** to change the status of the query as you have indicated and close the Change Query Status window.
 - 6.
- To change the status of a query using the Command line method use the following commands:
 1. To cancel the query, issue the CANCEL QUERY command..
 2. To run the query, issue the RUN HELD_QUERY command.
 3. To rerun the query in the background, issue the RUN IN BACKGROUND QUERY command.

Query Patroller variables

There are several DB2 registry variables that are used by Query Patroller to control query interception. Set these variables on the server using the db2set command.

DB2_QP_BYPASS_APPLICATIONS

- Operating system: All
- Default: NULL, Values: One or more applications, separated by a colon. Application names are case-sensitive.
- This variable causes Query Patroller to bypass all queries from the application or applications you specify.

Note: The qprunquery application should not be specified in this registry variable.

DB2_QP_BYPASS_USERS

- Operating system: All
- Default: NULL, Values: One or more user authorization IDs, separated by a colon. The user authorization ID should be specified in uppercase.
- This variable causes Query Patroller to bypass all queries submitted by the authorization ID or authorization IDs you specify.

Note: The instance owner should not be specified in this registry variable on the DB2 database server where Query Patroller is working.

DB2_QP_BYPASS_COST

- Operating system: All
- Default: NULL, Values: Cost, specified in timerons.
- This variable causes Query Patroller to bypass all queries from all users and all applications that have estimated cost lower than the specified number.

Viewing managed query details using Query Patroller

By viewing the properties of a query that has been managed by Query Patroller, you can see details such as information about the submitter of the query, the processing time, and the result table.

You must meet one of the following requirements:

- Have DBADM authority

- Be an operator whose profile has the MONITORING privilege with edit or view authority
- Be the submitter of the query

You can view the properties of a query using one of the following methods:

- Query Patroller Center method
- Command line method
- To view the properties of a query using the Query Patroller Center method
 1. Open the Managed Query Properties notebook:
 - a. From the Query Patroller Center, expand the object tree under the **Monitoring** folder to find the **Managed Queries** folder.
 - b. Click the **Managed Queries** folder. Any managed queries are displayed in the pane on the right side of the window (the contents pane).
 - c. In the contents pane, right-click the query that you want to work with, and click **Properties** in the pop-up menu. The Managed Query Properties notebook opens.

Note: You can also open the Managed Query Properties window by double-clicking the query that you want to work with.

2. To view general information about the query, click the **General** tab.
 - a. To view the SQL of the query in a new window, click **View SQL in Separate Window**.
 - b. To view the profile of the submitter who submitted the query, click **Submitter Properties**. You must have DBADM authority or be an operator whose profile has the USER ADMINISTRATION privilege with edit or view authority to open the Submitter Properties window.
 - c. If you want to change the status of the query, for example to cancel it, click **Change Status**.
 - d. Click the **Show Access Plan** if you want to launch Visual Explain to see more information about the query. This feature is available only if Visual Explain is installed on your system; if it is not installed, the button is disabled.
3. To view query execution information and result table details, click the **Results** tab.
 - a. To display the result table for the query, click **Show Results**.
 - b. To save the result table for the query, click **Save Results**.

Note: When saving the results of a query, a BLOB column value will not be saved. Instead it will be replaced with the keyword, "BLOB". A CLOB value will be truncated if it is greater than 32K in size.

- c. To delete the result table for the query, click **Drop Result Table**.
4. To view query event timestamps and processing durations, click the **Time** tab.
 5. To view authorization IDs, application information, and user information, click the **Other** tab.
 6. Click **Close** to close the Managed Query Properties notebook.
- To view the properties of a query using the Command line method:
 1. Issue the following commands:
 - To view the details of a query, issue the GET QUERY command.
 - To save the results of a query, issue the FILE RESULT command.

- To view the results of a query, issue the SHOW RESULT command.

Viewing the SQL of managed queries using Query Patroller

Use the SQL Statement window to view the SQL of a managed query. From this window, you can search for strings and SQL keywords in the statement, save the SQL statement to a file, print the SQL statement, and copy the SQL statement. You can paste the query into SQL Explain in order to find the access plan that the DB2 optimizer used for the SQL statement.

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile has the MONITORING privilege with edit or view authority
- Be the submitter of the query

About query interception and management in Query Patroller

To view the SQL of a query, use the Query Patroller Center method.

1. Open the SQL Statement window:
 - a. Open the Managed Query Properties notebook.
 - b. Click the **General** tab.
 - c. Click **View SQL in Separate Window**. The SQL Statement window opens.
2. Click **Copy Text** to copy the SQL statement to the clipboard.
3. Click **Find** to find a specific text string in the SQL statement. This action opens a Find window.
4. Click **Save As** to open a standard Save As window. From there you can specify a file and location to save the SQL statement.
5. Click **Print** to open a standard Print window. From there you can select a printer and choose to print the SQL statement.
6. Click **Close** to close the SQL Statement window.

Viewing the query submitter profiles

View the profile of a query submitter to find out who submitted a query, what charge-back account that submitter has, whether the submitter's database access is active or suspended, and what the resource limits of that submitter are.

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile has the USER ADMINISTRATION privilege with view or edit authority

To view a query submitter profile, use any of the following methods.

- Query Patroller Center method
- Command line method
- To view the query submitter profiles using the Query Patroller Center method, do the following:
 1. Open the Submitter Properties notebook.
 2. To view information about the submitter's profile type, user ID, and charge-back account, click the **General** tab.

3. To view the submitter's resource limits and other information, click the **Resources** tab.
 4. Click **OK** to close the Submitter Properties notebook.
This view of the Submitter Properties notebook is a read-only view of a submitter's profile.
- To view the query submitter profiles using the command line method, do the following:
 1. Issue the **GET SUBMITTER_PROFILE** command.

Filtering managed queries using Query Patroller

When you are using the Query Patroller Center to look at managed queries, there might be a significant number of rows in the Managed Queries report. Use the Filter Managed Queries notebook to display only the queries that meet conditions that you specify.

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile has the MONITORING privilege with view or edit authority
- Be the submitter of the query

To filter managed queries using the Query Patroller Center method:

1. Open the Filter Managed Queries notebook
 - a. From the Query Patroller Center, expand the object tree under the **Monitoring** folder until you find the **Managed Queries** folder.
 - b. Right-click the **Managed Queries** folder and click **Filter** in the pop-up menu. The Filter Managed Queries notebook opens.

This notebook will also open automatically when the **Managed Queries** folder is selected, if the number of objects in the folder exceeds the object count specified in the Filter Managed Queries notebook, and if the option to automatically display the filter is selected.

2. Optional: On the **Locate** tab, select an operator for the ID comparison, and enter a value to be compared to. The value must be numeric and must correspond to a query ID.
3. Select an operator for the Status comparison, and enter a value to be compared to. Possible values for query status are:
 - H - Held
 - Q - Queued
 - J - Rejected
 - R - Running
 - L - Released
 - A - Aborted
 - C - Canceled
 - D - Done
 - U - Unknown
4. Optional: Select an operator for the Submitter ID comparison, and enter a value to be compared to. The value must be a character string and must correspond to a submitter ID.

5. Optional: Select an operator for the Created comparison, and enter a value to be compared to. The value must be a timestamp representing the time the query was issued, for example 2003-07-29-00.00.00.
6. Optional: Select an operator for the Completed comparison, and enter a value to be compared to. The value must be a timestamp representing the time that processing was completed for that query, for example 2003-07-29-00.00.00.
7. Optional: Select an operator for the Query Class comparison, and enter a value to be compared to. The value must be numeric and must correspond to a query class ID.
8. To indicate that the filter should show only queries that match all the comparisons you have entered, click the **Meet all conditions** radio button. To indicate that the filter should show queries that match one or more of the comparisons, click the **Meet any conditions** radio button.
9. Optional: To view the WHERE clause that is generated by the comparisons you have entered, click the **Details** tab.
10. Optional: The Filter Managed Queries notebook will open automatically when the Managed Queries folder is selected and the number of objects listed exceeds the value specified in the **Object count** field, if you have selected the **Automatically display filter when object count is exceeded** check box. To adjust the number of objects that will cause the Filter Managed Queries notebook to open automatically, enter a new value in the **Object count** field.
11. To enable the filter, select the **Enable filter** check box, then click **OK** to close the Filter Managed Queries notebook.
12. Optional: If you do not want to enable the filter, clear the **Enable filter** check box, then click **OK** to close the Filter Managed Queries notebook. The values you have entered in the Filter Managed Queries notebook will be retained, and will appear when you reopen the notebook.

The filter that you define will affect the number of queries that you see in the report in the contents pane.

Result tables and result sets in Query Patroller

Each time Query Patroller executes a query, the rows returned form a result set. Query Patroller returns result sets to a result destination, which can be either the application through which the query was submitted or a result table.

A result table is a table that is created by Query Patroller to store the result set. A result table is created when:

- A query is held for any reason, then released and run to completion
- The submitter specifies in his submission preferences that a result table should be created
- A query is run in the background

When Query Patroller is installed, you must specify the table space where result tables will be stored. The name of the table space is stored in the Query Patroller system properties. You can later change or remove the table space name that is specified. If you remove the table space name specified in the Query Patroller system properties, Query Patroller will use a default table space.

The size of result tables can be limited for each submitter in the submitter's properties by restricting the number of rows that can be returned to the result table. Result tables consume storage space, so they should be removed periodically. You can remove result tables manually or you can schedule a purge job.

Query submitters can configure their submission preferences to handle result tables in certain ways:

- They can allow other Query Patroller users to view their result tables
- They can specify whether they want Query Patroller to return no results or truncated results if a query's result set is longer than the maximum allowed.

Note: If no results are returned, the query is considered to have failed (the query status is Aborted). If the query results are truncated, the query is considered to have completed successfully (the query status is Done; the result status is Truncated).

Viewing result tables using Query Patroller

You may want to view result tables if you suspect the data you are interested in has not changed significantly since you last ran a query, if you want to view the results of a query that was run in the past (resubmitting the query may return different results), or if a query has been run by another Query Patroller submitter, and you have been granted access to that submitter's result tables.

You must meet one of the following requirements:

- Have DBADM authority
- Be the submitter of the query
- Be granted access to the result table by the submitter of the query

A result table is created when:

- The query submitter's submission preferences specify that the application should be released after a query is submitted
- A query was held and then run
- A query was run in the background

Note: You cannot use the Query Patroller Center to view another submitter's results, even if you have been granted access to them. You must use the command line.

You can view a result table for a query using one of the following methods:

- Query Patroller Center method
- Command line method
- To view a result table for a query using the Query Patroller Center method
 1. Open the Show Results window:
 - a. Open the Managed Query Properties notebook for the query that you want to see the results for.
 - b. Click the **Results** tab.
 - c. Click **Show Results**. The Show Results window opens. The Show Results window allows you to view a result table 50 rows at a time.
 2. Click **Previous** or **Next** to move through the rows of the result table.
 3. Click **Close** to close the Show Results window.
- To view a result table for a query using the Command line method:
 1. Issue the SHOW RESULT command.

Running held queries at a scheduled time

You can set held queries to run at a scheduled time and specify the maximum time that held queries can run. This allows you to schedule held queries to run during off-peak hours when there are fewer demands on your system.

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile has the CONFIGURATION privilege with edit authority

To run held queries at a scheduled time, use one of the following methods.

- Query Patroller Center method
- Command line method
- To run held queries at a scheduled time using the Query Patroller Center method:
 1. Open the Query Patroller System Properties notebook:
 - a. From the Query Patroller Center, expand the object tree under the **Configuration** folder to find the **Query Patroller System** folder.
 - b. Click the **Query Patroller System** folder. A system property record is displayed in the pane on the right side of the window (the contents pane).
 - c. Right-click the system property record, and click **Properties** in the pop-up menu. The Query Patroller System Properties notebook opens.

Note: You can also open the Query Patroller System Properties notebook by double-clicking the entry in the contents pane when the Query Patroller System folder is displayed.

2. Click the **Held Queries** tab.
3. To set a time when held queries will be run:
 - a. Select the **Run held queries at a scheduled time** check box.
 - b. Click **Schedule Start Time** to open a window where you can manually schedule the time at which the queries will run. See “Scheduling the start time for running held queries” on page 88.

Note: There is no default schedule for Run held queries. If you do not create a schedule, no held queries will be run. If you create a schedule that has an end date, after the end date passes, no more held queries will be run.

4. In the **How long to run held queries** field, specify the length of time over which held queries will be run. After the specified amount of time passes, any held queries that have not been run will be held until the next start time. Any queries that are currently running are allowed to run to completion.
 5. Click **OK** to run held queries at a scheduled time.
- To run held queries at a scheduled time using the Command line method:
 1. Issue the UPDATE QP_SYSTEM command using the following parameters:
 - RUN_HELD_QUERIES
 - RUN_HELD_DURATION

Note: The UPDATE QP_SYSTEM command enables you to run held queries only if you have already manually set the schedule for running held queries through the Query Patroller Center. You cannot manually set schedules using the command line.

Scheduling the start time for running held queries

You can schedule held queries to run at a specific time, on a recurring basis. For instance, you can schedule held queries to run every Sunday at 1:00 in the morning.

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile has the CONFIGURATION privilege with edit authority

Note: If you are an operator, but you do not have edit authority, you can still view the schedule, but you cannot change the schedule.

Note: You can manually set schedules for running held queries only through the Query Patroller Center. You cannot manually set schedules using the command line.

To schedule running held queries using the Query Patroller Center method

1. Open the Schedule Start Time for Running Held Queries window:
 - a. Open the Query Patroller System Properties notebook.
 - b. Click the **Held Queries** tab.
 - c. Select the **Run held queries at a scheduled time** check box.
 - d. Click the **Schedule Start Time** push button. The Schedule Start Time for Running Held Queries window opens.
2. To specify how often you want to run held queries, use the **Interval** and **Frequency** fields. Depending on the interval you select, you may also need to use the **Details** field.
3. To specify when to begin running held queries, use the **Date** and **Time** fields.
4. Specify an end date for running held queries. If there is no end date, select the **No end date** check box. Otherwise, to specify an end date, use the **Date** and **Time** fields.

Note: The end date specified does not set the duration of a single task. It sets the end date for the recurring tasks. For instance, if you want to set held queries to run every Saturday for a duration of four hours, for the next seven weeks, use the **Date** and **Time** fields to specify an end date that is seven weeks from now.

5. Optional: To prevent the task from running, select the **Suspend schedule** check box.
6. Click the **Add** button to add the task to the **Schedule list**.
7. Click **OK** to return to the Query Patroller System Properties notebook, where you can set the duration of time that held queries will run.

Cost estimation in Query Patroller

Query Patroller evaluates each query that it intercepts based on the estimated size, or cost, of the query. The estimated cost of a query is the estimated total system resources that will be used in the execution of the query. A number of key Query Patroller thresholds are set based on the estimated cost of queries. On Query Patroller graphical user interfaces, this estimated cost is simply referred to as *cost*, for example, "Maximum cost of a query."

Query Patroller relies on DB2 for information on the estimated cost of queries. For any query submitted to DB2, there might be many different methods, or access plans, for retrieving the data from all of the required tables. An access plan is the data access strategy, which includes index usage, sort methods, locking semantics, and join methods.

The DB2 query optimizer analyzes the different access plans for a query to determine which of these plans is the most efficient in terms of estimated resource cost. In calculating the estimated cost of an access plan, the query optimizer considers both CPU cost (the number of instructions) and I/O (the number of seeks and page transfers). You can capture the details of the access plan that the query optimizer has selected to execute a query. These details are stored in explain tables, which can be queried or displayed using the DB2 explain facility.

The DB2 unit of measure for cost is the timeron. A timeron does not directly equate to any actual CPU execution time, but gives a relative measure of the estimated resources required to execute a particular query according to the selected access plan. Although the estimated cost of a query in timerons does not predict the execution time of that query, it can provide a basis for comparing two queries in the same environment under the same workload. For example, given the same external factors such as hardware capacity and concurrent workload, a very costly query involving complex joins and large table scans can be expected to have a longer execution time than a cheaper, simpler query.

DB2 passes on to Query Patroller the estimated cost of executing a query. Query Patroller uses this information to determine whether or not the query exceeds certain cost thresholds that have been defined for the system.

Note: To calculate the estimated cost of a query, the query optimizer relies on current table statistics. For this reason, it is recommended that you run the RUNSTATS command periodically to give the query optimizer the most accurate information on which to base its calculations.

Definition of your query management policy

You can use Query Patroller to implement a query handling policy, specifying exactly how queries are handled from the time they are submitted. As part of such a policy, you need to define exactly which queries are intercepted or managed and which queries are left to run without interference. You can also specify the circumstances under which queries are held or queued based on user or system resource thresholds.

Setting up a query handling policy involves several decision points:

1. Deciding which types of queries will be intercepted and managed by Query Patroller
2. Deciding whether to limit overall database workload

3. Deciding which queries will be held and when
4. Devising a policy for handling held queries
5. Deciding whether to use query classes to manage queries of different sizes

Decide which types of queries will be intercepted and managed by Query Patroller

Whenever Query Patroller intercepts a query to perform a cost estimate, this uses system resources. It also impacts the response time of the query. Similarly, once Query Patroller has intercepted a query, managing that query (assigning it to a query class, queuing it, checking the submitter priority, and so on) requires system resources and impacts the response time for the query. You need to weigh the benefits of intercepting and managing a query or group of queries against the impact that the overhead will have on performance and on the availability of system resources.

Given these considerations, you may not want Query Patroller to intercept certain kinds of queries:

- Queries that require almost instantaneous response, such as queries submitted by OLTP (online transaction processing) applications
- Queries that are submitted by users whose requirement for rapid response outweighs the benefits of interception and tracking by Query Patroller

You might want to roll out Query Patroller in stages and begin by intercepting queries from only a subset of submitters or applications. That way, you can test out your initial configuration without impacting all of your users.

Also, you may not want Query Patroller to manage queries that are too small to be a major impact on system performance. By setting a minimum query cost threshold for Query Patroller to manage queries, you can specify immediate processing of queries that are intercepted by Query Patroller but are estimated to be relatively small. This threshold is set for individual submitters or groups of submitters in their submitter profiles. Query Patroller provides a default value for this setting, but if you decide to customize this value for different submitters, you will need to determine an appropriate minimum query cost for management in your system.

Decide whether to limit overall managed query workload

You have the option of setting a limit on the total number of managed queries that run concurrently as well as the total aggregate cost of all managed queries running concurrently. By default, these Query Patroller settings are set to unlimited. If you decide to control the overall managed query workload, you need to determine what the optimum workload is for your system.

Decide which queries will be held and when

You can specify a cost limit on the size of queries submitted by individuals or groups of submitters. Queries exceeding this limit will be held by Query Patroller.

Query Patroller provides defaults for these cost limits on queries that you can use as a starting point, but if you decide to customize these settings, you will need to determine the appropriate query cost limits to set for the different submitters in your system.

Devise a policy for handling held queries

You need to decide how queries will be handled once they have been held. For example, you will need to decide whether held queries will be individually screened by an administrator or operator to see if it reasonable to run the queries despite their size or whether they will be canceled automatically for exceeding size thresholds. You also have the option to run held queries on a schedule.

If you decide to run held queries on a schedule, you need to decide what that schedule should be, taking into consideration factors such as periods of peak database usage and maintenance activities.

Decide whether to use query classes to manage queries of different sizes

Query classes can help you control the flow of queries against your database. If you need to limit the number of large queries running concurrently in your system, or you want to prevent smaller queries from being queued behind larger queries, you can create query classes to run different size queries.

Chapter 14. Managing space

Setting Query Patroller maintenance schedules for queries and result tables

Managed queries and result tables consume space, so they should be removed periodically. You can schedule Query Patroller to purge managed queries and result tables at regular intervals, on an ongoing basis.

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile has the CONFIGURATION privilege with edit authority

You can set maintenance schedules for queries and result tables using one of the following methods:

- Query Patroller Center method
- Command line method
- To set maintenance schedules for queries and result tables using the Query Patroller Center method
 1. Open the Query Patroller System Properties notebook:
 - a. From the Query Patroller Center, expand the object tree under the **Configuration** folder to find the **Query Patroller System** folder.
 - b. Click the **Query Patroller System** folder. A system property record is displayed in the pane on the right side of the window (the contents pane).
 - c. Right-click the system property record, and click **Properties** in the pop-up menu. The Query Patroller System Properties notebook opens.

Note: You can also open the Query Patroller System Properties notebook by double-clicking the entry in the contents pane when the Query Patroller System folder is displayed.

2. Click the **Options** tab.
3. Under **Managed query settings**, use the **How long to keep queries** fields to specify how long to keep queries in the Managed Queries folder. To keep queries indefinitely, leave the field blank.
4. Use the **How long to keep result tables** fields to specify how long to keep result tables after the completion of a query. To keep result tables indefinitely, leave the field blank.

Note: If you delete a managed query or a historical query, the corresponding result table is also deleted. Therefore the length of time specified in the **How long to keep queries** fields should be greater than or equal to the length of time specified in the **How long to keep result tables** fields.

5. Optional: In the **Table space for result tables** type the name of the table space where result tables will be stored. To use the default DB2 table space, leave the field blank.
6. Click **Schedule Purge Job** to open a window where you manually schedule purges of managed queries and result tables. .

Note: If you delete a managed query, the corresponding query in the Historical Analysis folder is not deleted. However, if you purge a historical query, and it also exists under Managed Queries, it gets deleted in both places.

7. Click **OK** to set the maintenance schedules.
- To set maintenance schedules for queries and result tables using the Command line method:
 1. Issue the UPDATE QP_SYSTEM command using the following parameters:
 - QUERY_PURGE_PERIOD
 - RESULT_PURGE_PERIOD
 - RESULT_TABLE_SPACE

Note:

- a. To schedule purges of queries and result tables you must set both the QUERY_PURGE_PERIOD parameter and the RESULT_PURGE_PERIOD parameter to a value other than -1 (the value -1 represents unlimited).
- b. When a query is purged, the corresponding result table is also deleted. Therefore, you should set the value of the QUERY_PURGE_PERIOD parameter to be greater than or equal to the value of the RESULT_PURGE_PERIOD parameter. Otherwise, result tables will be purged at the same time as their associated queries, even if the value of the RESULT_PURGE_PERIOD parameter is greater.
- c. You can use the UPDATE QP_SYSTEM command to schedule purges of queries and result tables only if a valid purge schedule exists. A default schedule for purging queries and result tables is created at installation. You can modify or delete the purge schedule using the Query Patroller Center. You cannot modify or delete the purge schedule using the command line.

Managing historical queries

Historical queries are used by Query Patroller to generate historical data. Historical queries consume space, so they should be removed periodically. However, once historical queries have been removed, they can no longer be used for generating historical data. You can schedule Query Patroller to purge historical queries at regular intervals, on an ongoing basis.

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile has the CONFIGURATION privilege with edit authority

To schedule purges of historical queries, use one of the following methods.

- To schedule purges of historical queries using the Query Patroller Center method:
 1. Open the Query Patroller System Properties notebook:
 - a. From the Query Patroller Center, expand the object tree under the Configuration folder to find the **Query Patroller System** folder.
 - b. Click the **Query Patroller System** folder. A system property record is displayed in the pane on the right side of the window (the contents pane).

- c. Right-click the system property record, and click **Properties** in the pop-up menu. The Query Patroller System Properties notebook opens.
2. Click the **Options** tab.
3. Under **Historical analysis**, specify the queries that you want to save by selecting the **Only managed queries** or the **All intercepted queries** radio button.
4. Under **Historical analysis**, use the **How long to keep queries** fields to specify how long to keep completed queries in the Historical Analysis view.
5. Click **Schedule Purge Job** to open a window where you manually schedule purges of historical queries. See “Scheduling purges of historical queries” on page 96

Note: If you delete a managed query, the corresponding query in the Historical Analysis folder is not deleted. However, if you purge a historical query, and it also exists under Managed Queries, it gets deleted in both places.

6. Click **OK** to schedule purges of historical queries.
- To schedule purges of historical queries using the Command line method:
 1. Issue the **UPDATE QP_SYSTEM** command using the following parameters:
 - QUERIES_TO_SAVE
 - HISTORY_PURGE_PERIOD

Note: You can use the **UPDATE QP_SYSTEM** command to schedule purges of historical queries only if a valid purge schedule exists. A default schedule for purging historical queries is created at installation. You can modify or delete the purge schedule using the Query Patroller Center. You cannot modify or delete the purge schedule using the command line.

Scheduling purges of managed queries and result tables

You can schedule purges of managed queries and result tables to run at a specific time, on a recurring basis. For instance, you can schedule purges of managed queries and result tables to run on the last Sunday of each month at 1:00 in the morning.

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile has the CONFIGURATION privilege with edit authority

Note:

- The same schedule is used for purging both managed queries and result tables; you do not have to create separate schedules.
- If you are an operator, but you do not have edit authority, you can still view the schedule, but you cannot change the schedule.
- You can manually set schedules purging managed queries only through the Query Patroller Center. You cannot manually set schedules using the command line.

To schedule purges of managed queries and result tables, use the following Query Patroller Center method.

1. Open the Schedule Managed Queries Purge Job window:

2. To specify how often you want to purge managed queries, use the **Interval** and **Frequency** fields. Depending on the interval you select, you may also need to use the **Details** field.
3. To specify when to begin purging managed queries, use the **Date** and **Time** fields.
4. Optional: To prevent the task from running, select the **Suspend Schedule** check box.
5. Specify an end date for purging managed queries. If there is no end date, select the **No end date** check box. Otherwise, to specify an end date, use the **Date** and **Time** fields.
6. Click the **Add** button to add the task to the **Schedule list**.
7. Click **OK** to schedule the task.

Scheduling purges of historical queries

You can schedule purges of historical queries to run at a specific time, on a recurring basis. For instance, you can schedule historical queries to run on the last Sunday of each month at 1:00 in the morning.

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile has the CONFIGURATION privilege with edit authority

Note:

- You can manually set schedules for purging historical queries only through the Query Patroller Center. You cannot manually set schedules using the command line.
- If you are an operator, but you do not have edit authority, you can still view the schedule, but you cannot change the schedule.

To schedule purges, use the following Query Patroller Center method:

1. Open the Schedule Historical Analysis Purge Job window:
 - a. Open the Query Patroller System Properties notebook.
 - b. Click the **Options** tab.
 - c. Under **Historical analysis**, click the **Schedule Purge Job** push button. The Schedule Historical Analysis Purge Job window opens.
2. To specify how often you want to purge historical queries, use the **Interval** and **Frequency** fields. Depending on the interval you select, you may also need to use the **Details** field.
3. To specify when to begin purging historical queries, use the **Date** and **Time** fields.
4. Optional: To prevent the task from running, select the **Suspend Schedule** check box.
5. Specify an end date for purging historical queries. If there is no end date, select the **No end date** check box. Otherwise, to specify an end date, use the **Date** and **Time** fields.
6. Click the **Add** button to add the task to the **Schedule list**.
7. Click **OK** to schedule the task.

Dropping result tables manually using Query Patroller

Drop a result table when you want to free up the space that the result table is consuming.

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile has the MONITORING privilege with edit authority
- Be the submitter of the query that produced the result table to be dropped

A result table is created when:

- The query submitter's submission preferences specify that the application should be released after a query is submitted
- A query was held and then run
- A query was run in the background

There are several ways to drop result tables:

- You can drop a result table immediately using the Query Patroller Center.
- You can schedule result tables to be dropped using the Query Patroller Center.
- You can drop a result table using the command line.

To drop a result table, use any of the following methods.

- Query Patroller Center methods
- Command line method
- To drop a result table using Query Patroller Center method #1
 1. Open the Managed Query Properties notebook:
 - a. From the Query Patroller Center, expand the object tree under the **Monitoring** folder to find the **Managed Queries** folder.
 - b. Click the **Managed Queries** folder. Any managed queries are displayed in the pane on the right side of the window (the contents pane).
 - c. In the contents pane, right-click the query that you want to work with, and click **Properties** in the pop-up menu. The Managed Query Properties notebook opens.

Note: You can also open the Managed Query Properties window by double-clicking the query that you want to work with.

2. Click the **Results** tab.
 3. Click **Drop Result Table**.
- To drop a result table using Query Patroller Center method #2
 1. Open the Query Patroller System Properties notebook:
 2. From the Query Patroller Center, expand the **Monitoring** folder to find the **Managed Queries** folder.
 3. Click the **Managed Queries** folder. Any managed queries are displayed in the pane on the right side of the window (the contents pane).
 4. In the contents pane, right-click the query that you want to drop the result table for, and click **Drop Result Table** in the pop-up menu.

Note: The Drop Result Table is visible only when query results are stored in the result table.

5. Optional: You can drop several result tables at once by holding the **Ctrl** button as you click the queries whose result tables you want to drop. Then right-click and click **Remove** in the pop-up menu.
- To drop a result table using Query Patroller Center method #3
 1. You can also schedule result tables to be dropped at a specific time, by scheduling purges of managed queries and result tables.
- To drop a result table using the Command line method:
 1. Issue the REMOVE RESULT command

Removing managed queries manually

Removing a managed query from Query Patroller removes the managed query information from the system. The associated historical query information is not removed. If a result table exists for the managed query being removed, the result table is dropped.

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile has the MONITORING privilege with edit authority

Instead of removing managed queries manually, you can also schedule queries to be purged.

To remove managed queries manually, use one of the following methods.

- **Query Patroller Center method**
 1. From the Query Patroller Center, expand the Monitoring folder to find the Managed Queries folder.
 2. Click the Managed Queries folder. Any managed queries are displayed in the pane on the right side of the window (the contents pane).
 3. In the contents pane, right-click the query that you want to remove, and click Remove in the pop-up menu.
 4. Optional: You can drop several queries at once by holding the Ctrl button while selecting the queries you want to drop. Then right-click and click Remove in the pop-up menu.
- **Command line method**

Issue the REMOVE QUERY_INFO command.

Removing historical queries manually

Removing a historical query from Query Patroller removes the historical query information from the system. If it exists, the associated managed query information is also removed. If a result table exists for the historical query being removed, the result table is dropped. Instead of removing historical queries manually, you can also schedule queries to be purged.

You must meet one of the following requirements:

- Have DBADM authority
- * Be an operator whose profile has the HISTORICAL ANALYSIS privilege with edit authority Procedure

To remove a historical query manually, use one of the following methods.

- **Query Patroller Center method**
 1. From the Query Patroller Center, expand the Historical Analysis folder to find the Queries folder.
 2. Click the Queries folder. Historical queries that fall within the time range specified in the historical analysis toolbar are displayed in the pane on the right side of the window (the contents pane). Use the historical analysis toolbar to specify the time range you are interested in.
 3. In the contents pane, right-click the query that you want to remove, and click Remove in the pop-up menu.
 4. Optional: You can drop several queries at once by pressing the Ctrl button and selecting the queries you want to drop. Then right-click and click Remove in the pop-up menu.
- To remove a historical query manually for a particular submitter:
 1. From the Query Patroller Center, expand the Historical Analysis folder to find the Submitters folder. Expand the Submitters folder and click on the folder for the particular submitter you want. Expand that submitter's folder to find the Queries folder.
 2. Click the Queries folder. The submitter's historical queries that fall within the time range specified in the historical analysis toolbar are displayed in the pane on the right side of the window (the contents pane). Use the historical analysis toolbar to specify the time range you are interested in.
 3. In the contents pane, right-click the query that you want to remove, and click Remove in the pop-up menu.
 4. Optional: You can drop several queries at once by pressing the Ctrl button and selecting the queries you want to drop. Then right-click and click Remove in the pop-up menu.
- **Command line method**
Issue the `REMOVE QUERY_INFO_HISTORY` command.

Removing orphaned result table aliases

Aliases created with `CREATE_RESULT_TABLE_ALIASES` option of the `UPDATE QP_SYSTEM` command are automatically dropped when a result table is dropped.

You must have DBADM authority.

However, there are two situations in which a result table may be dropped without the corresponding alias being dropped.

- When the result table is dropped manually without using the `qp` command line or Query Patroller Center.
- When the result table is dropped using the `qp` command line or Query Patroller Center under the authority of an operator who is not the submitter of the query and does not have DBADM authority.

This procedure is only available using the Command line. To clean up aliases that have no corresponding result tables, a new command, `REMOVE RESULT_TABLE_ALIASES`, has been created. This command is automatically executed whenever result tables are purged as part of the Query Patroller scheduled result table purging process. The `REMOVE RESULT_TABLE_ALIASES` command obtains the list of aliases to purge using the following query:

```
with a as (select tabschema, tabname from syscat.tables
           where type = 'A' and tabname like 'QUERY%_RESULTS'),
t as (select tabname from syscat.tables
```

```
        where type = 'T' and tabname like 'QUERY%_RESULTS')
select all tabschema, tabname from a
where not exists (select * from t where t.tabname=a.tabname)
```

Issue the REMOVE RESULT_TABLE_ALIASES command

This command removes all aliases that exist after their corresponding result tables have been dropped. The aliases were originally created by Query Patroller for result tables.

Chapter 15. Query class configuration

Query classes allow you to control the number of queries of a particular size that can run concurrently in your system. Before you create query classes for your system, you need to decide how many query classes to create and how to configure each query class.

Step 1: Characterize your workload

To determine the type and number of query classes to create, you first need to understand the characteristics of the workload that is typical for your system. In particular, you need to have a good sense of the size ranges of the different queries that run against your data warehouse, and the distribution or frequency of each size grouping.

If you do not have this information already, you can collect historical data to get a representative sample of database activity and use the Query Patroller Historical Analysis reports to determine what size groupings exist for your current workload.

Once you have collected the workload data, examine the "Query activity over time" report in the Historical Analysis folder of the Query Patroller Center. Sort the data by the Estimated Cost column by clicking on the column header.

This report will show you any natural size groupings that you can use as a guide in creating query classes.

For example, Table 8 shows the compiled figures for query activity in a given data warehouse over a two-week period.

Table 8. Distribution of queries in example workload

Query size (timerons)	Number of queries	Percent of total queries in workload
<8099	1588	88
10190–96444	206	11
>1030000	17	1

The queries in this workload fall naturally into three different size ranges: the estimated cost of the smallest (and most frequent) type of query fall below 8099 timerons. The next size range for queries in this workload is between 10190 and 96444 timerons. The largest and least frequent queries are larger than 1030000.

Based on these groupings, you can plan for your query classes. You may want to plan for a query class for each major size grouping that you identify.

Note: The creation of too many query classes can negatively impact performance so you will have to monitor your system performance after any configuration modifications.

Step 2: Decide on the maximum query cost for each query class

The maximum cost for queries in each query class should correspond to the upper end of the size range for the queries that the query class is intended to run.

Based on the data in Table 8 on page 101, you might create three query classes to correspond to the three size groupings identified. These query classes are shown in Table 9.

Table 9. Logical size groupings for query classes

Query class	Maximum size of query (timerons)
1	10000
2	100000
3	10000000

If your data reveals a continuous distribution of queries across the entire range of query sizes and there are no natural size-based divisions in your query set, you might need to decide on arbitrary divisions for your query classes that correspond to small, medium, and large queries.

Note: Each query class in your system must have a distinct value for the maximum cost of a query.

Step 3: Decide on the maximum number of queries for each query class

To make your query classes allocate the use of system resources more effectively, you can limit the number of queries that can run concurrently in each query class. The key to setting the maximum number of queries for a query class is to weigh the query execution time against the distribution of a particular size of query in the workload.

When deciding on the maximum number of queries for each query class, it might help to think of your system like a parking lot for queries that provides both short-term parking (for example, 15 minutes) and long-term parking (for example, up to 10 hours). Even if the short-term parking customers represent 50% of the total customers for the parking lot, you would not want to allocate 50% of the parking spaces for short-term parking because short-term customers only use their parking spot briefly. This means that there would likely be many unused short-term spots at any given time. Meanwhile, you might have long-term parking customers waiting for long-term spots to become available.

For example, imagine that the maximum number of queries (`MAX_TOTAL_QUERIES`) for the system shown in Table 8 on page 101 and Table 9 is set to 100. The queries that will run in query class 1 represent 88% of the overall workload. However, these queries take only moments to run. You might set the maximum number of queries (`MAX_QUERIES`) for this query class to 60. This number is large enough to ensure that small queries do not have to wait to run, but small enough that there will not be empty spots in this query class while larger queries in other query classes are queued. You might then set the maximum number of queries in query classes 2 and 3 to 30 and 10 queries respectively.

Another strategy in setting the maximum number of queries for query classes is to place limits on the number of larger queries only, leaving the number of small queries unlimited. By setting the maximum number of queries for the query class for larger queries you can limit the number of large queries running in the system. By setting the maximum number of queries for the smaller query class to unlimited, you ensure that smaller queries are not put in the queue even though sufficient resources are available to run them.

Step 4: Evaluate your query class definitions

You can evaluate the efficacy of your query class settings using the Managed Queries folder in the Query Patroller Center. In this folder, you can see how many queries are currently running in each query class, and how many queries are queued in each query class. If you check this view periodically as you run your workload, you can see if there is excessive queuing of queries in some classes and adjust the maximum number of queries for your query classes accordingly.

You can also use the Query Activity graph and report in the Historical Analysis folder to see whether queries have spent time in the queue during execution.

Configuring query classes

You can use query classes to control the number of managed queries of a particular size that can run concurrently in your system.

You must have read the topic "Query class configuration".

To configure query classes:

1. Using Query Patroller Historical Analysis reports, collect system and workload statistics.
2. To identify any natural size groupings for queries (based on estimated query cost) to use as a guide in creating query classes, examine the "Query activity over time" report which appears in the **Historical Analysis >> Queries** folder of the Query Patroller Center.
3. Sort the queries data by the **Estimated Cost** column by clicking on the column header.

Note: If the **Estimated Cost** column does not appear, click the **Customize columns** icon in the bottom toolbar.

4. Decide on the maximum query cost (in timerons) for each query class.

Note: Each query class in your system must have a distinct value for the maximum query cost. The maximum cost for queries in each query class should correspond to the upper end of the size range for queries that the query class is intended to run.

5. Decide on the maximum number of queries that can run concurrently in each query class.
6. Create query classes for Query Patroller.
7. Evaluate and refine your query class definitions. In the "Managed Query" report in the Query Patroller Center, you can see how many queries are currently running in each class and how many queries are queued in each query class.

Note: If the **Query class** column does not appear, click the **Customize columns** icon in the bottom toolbar.

If you have not already done so, you can further control the flow of queries against your database by setting or modifying the system thresholds for managed queries or by setting or modifying submitter resource limits.

Creating query classes for Query Patroller

By creating query classes, you can control the number of queries of various sizes that are running against the system at any one time. When you create a new query class, you specify its properties, and you can provide a useful description of that query class.

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile has the CONFIGURATION privilege with edit authority

When you create query classes, you should consider the following factors:

- The performance of your database
- The number of users submitting queries
- The average cost of the queries being submitted at any given time

Query classes can be created, removed, or modified while Query Patroller is started. Creation, changing the maximum query cost, or removal of a query class will take effect immediately unless there are queued or running queries. If there are queued or running queries, including newly submitted queries, the query class changes will take effect when these complete. If you do not want to wait for all queued and running queries to complete, a Query Patroller server restart is required.

Updating the maximum number of queries for a query class always takes effect immediately.

You can create a new query class using one of the following methods:

- Query Patroller Center method
- Command line method
- To create a new query class using the Query Patroller Center method:
 1. Open the Create Query Class window:
 - a. From the Query Patroller Center, expand the object tree under the **Configuration** folder to find the **Query Classes** folder.
 - b. Right-click the **Query Classes** folder, and click **Create** in the pop-up menu. The Create Query Class window opens.
 2. In the **Maximum number of queries** field, specify the maximum number of queries of this query class that can be run simultaneously. Queries belonging to this query class and submitted after this value is reached are queued by Query Patroller. This number must be less than or equal to the **Maximum number of queries** set in the Query Patroller System Properties notebook. If **Maximum number of queries** for the system is unlimited, you can create a query class with an unlimited maximum number of queries by leaving this field blank.
 3. In the **Maximum cost of a query** field, specify the maximum size for queries that belong to this query class. The query class that you are defining will contain queries that are smaller than this **Maximum cost of a query** amount, and larger than the **Maximum cost of a query** amount of the next smallest query class. You cannot define two query classes that have the same **Maximum cost of a query** amounts.

4. Optional: In the **Query class description** field, enter a comment or a description for the query class that you are creating.
 5. Click **OK** to create your query class and close the Create Query Class window. When you click **OK**, a Query class ID will be assigned by Query Patroller.
- To create a new query class using the Command line method:
 1. Issue the ADD QUERY_CLASS command using the following parameter:
 - Issue the ADD QUERY_CLASS command.

Note: The following warning message is returned when you add a new query class while the Query Patroller server is running.

```
DQP1024W Creation, change, or removal of a query class
will not take effect until the Query Patroller server is restarted.
```

This message is incorrect. The query class creation will take effect immediately unless there are queued or running queries. If there are queued or running queries, including newly submitted queries, the query class changes will take effect when the queued or running queries complete. If you do not want to wait for all queued and running queries to be completed, you must restart the Query Patroller server. Updating the maximum number of queries for a query class always takes effect immediately.

Removing query classes for Query Patroller

Remove a query class when you want to change the query classification that you are using. You might remove a query class if you discover that it is not being used, or if you want to resize all of your query classes.

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile has the CONFIGURATION privilege with edit authority

Query classes can be created, removed, or modified while Query Patroller is started. Creation, changing the maximum query cost, or removal of a query class will take effect immediately unless there are queued or running queries. If there are queued or running queries, including newly submitted queries, the query class changes will take effect when these complete. If you do not want to wait for all queued and running queries to complete, a Query Patroller server restart is required.

Updating the maximum number of queries for a query class always takes effect immediately.

You can remove a query class using one of the following methods:

- Query Patroller Center method
- Command line method
- To remove a query class using the Query Patroller Center method:
 1. From the Query Patroller Center, expand the **Configuration** folder to find the **Query Classes** folder.
 2. Click the **Query Classes** folder. Any existing query classes are displayed in the pane on the right side of the window (the contents pane).

3. In the contents pane, right-click the query class that you want to remove, and click **Remove** in the pop-up menu. A window opens where you can confirm that you want to remove the query class.
- To remove a query class using the Command line method:
 1. Issue the REMOVE QUERY_CLASS command.

Note: The following warning message is returned when you remove a query class while the Query Patroller server is running.

```
DQP1024W Creation, change, or removal of a query class
will not take effect until the Query Patroller server is restarted.
```

This message is incorrect. The query class removal will take effect immediately unless there are queued or running queries. If there are queued or running queries, including newly submitted queries, the query class changes will take effect when the queued or running queries complete. If you do not want to wait for all queued and running queries to be completed, you must restart the Query Patroller server. Updating the maximum number of queries for a query class always takes effect immediately.

Query Patroller query classes

Query classes are Query Patroller mechanisms for grouping and running queries according to size. Using query classes, you can control the flow of queries on a database so that system resources are shared among queries in the different size groupings. For example:

- You can ensure that smaller queries do not get stuck behind larger ones by specifying one query class for small queries and one for large queries. The small queries then run in their own query class and are queued separately from the large queries.
- You can prevent large queries from monopolizing system resources by specifying the number of queries in a single query class that can run at once, thus limiting the number of large queries that are processed concurrently.

Query classes categorize managed queries according to the estimated cost of each query. Each query class that is defined for a particular Query Patroller system accepts managed queries whose estimated cost falls within a specified range. For example, if query class 1 accepts queries costing up to 10 000 timerons, a query whose estimated cost is 7500 timerons will run in that class.

Each query class can be configured to limit the number of queries running concurrently in that query class. For example, query class 2 can be configured to allow 10 queries to run simultaneously by setting the maximum number of queries (MAX_QUERIES parameter) value for this query class to 10. Once this limit is reached, any additional queries that fit into that class are queued until one of the running queries in that class completes.

The upper limit for a query class is set by the maximum query cost (MAX_COST parameter) value for the query class. While you do not explicitly define a lower limit for a particular query class, a lower limit is enforced through the interaction of the following two factors:

1. A query always runs in the query class with the lowest estimated cost limit possible. For example, although a query with an estimated cost of 1200 timerons could run in either query class 1 with a MAX_COST value of 10 000

or query class 2 with a MAX_COST of 100 000, it will always run in query class 1. This means that query class 2 will never receive queries smaller than 10 001 timerons.

2. Each submitter profile contains a value for the minimum cost to manage a query (MIN_COST_TO_MANAGE parameter) that specifies the lowest cost a query can have in order for it to be managed by Query Patroller. For example, if the minimum query cost for the submitter profile defined for the managers group is 1000 timerons, any queries with an estimated cost smaller than 1000 timerons submitted by a member of this group will not be put into a query class. Instead, these queries will run without being queued.

Each query class defined within your system must have a unique value for the maximum cost of each query (MAX_COST) parameter.

Each query class is stored along with its parameter values as a row in the QUERY_CLASS control table.

Default query class

Every system will have a default query class even if no query classes have been defined by the administrator. The maximum query cost for this default query class is unlimited.

A query will run in the default query class under two conditions:

- The estimated cost of the query is less than or equal to the system maximum query cost (MAX_TOTAL_COST parameter) but larger than the maximum query cost for any user-defined query class. In cases such as this where query classes have been defined, only a single query can run at a time in the default query class.
- When no query classes have been defined by the administrator, all queries will run within the default query class. In cases such as this where no query classes have been defined, an unlimited number of queries can run at once in the default query class.

Note: The number of queries running simultaneously in this class can be indirectly constrained by the system-wide limit on the total number of queries in the system (MAX_TOTAL_QUERIES parameter).

Chapter 16. Query Patroller thresholds

Query Patroller relies on resource thresholds to determine the flow of the workload that can run against a database. Thresholds can be set to control the number and size of queries run by a particular submitter or group, to control the size of the overall system workload, or both. Depending on the characteristics of your particular system and workload, you might want to set some or all of these thresholds. The main types of thresholds are submitter thresholds and system thresholds.

Submitter thresholds

Submitter thresholds are set in submitter profiles, so thresholds can be set for an individual submitter or group of submitters. Each individual user can belong to multiple groups, therefore a user may have multiple group profiles as well as an individual submitter profile. When a user has multiple profiles that they can submit queries under, the appropriate threshold setting is chosen based upon a set of rules that select the effective profile for the submitter. For information on how Query Patroller chooses which submitter profile to use for a submitter, see the “Query Patroller submitter profiles” on page 65 topic.

The maximum cost threshold (the `MAX_COST_ALLOWED` parameter) for a submitter determines the maximum cost of query that a submitter can run. If a submitter tries to run a query whose estimated cost exceeds the maximum cost for his or her profile, then the query is held.

You might want to set this threshold if you have problems with runaway queries tying up system resources. By setting a size limit on individual queries, you can separate unreasonably large queries before they execute, and hold these queries so that you can decide to run them or cancel them on a case-by-case basis.

The maximum number of queries value for a submitter (`MAX_QUERIES_ALLOWED`) determines the number of queries that can be run simultaneously by a particular submitter. If a submitter tries to run a query when the maximum number of queries specified in his or her submitter profile is already running, the query will be queued until one of the queries completes.

You might want to set this threshold if you are allowing a submitter group to submit large queries, but you want to limit each individual submitter to a few queries. Setting this threshold can also address the problem of users resubmitting the same costly queries in rapid succession because the response time is slower than they expect.

You may want to set one or both submitter thresholds if you need to control the use of resources by particular users or groups. For example, if the data warehouse is primarily funded by one department, but another department is permitted limited use of the resources, you might want to set submitter thresholds so that the users or groups in the second department are restricted in the amount of resources they can tie up with their queries. The thresholds for the users and groups in the funding department can be more liberal in the amount of resources allowed for each submitter.

System thresholds

You can control the overall workload that executes against your database by setting system thresholds for cost (the `MAX_TOTAL_COST` parameter) and number of queries (the `MAX_TOTAL_QUERIES` parameter).

The maximum workload cost value for the system (`MAX_TOTAL_COST` parameter) determines the maximum size of the overall workload running against the database. The workload cost is calculated by adding the cost estimates of all the Query Patroller managed queries currently running in the system. If the execution of a new query will cause the aggregate cost of all queries currently running to exceed the maximum workload cost, the new query is placed in a queued state until the system can run the new query without exceeding the maximum workload cost.

You might want to set this threshold if you find that your CPU utilization is uncomfortably high when running your current workload, and you want to lower the risk of overloading the system by placing a restriction on workload size.

The maximum number of queries value for the system (`MAX_TOTAL_QUERIES` parameter) places a limit on the number of queries that can be running simultaneously in the system. When this threshold is reached, additional queries are placed in a queued state where they wait until the system can run the queries without exceeding the maximum number of queries value.

You might want to set this threshold if you have a need to limit the number of concurrent queries in order to avoid overloading the system.

Part 5. About Query Patroller historical analysis

You can use the historical analysis functions provided by Query Patroller to analyze various aspects of your data warehouse use over time. This allows you to gather information for business reports, performance tuning (through the optimization of Query Patroller system thresholds), and the identification of redundant database objects.

Important: With the new workload management features introduced in DB2 Version 9.5, Query Patroller and its components have been deprecated in Version 9.7 and might be removed in a future release. For more information, see the “DB2 Governor and Query Patroller have been deprecated” topic in the *What's New for DB2 Version 9.7* book.

You can find out whether certain tables, indexes, and columns are being used, and by which query submitters. From information such as the number of queries run and query execution time, you can see how resource usage varies by month, week, day, hour, or minute.

By using the Query Patroller historical analysis functions, you might be able to identify trends in data warehouse usage:

- The table, index, and column reports can help you identify which database objects are not being used and can therefore be removed to conserve storage space.
- Query reports by submitter, showing the number and size of queries, can help you identify which submitter groups are the most significant consumers of system resources.
- Reports on the number of queries run against the data warehouse over time can help you identify peak traffic times and off-peak times so that you can schedule maintenance tasks during off-peak times
- Reports on query execution over time can help you identify how various system, hardware, and configuration changes have impacted performance

By default, all successfully completed queries managed by Query Patroller are included in historical analysis. Additionally, you can include in your historical analysis successfully completed queries that were intercepted but not managed by Query Patroller.

In order to have the most recent information available for historical analysis, you must generate historical data. Generating historical data runs the SQL Explain facility against the queries that Query Patroller has saved for historical analysis and makes the information available in the Query Patroller historical analysis reports and graphs.

You can remove queries from historical analysis as you see fit by scheduling a regular purge of historical queries or by removing queries individually.

Chapter 17. Getting started with the Query Patroller historical analysis interface

Important: With the new workload management features introduced in DB2 Version 9.5, Query Patroller and its components have been deprecated in Version 9.7 and might be removed in a future release. For more information, see the “DB2 Governor and Query Patroller have been deprecated” topic in the *What’s New for DB2 Version 9.7* book.

Use the historical analysis functions of the Query Patroller Center to examine various aspects of data warehouse use over time.

The historical analysis reports include reports on queries, tables, indexes, and submitters, for example:

- Query Activity over Time report: Use this report to determine such things as peak database usage and which queries are most costly.
- Tables Not Hit report: Use this report to determine which tables to eliminate.
- Submitters report: Use this report to determine information about submitters, such as which submitters are using the most resources.

Authorities and privileges

To generate historical data, you must meet the following requirement:

- Have DBADM authority
- Have SETSESSIONUSER privilege on PUBLIC

To perform historical analysis tasks that modify Query Patroller data, such as removing historical queries manually, you must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile has the HISTORICAL DATA privilege with edit authority

To perform historical analysis tasks that involve viewing Query Patroller data, such as filtering tables, you must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile has the HISTORICAL DATA privilege with edit or view authority

Historical analysis tasks

Depending on your Query Patroller authority and privileges, you may be able to perform some or all of these tasks:

- Chapter 21, “Generating historical data using Query Patroller,” on page 123
- Chapter 23, “Viewing historical query details using Query Patroller,” on page 127
- Chapter 24, “Filtering tables for historical analysis using Query Patroller,” on page 129
- “Managing historical queries” on page 94

The Query Patroller Center historical analysis interface

There are two main types of views for Query Patroller historical analysis data, a table view and a graphical view.

Historical analysis toolbar



Use the historical analysis toolbar to display and work with the time intervals you are interested in, and to move between the graphical view and the table view.

Use the **Interval** and **End date** fields to establish a time range that determines which items appear in the folders under the **Historical Analysis** folder and in the historical analysis contents pane.


When you have specified an interval and an end date, you must click the **Apply** push button to refresh the view to show the historical analysis information with those settings.

Use the **Previous Interval** and **Next Interval** push buttons to move through the historical data displayed in the contents pane. Clicking **Previous Interval** moves the end date back by the interval specified in the **Interval** fields. For example, if the interval value is 5, the unit of time is days, and the end date is May 22, 2003, clicking **Previous Interval** will update the end date to May 17, 2003. Clicking **Next Interval** advances the end date by the interval specified in the **Interval** fields. For example, if the interval value is 5, the unit of time is days, and the end date is May 22, 2003, clicking **Next Interval** updates the end date to May 27, 2003. Clicking the **Previous Interval** and **Next Interval** push buttons automatically refreshes the information displayed in the contents pane and the items under the **Historical Analysis** folder.

When you are looking at the graphical view of a report, you can use the **Show Table** push button on the far right of the toolbar to switch to the table view. When you are looking at the table view, the **Show Table** push button becomes the **Show Graph** push button, so you can use it to switch back to the graphical view.

For historical analysis data in the **Queries** folder only, there is a histogram view instead of a graph. The **Show Graph** push button becomes the **Show Histogram** push button. In the histogram view, you can double-click the individual histogram bars to obtain information at a finer level. For example, if you are looking at a report that covers three months of queries, the report shows three histogram bars. Double-clicking one of the histogram bars displays the data for the weeks that comprise that month. If you have drilled down to obtain information at a finer level, you can then drill up to obtain information at a higher level by double-clicking with the right mouse button on any of the histogram bars. For example, if you have drilled down from a report showing three months of queries to a report showing four weeks of queries, double-click with the right mouse button to drill back up to the histogram that displays months.

Changing the information displayed in the historical analysis histogram



Use the Y-axis drop-down menu to specify the type of information you want to display about historical queries in the histogram. You can choose from the following six options:

- Queries Run - Average Time
- Queries Run - Total Time
- Queries Queued - Average Queue Time
- Queries Queued - Total Queue Time
- Queries Run - Number Run
- Queries Queued - Number Queued

Related tasks:

- Chapter 24, “Filtering tables for historical analysis using Query Patroller,” on page 129
- Chapter 25, “Filtering queries for historical analysis using Query Patroller,” on page 131

Chapter 18. Enabling collection of historical data

Query Patroller historical analysis can be a powerful tool for analyzing data warehouse usage. You can choose to collect historical data only for queries that were managed by Query Patroller, or for all queries intercepted by Query Patroller. Queries that are not intercepted cannot have historical data collected about them.

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile has the CONFIGURATION privilege with edit authority

To enable collection of historical data, use one of the following methods.

- Query Patroller Center method
- Command line method

The default behavior is that only managed queries will have historical data collected.

Note: If you are planning to collect data regarding query execution time, you need to ensure that the DB2 timestamp and statement monitor switches are set to 'ON'. If you are planning to collect data about how many rows are returned by queries, you need to ensure that the DB2 statement monitor switch is set to 'ON'.

- To enable collection of historical data using the Query Patroller Center method:
 1. Open the Query Patroller System Properties notebook.
 2. Click the **Options** tab.
 3. Under **Historical analysis**, specify the queries that you want to save by selecting the **Only managed queries** or the **All intercepted queries** radio button.
 4. Click **OK** to save your changes.
- To enable collection of historical data using the Command line method:
 1. Issue the UPDATE QP_SYSTEM command using the following parameter:
 - QUERIES_TO_SAVE

Chapter 19. Uses for historical analysis reports

Historical analysis reports in the Query Patroller Center provide much useful information. These reports can be accessed by clicking the **Historical Analysis** folder, which expands to show the following folders:

- **Queries**
- **Tables Hit**
- **Tables Not Hit**
- **Indexes Hit**
- **Indexes Not Hit**
- **Submitters**

When data is available, the **Tables Hit** and **Submitters** folders can be further expanded to show more reports.

Most historical analysis reports can be viewed in table form and in graphical form. Use the **Show Table** button and the **Show Graph** or **Show Histogram** button to switch between the two views.

To sort a report (when you are looking at the table view), click the header of a column to sort the report by that column.

Table 10, Table 11 on page 120 and Table 12 on page 120 provide a list of information you might want to know about your data warehouse, what you could use the information for, and how to find that information in the historical analysis reports. The reports can help you answer such questions as "How long do queries spend in the queue before they are executed?" and "How many queries were processed last month?"

Table 10. Uses for historical analysis reports for performance

Historical data	Use to	Reports to reference
Average query execution time over time	Identify trends or variations in query performance; assess impact of hardware changes or warehouse redesign	Queries (sorted by Execution Time) Queries histogram (displaying Queries Run -- Average Time in the Y-axis) Note: Compare the results of these two reports over time.
Time spent in queue	Tune Query Patroller thresholds to minimize query execution time	Queries histogram (displaying Queries Queued -- Total Queue Time in the Y-axis)
Number of queries processed in a given time unit	Understand data warehouse activity	Queries histogram (displaying Queries Run -- Number Run in the Y-axis)
Variation in number of queries processed in a given unit of time	Identify periods of high and low system activity to optimize maintenance schedules or to perform dynamic configuration of DB2	Queries histogram (displaying Queries Run -- Number Run in the Y-axis)

Table 10. Uses for historical analysis reports for performance (continued)

Historical data	Use to	Reports to reference
Query cost	Identify inefficient or problematic queries	Queries (sorted by Estimated Cost)

Table 11. Uses for historical analysis reports for submitter activity

Historical data	Use to	Reports to reference
Number of queries submitted by a user	Identify heavy users in order to tune submitter thresholds; identify submitters who require education about submitting queries	Submitters (sorted by Submitter ID)
Size of queries submitted by a user	Identify submitters of problematic queries; tune submitter thresholds	Queries (sorted by Estimated Cost)
Object usage by submitter or application	Determine the optimal location for specific objects depending on which groups of submitters are using the objects	Submitters → Submitter X → Tables Hit Submitters → Submitter X → Tables Hit → Table X → Columns Hit Submitters → Submitter X → Tables Hit → Table X → Indexes Hit

Table 12. Uses for historical analysis reports for object usage

Historical data	Use to	Reports to reference
Database objects hit (tables, columns, indexes)	Identify candidates for materialized query tables, indexes, or multidimensional clustering	Tables Hit Indexes Hit Tables Hit → Table X → Columns Hit
Database objects not hit (tables, columns, indexes)	Identify objects to be eliminated or to move to slower access devices	Tables Not Hit Indexes Not Hit Tables Hit → Table X → Columns Not Hit

Chapter 20. Collecting historical data with Query Patroller

You can use Query Patroller's historical analysis functions to collect and analyze data about your data warehouse workload. You can either collect data on a test workload or collect data about your actual production activity. You can use the information you collect about a workload to gain greater understanding of database and object usage, set Query Patroller thresholds, and create query classes.

You must decide whether you will collect data about a test workload or your regular production activity. If you want to run a test workload, create a submitter and specify that all queries submitted by the test submitter profile be tracked for historical analysis. All other queries should not be intercepted. If you are planning to collect data regarding query execution time, you need to ensure that the DB2 timestamp and statement monitor switches are set to 'ON'. If you are planning to collect data about how many rows are returned by queries, you need to ensure that the DB2 statement monitor switch is set to 'ON'.

About Query Patroller historical analysis

To collect historical data, use the following method:

1. Open the Query Patroller System Properties notebook:
2. Enable Query Patroller to intercept queries
3. Start Query Patroller by issuing the `qpstart` command.
4. Enable the collection of historical data
5. Leave all Query Patroller thresholds at their default settings, *or* set the Query Patroller thresholds to unlimited to allow queries to run unfettered. Leaving thresholds at their default settings or setting them to unlimited means that you will not be using the full query management features of Query Patroller, but you will be collecting data about the database workload.
6. Run your test workload, or allow database activity to proceed as normal.
7. Generate historical data to populate the historical analysis tables.
8. Open the Query Patroller Center and use the reports in the Historical Analysis folder to analyze your data.

Chapter 21. Generating historical data using Query Patroller

Generate historical data when you want to have current information available for historical analysis. Generating historical data runs the SQL Explain facility against the queries that Query Patroller has saved for historical analysis, and makes the information available in the Historical Analysis reports and graphs in the Query Patroller Center.

You must have DBADM authority and SETSESSIONUSER privilege on PUBLIC to generate historical data. When you run the historical data generator for Query Patroller, if the Explain tables do not already exist, the generator creates them for you. However, it is strongly recommended that you create the Explain tables before running the historical data generator. When you create the Explain tables be sure you create them on the same database partition. Actively creating the Explain tables on the same database partition improves the performance of the Explain facility. This improvement increases the performance of the historical data generator. When the historical data generator creates explain tables it automatically creates them on the same database partition.

About Query Patroller historical analysis

Other than the **Query Activity Over Time** report, the **Historical Analysis** folder only contains information about the queries for which you have generated historical data. To decide whether you should generate historical data, you can determine when it was last generated, and for which queries.

If SQL Explain was run unsuccessfully on a query, it will not be rerun when historical data is next generated, even if that query falls within the specified time range.

To generate historical data for all outstanding queries (queries that have not yet had historical data generated for them) use one of the following methods:

- Query Patroller Center method
- Command line method

:

- To generate historical data using the Query Patroller Center method:
 1. From the Query Patroller Center, find the **Historical Analysis** folder in the object tree.
 2. Right-click the **Historical Analysis** folder and click **Generate Historical Data** in the pop-up menu. The Generate Historical Data window opens.
 3. Click the **Generate data for all outstanding queries** radio button.
 4. Click **OK** to generate historical data and close the window.
- To stop the generation of historical data using the Query Patroller Center method:
 1. Click **Terminate** in the Progress window. The Progress window opens after you click **OK** in the Generate Historical Data window.
- To generate historical data using the Command line method:
 1. Issue the `GENERATE HISTORICAL_DATA` command.
- To stop generation of historical data using the Command line method:

1. Issue the `GENERATE HISTORICAL_DATA` command using the `STOP` parameter.

Usage notes

1. It is strongly recommended that you run the historical data generator (using the `GENERATE HISTORICAL_DATA` command) during periods of minimal database usage. Running this command during these off-peak hours minimizes the risk of a performance impact on the database.
2. It is recommended that you generate historical data on a regular basis to reduce the number of queries that data is being collected for at one time.
3. If you run the historical data generator and shut it down in an abnormal way, you will receive an error the next time you attempt to run the historical data generator. Examples of abnormal shutdown include:
 - DB2 stops unexpectedly
 - Issuing a `db2stop force` command
 - Issuing a `killdb2` command

When the historical data generator shuts down abnormally, you must issue the following command before attempting to rerun the historical data generator:

```
qp -d database generate historical_data stop
```

where `database` identifies the database that the command is being run against.

Chapter 22. Determining when historical data was last generated

The reports and graphs under the **Historical Analysis** folder only contain information about queries that have had historical data generated for them.

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile has the HISTORICAL ANALYSIS privilege with view or edit authority

About Query Patroller historical analysis

To determine whether you should generate historical data, you can check:

- when historical data was last generated for all outstanding queries
- which queries have had historical data generated

To determine when historical data was last generated for all outstanding queries, use one of the following methods:

- Query Patroller Center method
- Command line method
- To determine when historical data was last generated using the Query Patroller Center method 1:
 1. From the Query Patroller Center, find the **Historical Analysis** folder in the object tree.
 2. Right-click the **Historical Analysis** folder and click **Generate Historical Data** in the pop-up menu. The **Generate Historical Data** window opens.
 3. The date and time that historical data was last generated is displayed in the **Date data was last generated for all outstanding queries** field.

Note: The **Date data was last generated for all outstanding queries** field in the **Generate Historical Data** window is not updated when you choose to generate data for a specific period of time. It is updated only when you choose to generate historical data for all outstanding queries.

 4. Click **Cancel** to close the **Generate Historical Data** window.
- To determine when historical data was last generated using the Query Patroller Center method 2:
 1. From the Query Patroller Center, find the **Historical Analysis** folder in the object tree.
 2. Click the **Historical Analysis** folder. The **Query Activity Over Time** report opens in the contents pane.
 3. The **Explain Run** column in the **Query Activity Over Time** report displays whether SQL Explain has been run on a query.
 - If the **Explain Run** column for a query shows a status of **Ran successfully**, historical data has been generated for that query, and will appear in the Historical Analysis reports and graphs. .
 - If the **Explain Run** column for a query shows a status of **Not yet run**, historical data has not been generated for that query.

- If the **Explain Run** column for a query shows a status of **Ran unsuccessfully** , historical data has not been generated for that query, and therefore will not appear in any historical analysis reports or graphs. To determine why it was unsuccessful, you can examine the qpuser.log and the qpdiaq.log files
- 4. Optional: You can sort the queries in the **Query Activity Over Time** report by clicking Explain Run .
- To determine when historical data was last generated using the Command line method:
 1. Issue the GET QP_SYSTEM command.:

Note: The command output provides details of when historical data was last generated, but it does not indicate which queries have had historical data generated. To determine which queries have had historical data generated you must use the Query Patroller Center method.

Chapter 23. Viewing historical query details using Query Patroller

Viewing the details of a historical query allows you to see information such as the SQL of the query, the processing time, user information, application information, and package information. By looking at the details of specific historical queries, you can examine information about which queries took a long time to complete and which queries were submitted at busy times.

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile has the HISTORICAL ANALYSIS privilege with view or edit authority

About Query Patroller historical analysis

To view the details of a historical query using the Query Patroller Center method:

1. Open the Historical Query Properties notebook.
To open the Historical Query Properties notebook for a particular query:
 - a. From the Query Patroller Center, expand the object tree under the **Historical Analysis** folder to find the **Queries** folder.
 - b. Click the **Queries** folder. The Query Activity over Time report is displayed in the pane on the right side of the window (the contents pane).
 - c. In the contents pane, right-click the query that you want information for, and click **Properties** in the pop-up menu. The Historical Query Properties notebook opens.
2. To view the SQL statement and general information about the submitter and query results, click the **General** tab. Click the **Show Access Plan** if you want to launch Visual Explain to see more information about the query.
3. To view query event time stamps and processing durations, click the **Time** tab.
4. To view package information, click the **Package** tab.
5. To view information about the statement, application, and end user, click the **Other** tab.
6. Click **Close** to close the Historical Query Properties notebook.

Chapter 24. Filtering tables for historical analysis using Query Patroller

When you are using the Query Patroller Center to look at historical analysis reports, there might be a significant number of rows in the Tables Hit and Tables Not Hit reports. Use the Filter window to display only the tables that meet conditions based on the name of the table and the schema.

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile has the HISTORICAL DATA privilege with view or edit authority

About Query Patroller historical analysis

The filter that you define will affect the number of objects that you see in the object tree and the number of reports that you see in the contents pane.

Query Patroller Center method

1. Open the Filter notebook.
2. Optional: On the **Locate** tab, select an operator for the table name comparison, and enter a value for the table name to be compared to.
3. Optional: Select an operator for the table schema comparison, and enter a value for the table schema to be compared to.
4. To indicate that the filter should show only tables that match both comparisons you have entered, click the **Meet all conditions** radio button. To indicate that the filter should show tables that match either the table name comparison or the table schema comparison, click the **Meet any conditions** radio button.
5. Optional: To view the WHERE clause that is generated by the table name and table schema comparisons you have entered, click the **Details** tab.
6. Optional: The Filter notebook will open automatically when a folder is selected and the number of objects listed exceeds the value specified in the **Object count** field, if you have selected the **Automatically display filter when object count is exceeded** check box. To adjust the number of objects that will cause the Filter notebook to open automatically, enter a new value in the **Object count** field.
7. To enable the filter, select the **Enable filter** check box, then click **OK** to close the Filter notebook.
8. Optional: If you do not want to enable the filter, clear the **Enable filter** check box, then click **OK** to close the Filter notebook. The values you have entered in the Filter notebook will be retained, and will appear when you reopen the notebook.

Chapter 25. Filtering queries for historical analysis using Query Patroller

When you are using the Query Patroller Center to look at historical analysis reports, there might be a significant number of rows in the Queries report. Use the Filter Queries notebook to display only the queries that meet conditions based on the query ID, the submitter ID, or whether SQL Explain has been run.

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile has the HISTORICAL DATA privilege with view or edit authority

About Query Patroller historical analysis

To filter queries for historical analysis using the Query Patroller Center method:

1. Open the Filter Queries notebook:
 - a. From the Query Patroller Center, expand the object tree under the **Historical Analysis** folder until you find the **Queries** folder.
 - b. Right-click the **Queries** folder and click **Filter** in the pop-up menu. The Filter Queries notebook opens.

The Filter Queries notebook will also open automatically when the **Queries** folder is selected, if the number of objects in the folder exceeds the object count specified in the Filter Queries notebook, and if the option to automatically display the filter is selected.
2. Optional: On the **Locate** tab, select an operator for the Query ID comparison, and enter a value to be compared to. The value must be numeric and must correspond to a query ID.
3. Optional: Select an operator for the Submitter ID comparison, and enter a value to be compared to. The value must be a character string and must correspond to a submitter ID.
4. Optional: Select an operator for the Explain Run comparison, and enter a value to be compared to. The value must be a character representing whether the Historical Analysis Data Generator has been run on this query. Possible values are:
 - N - Not yet run
 - S - Ran successfully
 - F - Run failed
5. To indicate that the filter should show only queries that match all the comparisons you have entered, click the **Meet all conditions** radio button. To indicate that the filter should show queries that match one or more of the comparisons, click the **Meet any conditions** radio button.
6. Optional: To view the WHERE clause that is generated by the comparisons you have entered, click the **Details** tab.
7. Optional: The Filter Queries notebook will open automatically when the Queries folder is selected and the number of objects listed exceeds the value specified in the **Object count** field, if you have selected the **Automatically display filter when object count is exceeded** check box. To adjust the number

of objects that will cause the Filter Queries notebook to open automatically, enter a new value in the **Object count** field.

8. To enable the filter, select the **Enable filter** check box, then click **OK** to close the Filter Queries notebook.
9. Optional: If you do not want to enable the filter, clear the **Enable filter** check box, then click **OK** to close the Filter Queries notebook. The values you have entered in the Filter Queries notebook will be retained, and will appear when you reopen the notebook.

The filter that you define will affect the number of queries that you see in the report in the contents pane.

Chapter 26. Viewing index details using Query Patroller

When you are examining Query Patroller historical queries, you might want more information about the indexes that were used by those queries, as well as information about the indexes that were not used. Viewing the index details allows you to see the columns that the index is based on, the index definer, the index type, whether the index is required by the system, and more.

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile has the HISTORICAL ANALYSIS privilege with view or edit authority

About Query Patroller historical analysis

The information displayed in the Index Properties window comes from DB2 system catalog tables. The SYSCAT.INDEXES table is the source for the information in the following fields of the Index Properties window:

- Index name
- Index schema
- Table name
- Table schema
- Index definer
- Index type
- Index defined by user
- Index required by system
- Comments

The DB2 SYSCAT.INDEXCOLUSE table is the source for the information in the **Columns** table in the Index Properties window.

To view index details using the Query Patroller Center method:

1. Open the Index Properties window. Index details are displayed in the Index Properties window.
 - To open the Index Properties window for indexes that are being used:
 - a. From the Query Patroller Center, expand the object tree under the **Historical Analysis** folder to find the **Indexes Hit** folder.
 - b. Click the **Indexes Hit** folder. The **Indexes Hit** report is displayed in the pane on the right side of the window (the contents pane).
 - c. In the contents pane, right-click the index for which you would like to obtain more information, and click **Properties** in the pop-up menu. The Index Properties window opens.
 - To open the Index Properties window for indexes that are *not* being used:
 - a. From the Query Patroller Center, expand the object tree under the **Historical Analysis** folder to find the **Indexes Not Hit** folder.
 - b. Click the **Indexes Not Hit** folder. The **Indexes Not Hit** report is displayed in the pane on the right side of the window (the contents pane).

- c. In the contents pane, right-click the index for which you would like to obtain more information, and click **Properties** in the pop-up menu. The Index Properties window opens.
2. Optional: Click on a column heading in the **Columns** table to sort by that heading.
3. Click **Close** to close the Index Properties window.

Part 6. Using Query Patroller with other DB2 components

Using the DB2 governor with Query Patroller

The main function of Query Patroller is to help database administrators manage queries against a database. The main function of the DB2 governor is to help administrators manage applications running against a database.

Important: With the new workload management features introduced in DB2 Version 9.5, Query Patroller and its components have been deprecated in Version 9.7 and might be removed in a future release. For more information, see the “DB2 Governor and Query Patroller have been deprecated” topic in the *What's New for DB2 Version 9.7* book.

The DB2 governor allows you to place limits on resources such as the number of locks, the amount of idle time, and the amount of CPU used by an application. The DB2 governor can be used together with Query Patroller to provide a tremendous amount of administrative control. However, in order to use them together effectively, you must understand how they interact.

Query Patroller is a system of collaborative applications that run against a database. Since the governor can act against these applications the way it does against other applications, there are some guidelines to follow when specifying rules in the governor configuration file.

Specifically, it is important to avoid including the processes used by Query Patroller in the rules of the governor configuration file. Query Patroller uses javaw.exe, java.exe, db2fmp.exe, and qp.exe on Windows and java, db2fmp, and qp on UNIX to carry out its operations. To prevent the governor from acting against Query Patroller, do not include these processes in the governor configuration file. You should also make sure you do not have a general rule that intercepts all applications by default. Instead, explicitly include the list of applications to be intercepted by the DB2 governor.

Note: In addition to rules acting against Query Patroller processes, there may be other rules in the governor configuration file that can cause Query Patroller to be intercepted.

If the processes used by Query Patroller cannot be excluded from interception by the DB2 governor, then you should follow the following guidelines in writing your governor configuration file rules.

- The Query Controller uses javaw.exe and db2fmp.exe on Windows, and java and db2fmp on UNIX. If you must intercept these processes, set the rowsel and rowsread limits to a value that is considerably higher than the larger of the number of records in SYSCAT.DBAUTH where (DBADMAUTH='Y' and GRANTEETYPE='U') and the number of records where (DBADMAUTH='Y' and GRANTEETYPE='G'). If the DB2 governor still intercepts the Query Controller, retry setting the rowsel and rowsread limits with higher values.
- Query Patroller Center uses javaw.exe, java.exe and db2fmp.exe on Windows, and java and db2fmp on UNIX. If you must intercept these processes, set the rowsel and rowsread limits to a value that is considerably higher than the larger of the number of records in SYSCAT.DBAUTH where (DBADMAUTH='Y'

and GRANTEETYPE='U') and the number of records where (DBADMAUTH='Y' and GRANTEETYPE='G'). Also, note that if rowsel and rowsread limits are smaller than the number of records in the largest Query Patroller control table, Query Patroller Center or Query Patroller Command Line Processor will be acted upon by the DB2 governor. There is no way to predetermine the maximum number of records in the Query Patroller control tables as they are mostly dynamic. If desired, increase the rowsel and rowsread limits to values that are higher than the current maximum number of records or set them to unlimited. Restrictiveness of other limits, including CPU usage and idle time, may also cause the DB2 governor to act against Query Patroller processes depending on the amount of time and resources used by Query Patroller to operate on Query Patroller control tables. Once again, this amount cannot be predetermined as it depends on hardware capability and data size. If desired, increase the limits to higher values to prevent the DB2 governor from acting against Query Patroller.

Effects of DB2 governor actions on Query Patroller processes

If the priority or schedule actions apply to the Query Patroller processes, Query Patroller will continue to run with reduced system resources. However, if a force action applies to a Query Patroller process, the process can be terminated. The force action might terminate a Query Patroller process normally, returning a SQL1224N return code, or it might cause an application error or an abnormal termination of the DARI process (SQL1131N) if the db2fmp process was started before the force was issued. Query Patroller cannot stop the db2fmp process after it has started. The db2fmp process will try to complete execution even after Query Patroller has shut down the database connection that the db2fmp process requires for successful execution.

For more information about the db2fmp process, see the *Application Development Guide: Programming Client Applications*.

Running Query Patroller and the DB2 governor against the same submitter applications

Both Query Patroller and the DB2 governor can be used against the same query submitter applications. For example, a submitting application such as DB2 CLP (db2bp.exe on Windows and db2bp on UNIX) can be listed as an application that is intercepted by Query Patroller as well as included in the governor configuration file.

Query Patroller intercepts queries at submission time, while the DB2 governor intercepts applications at query execution time. Since query submission occurs before query execution, Query Patroller will always intercept queries before the DB2 governor does. This means that if Query Patroller holds or queues a query, the DB2 governor must wait until the query is executed before intercepting the application that submitted the query.

A query that is intercepted by Query Patroller can be executed either by the submitting application or by another application called qprunquery.exe on Windows and qprunquery on UNIX. If the submitter's submission preferences specify that the submitting application has to wait until the query results are returned before releasing the application, then it is the submitting application that will execute the query. If the submitting application is listed in the DB2 governor configuration file, then the DB2 governor will intercept the submitting application when it executes the query.

If the submitter's submission preferences specify that the submitting application should be released and the query results sent to a result table, then the query is executed by qprunquery. In this situation, the DB2 governor will only intercept the application if qprunquery is included in the DB2 governor configuration file.

Using Query Patroller with the DB2 connection concentrator

When Query Patroller places a query in the queue, that query blocks the application the entire time it is in the queue, until the query is run.

When the DB2 connection concentrator is not activated, every application gets its own agent to manage the database connection until the application disconnects. When the concentrator is activated, all applications share a pool of agents which are switched between applications on transaction boundaries. This means that if the concentrator is enabled and Query Patroller queues queries, it ties up those agents until the queries are run. This would have the effect of reducing the pool of available agents and affect the performance of DB2 since applications would not be able to connect or execute a request due to their inability to acquire the services of an agent. For this reason, when the connection concentrator is activated, Query Patroller will not queue queries; instead, it will, by default, reject the queries which are supposed to be queued with an sqlcode 29009, reason code 6.

To prevent queries from being rejected when they are chosen to be queued, you can choose to allow Query Patroller to run queries instead of rejecting them when the concentrator is activated by setting the option BLOCK_OPTION at the system level, using the UPDATE QP_SYSTEM command, or at the user level, using the UPDATE SUBMITTER_PROFILE command. By default, BLOCK_OPTION is set to 'reject' ('R') specifying that queries are to be rejected rather than queued when the concentrator is activated. To specify that Query Patroller should run queries rather than reject them when the concentrator is activated, set BLOCK_OPTION to 'proceed' ('P').

For example, to allow Query Patroller to run queries against the database "sample" that would otherwise be rejected when the concentrator is activated, set BLOCK_OPTION option to 'P' as follows:

```
qp -d sample -u userid -p password "UPDATE QP_SYSTEM USING BLOCK_OPTION 'P'"
```

To allow Query Patroller to run queries submitted under the profile "STEVED" that would otherwise be rejected when the concentrator is activated, set BLOCK_OPTION for this profile to 'P' as follows:

```
qp -d sample -u userid -p password "UPDATE SUBMITTER_PROFILE  
FOR USER'STEVED' USING BLOCK_OPTION 'P'"
```

The values for BLOCK_OPTION are stored in the QP_SYSTEM and SUBMITTER_PROFILE tables for the database.

The BLOCK_OPTION setting for QP_SYSTEM is not nullable; the BLOCK_OPTION setting for SUBMITTER_PROFILE is nullable. If the BLOCK_OPTION is set both for QP_SYSTEM and for a user's submitter profile, the value for the submitter profile takes precedence for that user. For all other users, the BLOCK_OPTION setting for QP_SYSTEM applies. To ensure that the BLOCK_OPTION setting for QP_SYSTEM applies to a particular user, set the BLOCK_OPTION for that user's SUBMITTER_PROFILE to NULL.

Part 7. Tuning Query Patroller

From time to time you may find that you experience some problems with your query workload. The following table describes some of the problems that you may encounter and suggests some possible causes that you can explore to remedy the situation.

Table 13. List of query management problems and possible causes

Problem	Possible cause
Too many queries are being held.	Maximum query cost for submitter profiles (MAX_COST_ALLOWED) is set too low.
Queries are queued too long or too often.	Maximum number of queries for submitters (MAX_QUERIES_ALLOWED) or the system (MAX_TOTAL_QUERIES) is set too low. If you have query classes defined in your system, you may have set the maximum number of queries (MAX_QUERIES) too low for one or more of your query classes or you may have created too many query classes.
Small queries are running too slowly.	Submitter's minimum query cost to manage (MIN_COST_TO_MANAGE) is set too low.
Scheduled jobs are not completing overnight.	The setting for the duration of time to run held queries (RUN_HELD_DURATION) is too low.
Performance of time-critical queries is impacted after installing Query Patroller.	Consider allowing certain applications or submitters to bypass Query Patroller interception.
Too many queries are running unmanaged.	Submitter's minimum query cost to manage (MIN_COST_TO_MANAGE) is set too high or certain applications are not being intercepted by Query Patroller that should be intercepted.
Query cost estimates seem inaccurate.	Database statistics not accurate. Ensure that RUNSTATS is run periodically and after all significant database changes.
Database performance is impacted when using submitter profiles to bypass queries.	Instead of using submitter profiles to bypass queries, use one or more of the Query Patroller registry variables (DB2_QP_BYPASS_APPLICATIONS, DB2_QP_BYPASS_USERS or DB2_QP_BYPASS_COST).

Part 8. Query Patroller commands

Chapter 27. Query Patroller command line support

Query Patroller commands can be invoked from a command line window, preceded by the keyword `qp`. The general syntax for entering Query Patroller commands is described below.

Note: For languages other than English, Query Patroller commands must be run from a DB2 CLP window for the correct codepage to apply.

Command syntax

```
qp [-u username -p password] -d database qp-command
```

Command parameters

-u username

Specifies the user ID username under which to connect to the database and execute the Query Patroller command.

-p password

Specifies the password used to authenticate the user name. If the password is omitted, the user is prompted to enter it.

-d database

Identifies the database that the command is being run against.

qp-command

Specifies the Query Patroller command string to be executed. The Query Patroller command string can consist of a single Query Patroller command plus any necessary options.

The Query Patroller commands are listed under this topic in the Information Center navigation tree.

You can also view the list of all Query Patroller commands by entering one of the following commands:

```
qp -d db-name ?
```

or

```
qp -d db-name help
```

where `db-name` is the name of a database.

If the client running the Query Patroller command is a remote client, then the user and password options are also required, as follows:

```
qp -d db-name -u user_id -p password ?
```

or

```
qp -d db-name -u user_id -p password help
```

Examples

The following command cancels query 854 running against the SAMPLE database:

```
qp -d sample cancel query 854
```

The following command, issued under the username "testuser," lists all of the query classes defined for the TESTDB database: `qp -u testuser -p testpw -d testdb list query_classes`

Usage Notes

1. The `qp` command line returns an exit code of 0 when the command is successful, and an exit code of -1 if there is a failure. On Linux or UNIX platforms where exit code is limited to a 8-bit value, the failure exit code is effectively 255.
2. All parameter values whose corresponding SQL types are char or varchar must have single quotes around them. If the parameter values themselves contain a single quote, an escape character is required. For example, the username Mike O'Connell must be specified as 'Mike O'Connell'.
3. On UNIX platforms, if a `qp` command string contains any single quote, the entire `qp` command must be surrounded by double quotes. For example, to create a new submitter profile for "testuser" you would enter the command: `qp -d wsdb "add submitter_profile 'TESTUSER' using default"`
4. A database connection is established when the command is executed.

Chapter 28. ADD OPERATOR_PROFILE

Adds a new operator profile to the set of Query Patroller operator profiles defined in the OPERATOR_PROFILE table.

Authorization

You must meet one of the following requirements:


- Have DBADM authority

Required connection

None. This command establishes a database connection.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see Chapter 27, "Query Patroller command line support," on page 143



Command parameters

USER username

Specifies the username for which the operator profile is created. The user ID username must also exist as a DB2 authorization ID.

GROUP groupname

Specifies the group name for which the operator profile is created. The groupname must be a DB2 user ID and be an existing ID at the operating system or Kerberos level.

DEFAULT

Specifies that the operator profile should be created with the default values for all parameters. For any individual parameter that has a default value, the parameter can be set to the default by entering the parameter with DEFAULT as the value. For example, to set the MONITORING privilege for user A to the default value, enter the following:

```
qp -d sample "ADD OPERATOR_PROFILE FOR USER 'USERA' USING  
MONITORING DEFAULT"
```

LIKE USER username

Specifies that the new operator profile is to have the same settings as the operator profile for USER username.

LIKE GROUP groupname

Specifies that the new operator profile is to have the same settings as the operator profile for GROUP groupname.

Setting-keyword**CONFIGURATION**

Specifies the authority level of the CONFIGURATION privilege to be assigned to the operator.

- E** Allows the operator to add, edit and view Query Patroller query classes and system settings.
- V** Allows the operator to view Query Patroller query classes and system settings.
- N** Renders the operator unable to view or edit Query Patroller query classes and system settings. This is the default value.

HISTDATA

Specifies the authority level of the HISTORICAL DATA privilege to be assigned to the operator.

- E** Allows the operator to view and remove historical data.

Note: You must have DBADM authority and SETSESSIONUSER privilege on PUBLIC to generate historical data.

- V** Allows the operator to view historical data.
- N** Renders the operator unable to view or remove historical data. This is the default value.

MONITORING

Specifies the authority level of the MONITORING privilege to be assigned to the operator.

- E** Allows the operator to view and manage queries. This includes the ability to monitor, remove and change the status of managed queries, view query details, and delete result tables.
- V** Allows the operator to view the details of queries managed by Query Patroller.
- N** Renders the operator unable to view or manage queries. This is the default value.

USERADMIN

Specifies the authority level of the USER ADMINISTRATION privilege to be assigned to the operator.

- E** Allows the operator to edit and view Query Patroller submitter profiles and submission preferences, and to view operator profiles.
- V** Allows the operator to view Query Patroller operator and submitter profiles and other users' submission preferences.
- N** Renders the operator unable to edit or view Query Patroller operator and submitter profiles and other users' submission preferences. This is the default value.

SUSPENDED

Indicates whether or not privileges for this operator profile are suspended.

N Privileges are not suspended for this operator profile. This is the default value.

Y Privileges are suspended for this operator profile.

Examples

The following command creates an operator profile for the group HELPDESK on the sample database. The helpdesk group will be administering users and managing queries, but they do not need to view or edit system settings or view historical data. The helpdesk group operator profile gives members monitoring privilege with edit authority, and the user administration privilege with edit authority.

Usage notes

```
qp -d sample "ADD OPERATOR_PROFILE FOR GROUP 'HELPDESK' USING HISTDATA 'N'  
CONFIGURATION 'N' MONITORING 'E' USERADMIN 'E' SUSPENDED 'N'"
```

1. Operator profiles do not need to be created for users with DBADM authority on a database. Such users already possess the maximum level of operator privileges, therefore adding operator profiles for them is redundant. It may also be misleading to create an operator profile for a user with DBADM authority since the user can automatically perform all Query Patroller tasks despite any restrictions on the operator privileges associated with the profile.

Chapter 29. ADD QUERY_CLASS

Adds a new query class to the list of query classes defined for the database.

Authorization

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile includes CONFIGURATION privilege with edit authority.

Required connection

None. This command establishes a database connection.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see . Query Patroller command line support.

```
▶▶—ADD QUERY_CLASS—USING—————▶▶
▶—MAX_QUERIES—max-number-of-queries—MAX_COST—max-query-cost————▶▶
▶—DESCRIPTION—'description'—▶▶
```

Command parameters

MAX_QUERIES *max-number-of-queries*

The maximum number of queries that can run simultaneously for this query class. When this threshold of running queries is reached, additional queries for this query class are queued until resources become available. The value must be greater than 0 and less than or equal to the value of MAX_TOTAL_QUERIES specified in the Query Patroller system settings.

MAX_COST *max-query-cost*

The maximum cost for a single query that this query class will accept. The value must be less than or equal to the value of MAX_TOTAL_COST specified in the Query Patroller System settings.

Note: All query classes defined in a system must have distinct values for MAX_COST.

DESCRIPTION *description*

Specifies a description for the query class. This description must be 256 characters or less enclosed in single quotes. This parameter is nullable.

Examples

The following example creates a new query class for the SAMPLE database. A maximum of 10 queries will be allowed to run simultaneously in this query class, and the maximum size for any one query will be 1000 timerons.

```
qp -d sample "ADD QUERY_CLASS USING MAX_QUERIES 10 MAX_COST 1000  
DESCRIPTION 'This query class runs small sales queries.'"
```

Usage notes

1. You can create a maximum of 99 query classes in a system.
2. Query classes can be created, removed, or modified while Query Patroller is started. Creation, changing the maximum query cost, or removal of a query class will take effect immediately unless there are queued or running queries. If there are queued or running queries, including newly submitted queries, the query class changes will take effect when these complete. If you do not want to wait for all queued and running queries to complete, a Query Patroller server restart is required. Updating the maximum number of queries for a query class always takes effect immediately.

Chapter 30. ADD SUBMISSION_PREFERENCES

Creates a submission preferences file for a specified submitter.

Authorization

You must meet one of the following requirements:


- Have DBADM authority
- Be an operator whose profile includes the USER ADMINISTRATION privilege with edit authority
- Be the owner of the submission preferences file being created.

Required connection

None. This command establishes a database connection.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see Chapter 27, “Query Patroller command line support,” on page 143.



Command parameters

USER username

Specifies the username of the submitter for whom the submission preferences file is created. If no username is specified, the current login ID is used. The user ID *username* must also exist as a DB2 authorization ID.

LIKE USER username

Specifies that the new submission preferences file is to have the same settings as the submission preferences for USER *username*.

DEFAULT

Specifies that new submission preferences be created for the specified user with default values for all parameters (for example, `RESULT_EXCEEDED_ACTION` will be 'A', `RESULT_ACCESSIBILITY` will be 'S', and so on).

parameter-keyword value

The following parameters can be set using this command:

SUBMITTER_PROFILE_NAME profile-name

SUBMITTER_PROFILE_TYPE profile-type

Specifies the group profile to use to establish submitter settings,

such as thresholds, if the submitter has no profile of type user. The value of SUBMITTER_PROFILE_TYPE must be 'G'. If multiple group profiles exist for this submitter and no group profile is specified here, queries will be submitted using the group profile with the lowest resource thresholds.

RESULT_EXCEEDED_ACTION

Specifies the action to take when the query results to be stored in a result table exceed the maximum result rows specified in the submitter's profile.

- 'A' Specifies that no results are stored in the result table if the result set is longer than the limit defined in the submitter's profile. This option is the default.
- 'T' Specifies that a truncated set of results is stored in the result table if the result set is longer than the limit defined in the submitter's profile.

RESULT_ACCESSIBILITY

Specifies whether the result table containing query results will be available to more users than just the submitter.

- 'O' Specifies that the result table is accessible by the DB2 IDs listed in the value of the OTHER_GRANTEES parameter. The DB2 IDs listed must have access to the database where the query was submitted.
- 'S' Specifies that access to the result table is restricted to the query submitter. This option is the default.

OTHER_GRANTEES grantees

Specifies the DB2 user IDs or group IDs that can access the result table. Up to 1024 alphanumeric characters are accepted. Multiple IDs must be separated with commas.

RESULT_DESTINATION result-destination-id

Specifies whether the submitting application will wait for query results to return, or will be freed up for further activity.

- 'A' Specifies that the application that submitted the query will wait for the result set to return while Query Patroller manages the query. When this option is selected, the application that submitted the query may become unresponsive until the result set is returned. This option is the default.
- 'T' Specifies that the result set is stored in a DB2 table. When the query is submitted, the application that submitted the query becomes free for further processing.

EMAIL_ADDRESSES email-addresses

Specifies the e-mail address or addresses to receive notification regarding queries submitted by this submitter.

Note: This notification only applies if e-mail notification is enabled in the QP_SYSTEM settings.

The value for this parameter may be up to 1024 characters. Multiple e-mail addresses must be separated by a comma.

Chapter 31. ADD SUBMITTER_PROFILE

Adds a new submitter profile to the SUBMITTER_PROFILE table.

Authorization

You must meet one of the following requirements:


- Have DBADM authority
- Be an operator whose profile includes the USER ADMINISTRATION privilege with edit authority.

Required connection

None. This command establishes a database connection.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see Query Patroller command line support.



Command parameters

USER username

Specifies the name of the user associated with the profile. The user ID *username* must also exist as a DB2 authorization ID.

GROUP groupname

Specifies the name of the group associated with the profile. The group specified must be a DB2 group ID and exist as a group at the operating system level or Kerberos level.

USING DEFAULT

Specifies that the submitter profile should be created with the default values for all parameters. For any individual parameter that has a default value, the parameter can be set to the default by entering the parameter with `DEFAULT` as the value. For example, to set the `PRIORITY` for `USERA` to the default value, enter the following:

```
qp -d sample "ADD SUBMITTER_PROFILE FOR USER 'USERA' USING PRIORITY  
DEFAULT"
```

LIKE USER username

Specifies that the new submitter profile is to have the same settings as the submitter profile for `USER username`.

LIKE GROUP *groupname*

Specifies that the new submitter profile is to have the same settings as the submitter profile for GROUP *groupname*.

profile-parameter

Specifies the parameter values to be assigned to the profile. The following parameters can be set:

PRIORITY *priority*

Specifies the priority level assigned to queries submitted under this profile. The higher the value for this parameter, the higher the priority that is assigned to the submitter's queries.

- Values must be an integer between 0 and 999 inclusive
- The default value is 500

MAX_QUERIES_ALLOWED *max-queries*

Specifies the maximum number of queries that a submitter is able to run simultaneously. Queries submitted after this limit is reached are placed in a queued state until other submitted queries complete. When creating a submitter profile for a group, note that the value set for this parameter applies to each user. For example, if this value were set to 10 for Group A, then each user belonging to Group A has the authority to run 10 queries simultaneously.

- A value of "-1" indicates that users with this profile can have unlimited queries running simultaneously (up to the value of MAX_TOTAL_QUERIES specified in the QP_SYSTEM table.)
- The default value for this parameter is 100.

MAX_RESULT_ROWS *max-number-of-result-rows*

Specifies the maximum number of result rows that will be stored in a result table for a single query submitted under this profile. Only queries whose results are to be stored in a result table are subject to this limit.

- A value of "-1" indicates that users with this profile can store results with as many rows as necessary to accommodate the complete result set.
- The default value for this parameter is 1,000,000 rows.

MAX_COST_ALLOWED *max-query-cost*

Specifies the maximum query cost for a submitter under this profile. If the estimated cost of a query submitted under this profile exceeds this value, the query is placed in a held state.

- A value of "-1" indicates that users with this profile can run queries of any size (up to the value of MAX_TOTAL_COST specified in the QP_SYSTEM table.)
- The default value for this parameter is 10,000,000.

MIN_COST_TO_MANAGE *min-query-cost*

Specifies the minimum cost of a query managed by Query Patroller. Queries whose estimated cost is lower than this value are not managed by Query Patroller. A query that falls below this minimum cost is still tracked for historical analysis, provided the value of the QUERIES_TO_SAVE parameter in the QP_SYSTEM table is set to A (all queries). The default value is 15,000.

ACCOUNT_ID *account-id*

Specifies an alphanumeric ID to use for account tracking purposes.

Up to 128 characters will be accepted. You can use this parameter to sort submitters into logical groupings to track usage costs. This parameter is nullable.

SUSPENDED Y/N

Specifies whether or not a submitter is prohibited from submitting queries. The default value for this parameter is "N". The character value for this parameter must be enclosed in single quotes.

INTERCEPT Y/N

Specifies that Query Patroller is to intercept or manage queries submitted by this submitter. If queries are not intercepted, Query Patroller will not evaluate the cost of the query or track it for historical analysis. The default value for this parameter is "Y". The character value for this parameter must be enclosed in single quotes.

Examples

The following example adds a submitter profile for the `MARKETING` group to use the `SALES` database. Because users in this group tend to submit large queries, the profile defined for this group allows queries with a large estimated cost and allows a large number of result rows to be stored in result tables. Users in this group also tend to submit fewer queries at once, so the value for `MAX_QUERIES_ALLOWED` is set to 100. The remaining settings for this profile are left as default values.

```
qp -d sales "ADD SUBMITTER_PROFILE FOR GROUP 'MARKETING' USING
MAX_QUERIES_ALLOWED 100 MAX_RESULT_ROWS 1000000000 MAX_COST_ALLOWED
20000000"
```

Usage notes

1. To print or view all of the settings for a particular submitter profile, use the `GET SUBMITTER_PROFILE` command.
2. To list all of the submitter profiles currently defined for a database, use the `LIST SUBMITTER_PROFILES` command.

Chapter 32. CANCEL QUERY

Cancels the specified query.

Authorization

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile includes MONITORING privilege with edit authority
- Be the submitter of the query to be canceled.

Required connection

None. This command establishes a database connection.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see Chapter 27, "Query Patroller command line support," on page 143.

►►—CANCEL QUERY—*query-id*—◄◄

Command parameters

QUERY *query-id*

Specifies the ID of the query to be canceled.

This command cancels managed queries only; it cannot cancel queries that are tracked but not managed.

Chapter 33. GENERATE HISTORICAL_DATA

Collects data from the TRACK_QUERY_INFO table, runs SQL Explain facility against this data, and populates the QUERY_ANALYSIS table with the results. These results are used to generate the Historical Analysis reports and graphs in the Query Patroller Center.

Authorization

You must meet all of the following requirements:

- Have DBADM authority
- Have SETSESSIONUSER privilege on PUBLIC

If you have DBADM authority but do not have SETSESSIONUSER privilege on PUBLIC, the command will proceed but it will fail to process most queries.

Required connection

None. This command establishes a database connection.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see Chapter 27, “Query Patroller command line support,” on page 143.

```
▶▶—GENERATE HISTORICAL_DATA—▶▶
    —STOP—
    —FROM START—start-time—END—end-time—▶▶
```

Command parameters

STOP Stops the generation of historical data. When this option is specified, an update is made to the QP_SYSTEM control table indicating that a stop command has been issued. When the historical data generator is running, it checks the QP_SYSTEM table every 20 queries to see if a stop command has been issued. If so, generation of historical data will stop.

FROM START *start-time* **END** *end-time*

Specifies that historical data should be generated for database activity between *start-time* and *end-time*. The format for the *start-time* and *end-time* values is 'YYYY-MM-DD HH24:MI:SS' where:

- YYYY represents a 4-digit year value
- MM represents a 2-digit month value from 1–12 (for example, January is represented by '01')
- DD represents a 2-digit day of the month value from 1–31
- HH24 represents an hour of the day value from 00–24

Note: If the value for the hour is '24' the minutes and seconds values must be '00'.

- MI represents a 2-digit minutes value between 00–59

- SS represents a 2-digit seconds value between 00–59.

Usage Notes

1. The results of this command can be viewed in the Query Patroller Center historical analysis reports and graphs.
2. It is strongly recommended that you run the historical data generator during periods of minimal database usage. Running this command during these off-peak hours minimizes the risk of a performance impact on the database.
3. It is recommended that you run the `GENERATE HISTORICAL_DATA` command on a regular basis to reduce the number of queries that data is collected for at one time.
4. When running the historical data generator for Query Patroller, if the Explain tables do not already exist, the generator will create them for you. However, it is strongly recommended that you create the Explain tables before running the historical data generator. When you create the Explain tables be sure you create them on the same partition. Actively creating the Explain tables on the same partition improves the performance of the Explain facility. This improvement increases the performance of the historical data generator. When the historical data generator creates explain tables it automatically creates them on the same partition.
5. If you do not specify `FROM START start-time END end-time` (that is, if you just specify `GENERATE HISTORICAL_DATA` without any options), then historical data is generated for all queries that were run between the last time the `GENERATE HISTORICAL_DATA` command was run without the `FROM START start-time END end-time` being specified and the current time.

Usage Notes

1. When saving or filing the results of a query, a BLOB column value will not be saved. Instead it will be replaced with the keyword, "BLOB". A CLOB value will be truncated if it is greater than 32K in size.

Chapter 35. GET OPERATOR_PROFILE

Retrieves the settings for a specified Query Patroller operator profile.

Authorization

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile includes the USER ADMINISTRATION privilege with view or edit authority.

Required connection

None. This command establishes a database connection.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see Chapter 27, “Query Patroller command line support,” on page 143.

```
▶▶ GET OPERATOR_PROFILE FOR USER 'username' GROUP 'groupname' ▶▶
```

Command parameters

USER *username*

Specifies the name of the user whose operator profile is to be retrieved.
The user ID *username* must also exist as a DB2 authorization ID.

GROUP *groupname*

Specifies the name of the group whose operator profile is to be retrieved.
This group name must also exist as a DB2 authorization ID.

Examples

The following command retrieves the profile settings for user jsmith for the TESTDB database:

```
qp -d testdb "GET OPERATOR_PROFILE FOR USER 'JSMITH'"
```

Chapter 36. GET QP_SYSTEM

Retrieves the system settings for a Query Patroller enabled database.

Authorization

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile includes the CONFIGURATION privilege with edit or view authority.

Required connection

None. This command establishes a database connection.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see Chapter 27, “Query Patroller command line support,” on page 143.

▶▶—GET QP_SYSTEM—▶▶

Chapter 37. GET QUERY

Retrieves the details of a specified query.

Authorization

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile includes the MONITORING privilege with view or edit authority
- Be the submitter of the query to be retrieved.

Required connection

None. This command establishes a database connection.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see Chapter 27, “Query Patroller command line support,” on page 143.

```
▶▶ GET QUERY query-id [SHOW FULL QUERY]
```

Command parameters

query-id

Specifies the ID of the query to be retrieved.

This command retrieves information about managed queries only; it cannot retrieve information about queries that are tracked but not managed.

SHOW FULL QUERY

Specifies that the full query text is to be displayed or printed. By default, only the first 1KB of query text is returned unless this keyword is specified.

Chapter 38. GET QUERY_CLASS

Retrieves the settings for a specified Query Patroller query class.

Authorization

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile includes CONFIGURATION privilege with view or edit authority.

Required connection

None. This command establishes a database connection.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see Chapter 27, “Query Patroller command line support,” on page 143.

▶▶—GET QUERY_CLASS—*query-class-id*—————▶▶

Command parameters

query-class-id

The ID for the query class to be retrieved.

Chapter 39. GET SUBMISSION_PREFERENCES

Retrieves the submission preferences for a specified Query Patroller user.

Authorization

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile includes the USER ADMINISTRATION privilege with view or edit authority
- Be the owner of the submission preferences being retrieved.

Required connection

None. This command establishes a database connection.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see Chapter 27, “Query Patroller command line support,” on page 143.

```
▶▶ GET SUBMISSION_PREFERENCES [FOR USER 'username'] ▶▶
```

Command parameters

FOR username

The name of the user whose submission preferences are to be retrieved. The user ID *username* must also exist as a DB2 authorization ID. If this is not specified, then the submission preferences associated with the current user are returned.

Usage notes

1. If the specified username does not have any submission preferences defined, this command returns the values for the default PUBLIC submission preferences.

Chapter 40. GET SUBMITTER_PROFILE

Retrieves the setting details for a specified Query Patroller submitter profile.

Authorization

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile includes the USER ADMINISTRATION privilege with edit or view authority.

Required connection

None. This command establishes a database connection.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see Chapter 27, “Query Patroller command line support,” on page 143.

```
▶▶ GET SUBMITTER_PROFILE FOR USER 'username'
                             GROUP 'groupname' ▶▶
```

Command parameters

USER *username*

Specifies the name of the user whose submitter profile is to be retrieved. The user ID *username* must also exist as a DB2 authorization ID.

GROUP *groupname*

Specifies the name of the group whose submitter profile is to be retrieved. The group name must also exist as a DB2 authorization ID.

Chapter 42. LIST QUERIES

Lists the queries for a specified database. This list is displayed in descending order by query id.

This command lists managed queries only; it does not list queries that are tracked but not managed.

Authorization

You must meet one of the following requirements:


- Have DBADM authority
- Be an operator whose profile includes the MONITORING privilege with edit or view authority
- Be the submitter who owns the queries being listed.

Required connection

None. This command establishes a database connection.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see Chapter 27, "Query Patroller command line support," on page 143.



Command parameters

FOR USER ALL

Specifies that managed queries for all users are listed.

FOR USER *username*

Specifies that queries submitted by a particular submitter are listed. If the USER parameter is not specified, the default is to list the queries belonging to the current user.

FOR QUERY_CLASS *query-class-id*

Specifies that queries running in the query class identified by *query-class-id* are listed.

WITH STATUS status

Specifies that queries with a particular status are listed. The values for this parameter may be any of the following:

- A** Specifies that aborted queries are listed.
- C** Specifies that cancelled queries are listed.
- D** Specifies that completed ("done") queries are listed.
- U** Lists the queries whose status is unknown.
- H** Specifies that held queries are listed.
- J** Specifies that rejected queries are listed.
- L** Specifies that released queries are listed.
- Q** Specifies that queued queries are listed.
- R** Specifies that running queries are listed.

CREATED BETWEEN start AND stop

Specifies that queries created between the specified start and stop time are listed. The format for the start and stop time is 'YYYY-MM-DD HH24:MI:SS' where:

- YYYY represents a 4-digit year value
- MM represents a 2-digit month value from 1–12 (for example, January is represented by '01')
- DD represents a 2-digit day of the month value from 1–31
- HH24 represents an hour of the day value from 00–24

Note: If the value for the hour is '24' the minutes and seconds values must be '00'.

- MI represents a 2-digit minutes value between 00–59
- SS represents a 2-digit seconds value between 00–59.

COMPLETED BETWEEN start AND stop

Specifies that queries completed between the specified start and stop time are listed. The format for the start and stop time is 'YYYY-MM-DD HH24:MI:SS' where:

- YYYY represents a 4-digit year value
- MM represents a 2-digit month value from 1–12 (for example, January is represented by '01')
- DD represents a 2-digit day of the month value from 1–31
- HH24 represents an hour of the day value from 00–24

Note: If the value for the hour is '24' the minutes and seconds values must be '00'.

- MI represents a 2-digit minutes value between 00–59
- SS represents a 2-digit seconds value between 00–59.

FROM query-id1 TO query-id2

Specifies that queries with query ids between query-id1 and query-id2 are listed.

SHOW ALL

Specifies that all queries matching the specified criteria are displayed.

SHOW number-of-queries-to-list

Specifies the maximum number of queries matching the specified criteria to list. This number must be a positive integer.

Usage notes

1. If the LIST QUERIES command is entered without any parameters, a list of all managed queries issued by the current user is returned.
2. A submitter without DBADM authority or without an operator profile with MONITORING privilege is able to list only the queries owned by that submitter.

Chapter 43. LIST QUERY_CLASSES

Lists the settings for the Query Patroller query classes defined for a particular database.

Authorization

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile includes the CONFIGURATION privilege with edit or view authority.

Required connection

None. This command establishes a database connection.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see Chapter 27, “Query Patroller command line support,” on page 143.

```
▶▶—LIST QUERY_CLASSES—▶▶  
└─SHOW—┬─ALL—┬─  
         └─number-of-query-classes-to-list—┘
```

Command parameters

ALL Specifies that all of the query classes currently defined for the database are to be listed.

number-of-query-classes-to-list

Specifies the maximum number of query classes to be listed. This value must be a positive integer.

Usage notes

1. If the LIST QUERY_CLASSES command is entered without any parameters, all of the query classes for the system are displayed.

Chapter 44. LIST SUBMISSION_PREFERENCES

Lists the submission preferences for all Query Patroller submitters or a specified number of submitters for a database.

Authorization

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile includes the USER ADMINISTRATION privilege with edit or view authority.

Required connection

None. This command establishes a database connection.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see Chapter 27, “Query Patroller command line support,” on page 143.

```
▶▶—LIST SUBMISSION_PREFERENCES—————▶▶
▶
└─SHOW—┬─ALL—————┬─
          └─number-of-submission-preferences-to-list—┘
```

Command parameters

ALL Specifies that all of the submission preferences currently defined for the database are listed.

number-of-submission-preferences-to-list
Specifies the maximum number of submission preferences that are listed. This value must be a positive integer.

Usage notes

1. If the LIST SUBMISSION_PREFERENCES command is entered without any parameters, all of the submission preferences for the system are displayed.

Chapter 45. LIST SUBMITTER_PROFILES

Lists the settings for all of the Query Patroller submitter profiles or a specified number of submitter profiles for a database.

Authorization

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile includes the USER ADMINISTRATION privilege with edit or view authority.

Required connection

None. This command establishes a database connection.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see Chapter 27, “Query Patroller command line support,” on page 143.

```
▶—LIST SUBMITTER_PROFILES—————▶
▶
└─SHOW—┬─ALL—————┐
          └─number-of-submitter-profiles-to-list—┘
```

Command parameters

ALL Specifies that all of the submitter profiles currently defined for the database are to be listed.

number-of-submitter-profiles-to-list

Specifies the maximum number of submitter profiles to be listed. This value must be a positive integer.

Usage notes

1. If the LIST SUBMITTER_PROFILES command is entered without any parameters, all of the submitter profiles for the system are displayed.

Chapter 46. qpcenter - Start Query Patroller Center

Starts the Query Patroller Center. qpcenter can be issued from an operating system command prompt.

Authorization

None.

Required connection

None. This command establishes a database connection.

Command syntax

```
►► qpcenter [-u username -p password] -d database ◀◀
```

Command parameters

-u username

Specifies the user name under which to connect to the database.

-p password

The password used to authenticate the user name.

-d database

Identifies the database with which to establish a connection.

If all three of parameters, -u, -p and -d, are specified, Query Patroller automatically connects to the specified database using the given userid and password. If one or more parameters is missing, Query Patroller displays a login window with the supplied information filled in. You must provide the missing information and click OK to connect to the database.

Chapter 47. qpsetup - Set up Query Patroller server

The qpsetup command can do the following:

- Set up a Query Patroller server on a specified DB2 database.
- Replace Query Patroller control tables on a specified control table space.

Authorization


The following section applies to setting up a Query Patroller server on a specified DB2 database:

- You require SYSADM authority to run the qpsetup command.

Required connection

None. This command establishes a database connection.

Command Syntax



Command parameters

USER *user-name*

This parameter specifies the name under which qpsetup is run.

PASSWORD *password*

This parameter specifies the corresponding password of the above mentioned user. This parameter is required if the USER parameter is specified. The above mentioned user will be prompted for a password if the user does not supply one.

DB *db-name*

This parameter specifies the name of the database on which you want to set up DB2 Query Patroller. This parameter is required.

CONTROL_TABLESPACE *tablespace-name*

This parameter specifies the name of the table space on which you want to

create DB2 Query Patroller control tables. This parameter is required. If the table space does not exist, then it will be created. For increased performance and high availability, it is recommended that the control table space be in a single database partition group and at the same database partition where the DB2 Query Patroller server is launched.

CONTROL_DBPARTITIONGROUP *db-partition-group-name*

This parameter specifies the name of the database partition group on which you want to create the control table space. This parameter is optional. If the specified database partition group does not exist, then it will be created. If it is not specified, then the table space will be created on the default database partition group, IBMDEFAULTGROUP. This parameter is ignored if the control table space already exists. For increased performance and high availability, it is recommended that the control tablespace be in a single database partition group and on the same database partition where the DB2 Query Patroller server is launched.

CONTROL_DBPARTITIONNUM *db-partition-group-number*

This parameter specifies the database partition numbers on which you want to create the control database partition group. This parameter is optional. If it is not specified, then the database partition group will be created on all database partitions. This parameter will be ignored if the CONTROL_DBPARTITIONGROUP parameter is not specified, or if the database partition group already exists.

CONTROL_PATH *tablespace-path*

This parameter specifies the path of the table space containers for the control table space. This parameter is required if the control table space needs to be created. It will be ignored otherwise.

CONTROL_DMS YES | NO

This parameter specifies if the control table space to be created is a DMS table space or an SMS table space. This parameter is optional. The default value is NO. This parameter is ignored if the control table space already exists.

CONTROL_DMS_CONTAINER FILE | DEVICE

This parameter specifies if the container type is FILE or DEVICE. This parameter is optional. The default value is FILE. This parameter is ignored if the control table space already exists and if the CONTROL_DMS parameter is not specified or has NO as its value.

CONTROL_DMS_NUMPAGES *number-of-4K-pages*

This parameter specifies the number of pages to be created for the control table space. This parameter is required if the CONTROL_DMS parameter value is YES. This parameter is ignored if the control table space already exists.

RESULT_TABLESPACE *tablespace-name*

This parameter specifies the name of the table space on which you want to create the result tables. This parameter is required. If the table space does not exist, it will be created.

RESULT_DBPARTITIONGROUP *db-partition-group-name*

This parameter specifies the name of the database partition group on which you want to create the result table space. This parameter is optional. If it is not specified, then the table space will be created on the default database partition group, IBMDEFAULTGROUP. This parameter is ignored if the result table space already exists.

RESULT_DBPARTITIONNUM *db-partition-group-number*

This parameter specifies the database partition numbers on which you want to create the result database partition group. This parameter is optional. If it is not specified, then the database partition group will be created on all database partitions. This parameter will be ignored if the RESULT_DBPARTITIONGROUP parameter is not specified, or if the specified database partition group already exists.

RESULT_PATH *tablespace-path*

This parameter specifies the path of the table space containers for the result table space. This parameter is required if the result table space needs to be created. This parameter will be ignored otherwise.

RESULT_DMS YES | NO

This parameter specifies if the result table space to be created is a DMS table space or an SMS table space. This parameter is optional. The default value is NO. This parameter is ignored if the result table space already exists.

RESULT_DMS_CONTAINER FILE | DEVICE

This parameter specifies if the container type is FILE or DEVICE. This parameter is optional. The default value is FILE. This parameter is ignored if the result table space already exists and if the RESULT_DMS parameter is not specified or has NO as its value.

RESULT_DMS_NUMPAGES *number-of-4K-pages*

This parameter specifies the number of 4K pages to be created for the result table space. This parameter is required if the RESULT_DMS parameter value is YES. This parameter is ignored if the result table space already exists.

REPLACE

This parameter specifies whether or not the DB2 Query Patroller control tables are to be replaced with new ones, if they already exist. This parameter is optional.

HELP This parameter specifies that the qpsetup command syntax is to be displayed.

Authorization

The following section applies to replacing Query Patroller control tables on a specified control table space:

- You require DBADM authority to make specifications on existing table spaces using the qpsetup command.

Required connection

None. This command establishes a database connection.

Command syntax

Note: If the table space does not exist, it will NOT be created and an error message will be returned.

```
▶▶ qpsetup CONTROL_TABLES_ONLY [USER user-name PASSWORD password]
```

►DB—*db-name*—CONTROL_TABLESPACE—*tablespace-name*—REPLACE—►

Command parameters

CONTROL_TABLES_ONLY

This parameter specifies that qpsetup will only create the DB2 Query Patroller control tables, without creating table spaces nor binding packages.

USER *user-name*

This parameter specifies the name under which qpsetup is run.

PASSWORD *password*

This parameter specifies the corresponding password of the above user. This parameter is required if USER is specified. The above mentioned user will be prompted for a password if the user does not supply one.

DB *db-name*

This parameter specifies the name of the database on which you want to setup DB2 Query Patroller. This parameter is required.

CONTROL_TABLESPACE *tablespace-name*

This parameter specifies name of the table space on which you want to create DB2 Query Patroller control tables. This parameter is required. If the table space does not exist, then an error message will be returned.

REPLACE

This parameter specifies whether or not the DB2 Query Patroller control tables are to be replaced with new ones, if they already exist. This parameter is optional.

Usage notes

- For each database that you want to use Query Patroller with, you need to run the qpsetup command. For each database, this will create a set of Query Patroller control database objects, such as control tables, views and triggers associated with the tables, as well as user defined functions and procedures required for QP execution. The control tables contain information such as configuration settings, user profiles, and historical query data.

Chapter 48. qptest - Start Query Patroller

Starts Query Patroller. qptest can be issued from an operating system command prompt.

Authorization

You must be the owner of the instance containing the database that you want to run Query Patroller on. You must also have SETSESSIONUSER privilege on PUBLIC.

Required connection

None. This command establishes a database connection.

Command syntax

►►—qptest—*db-name*—————►►

Command parameters

db-name

Specifies the name of the database for which you want Query Patroller to manage queries.

Usage notes

- Start DB2 before starting Query Patroller.
- The configuration parameter *dyn_query_mgmt* must be set to ENABLE for the database for which you want Query Patroller to manage queries.

Chapter 49. qptest - Stop Query Patroller

Stops Query Patroller. `qptest` can be issued from an operating system command prompt.

Authorization

You must have DBADM authority.

Required connection

None. This command establishes a database connection.

Command syntax

```
▶▶ qptest db-name [force] ▶▶
```

Command parameters

db-name

Specifies the name of the database for which you want Query Patroller to stop managing queries.

force Specifies that Query Patroller should stop even when there are active (that is, running or queued) queries. Active queries will be in an inconsistent state until Query Patroller is restarted and query recovery is completed.

Chapter 50. REMOVE OPERATOR_PROFILE

Deletes the specified operator profile from the set of Query Patroller operator profiles.

Authorization

You must meet one of the following requirements:

- Have DBADM authority

Required connection

None. This command establishes a database connection.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see Chapter 27, "Query Patroller command line support," on page 143.

```
▶▶ REMOVE OPERATOR_PROFILE FOR [ USER—'username' ] [ GROUP—'groupname' ] ▶▶
```

Command parameters

USER *username*

Specifies the name of the user whose operator profile is to be deleted. The user ID *username* must also exist as a DB2 authorization ID.

GROUP *groupname*

Specifies the name of the group whose operator profile is to be deleted. This group name must also exist as a DB2 authorization ID.

Examples

The following command removes the operator profile for "sdiniro" in the SAMPLE database:

```
qp -d sample "REMOVE OPERATOR_PROFILE FOR USER 'SDINIRO'"
```

Chapter 51. REMOVE QUERY_CLASS

Removes a query class definition from the Query Patroller QUERY_CLASS control table.

Authorization

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile includes the CONFIGURATION privilege with edit authority.

Required connection

None. This command establishes a database connection.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see Chapter 27, “Query Patroller command line support,” on page 143.

►► REMOVE QUERY_CLASS—*query-class-id*—————►►

Command parameters

query-class-id

The identifier for the query class to be removed.

Examples

The following example removes query class 5 from the list of query classes defined for the TESTDB database:

```
qp -d testdb "REMOVE QUERY_CLASS 5"
```

Usage Notes

1. Query classes can be created, removed, or modified while Query Patroller is started. Creation, changing the maximum query cost, or removal of a query class will take effect immediately unless there are queued or running queries. If there are queued or running queries, including newly submitted queries, the query class changes will take effect when these complete. If you do not want to wait for all queued and running queries to complete, a Query Patroller server restart is required. Updating the maximum number of queries for a query class always takes effect immediately.

Chapter 52. REMOVE QUERY_INFO

Deletes the information for a particular query or set of queries from the `MANAGE_QUERY_INFO` control tables. This means that the information is not available in the Managed query folders in the Query Patroller Center interface or through the `GET QUERY` or `LIST QUERIES` commands.

Authorization

You must meet one of the following requirements:

- Have `DBADM` authority
- Be an operator whose profile includes the `MONITORING` privilege with edit authority.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see Chapter 27, “Query Patroller command line support,” on page 143.

```
▶▶ REMOVE QUERY_INFO [FOR QUERY query-id | FOR ALL | OLDER THAN value [MINUTES | HOURS | DAYS]] ▶▶
```

Command parameters

FOR QUERY *query-id*

Specifies the ID of the query whose information is removed.

FOR ALL

Specifies that all query information in the `MANAGE_QUERY_INFO` control table should be deleted.

OLDER THAN *value* *time-unit*

Indicates that queries older than the specified time are deleted from the `MANAGE_QUERY_INFO` table.

time-units

MINUTES

Queries older than the specified number of minutes are deleted.

HOURS

Queries older than the specified number of hours are deleted.

DAYS

Queries older than the specified number of days are deleted.

Usage notes

1. When a query is deleted, any related result table or result set information is deleted.

2. Queries in Queued or Running state cannot be removed.
3. Held queries cannot be removed. If you want to remove a Held query, cancel it first; then you can remove the cancelled query.

Chapter 53. REMOVE QUERY_INFO_HISTORY

Deletes the information for a particular query or set of queries from the TRACK_QUERY_INFO control tables. This means that, once this command is run, the information is not available in the historical analysis views in the Query Patroller Center interface.

When the information for a query has been deleted from the TRACK_QUERY_INFO control table, it is also deleted from the MANAGE_QUERY_INFO table, and any result set information corresponding to that query is deleted from the RESULT_INFO table. This means that the query is not viewable from the managed queries views in the Query Patroller Center, and result tables for the query are not available.

Authorization

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile includes the HISTORICAL DATA privilege with edit authority.

Required connection

None. This command establishes a database connection.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see Chapter 27, “Query Patroller command line support,” on page 143.

```
▶▶ REMOVE QUERY_INFO_HISTORY — FOR QUERY query-id —————▶▶
                                |
                                | FOR ALL
                                |
                                | OLDER THAN value — DAYS
                                |                   |
                                |                   | YEARS
```

Command parameters

FOR QUERY *query-id*

Specifies the ID of the query whose information is removed.

FOR ALL

Specifies that all query information in the TRACK_QUERY_INFO control table should be deleted.

OLDER THAN *value* **time-unit**

Indicates that historical information for queries older than the specified time is deleted from the TRACK_QUERY_INFO table.

time-units

DAYS Historical information for queries older than the specified number of days is deleted.

YEARS

Historical information for queries older than the specified number of years is deleted.

Usage notes

1. When a query is deleted from the TRACK_QUERY_INFO table, any corresponding managed query in the MANAGE_QUERY_INFO table, any result table or any result information is also deleted.
2. Only completed queries can be removed. In particular, held queries cannot be removed. If you want to remove a held query, you must first cancel it, then you can remove the canceled query.

Chapter 54. REMOVE RESULT

Drops the DB2 table containing the result set for a specified query or set of queries.

Authorization

You must meet one of the following requirements:


- Have DBADM authority
- Be an operator whose profile includes MONITORING privilege with edit authority
- Be the submitter of the query or queries that generated the results.

Required connection

None. This command establishes a database connection.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see Chapter 27, “Query Patroller command line support,” on page 143.



Command parameters

QUERY *query-id*

The ID of the query that generated the results to be deleted.

ALL Indicates that all of the result sets for the specified database are to be deleted.

USER *username*

Specifies that all of the result sets for the queries submitted by USER *username* are to be deleted.

OLDER THAN *value* **time-unit**

Indicates that result tables older than the specified time are deleted.

time-units

MINUTES

Result tables older than the specified number of minutes are deleted.

HOURS

Result tables older than the specified number of hours are deleted.

DAYS

Result tables older than the specified number of days are deleted.

Examples

The following example drops the DB2 table containing the results for query 958 executed on the SAMPLE database:

```
qp -d sample "REMOVE RESULT FOR QUERY 958"
```

Chapter 55. REMOVE RESULT_TABLE_ALIASES Command

Removes all aliases that still exist after their corresponding result tables have been dropped. The aliases were originally created by Query Patroller for result tables.

Authorization

You must have DBADM authority.

Required connection

None. This command establishes a database connection.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see Chapter 27, “Query Patroller command line support,” on page 143.

▶▶—REMOVE RESULT_TABLE_ALIASES—————▶▶

Command parameters

None

Chapter 56. REMOVE SUBMISSION_PREFERENCES

Deletes the submission preferences file for a particular Query Patroller submitter.

Authorization

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile includes the USER ADMINISTRATION privilege with edit authority
- Be the owner of the submission preferences to be removed.

Required connection

None. This command establishes a database connection.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see Chapter 27, "Query Patroller command line support," on page 143.

```
▶▶—REMOVE SUBMISSION_PREFERENCES FOR—USER—'username' —————▶▶
```

Command parameters

USER *username*

Specifies the name of the user whose submission preferences file is to be deleted. The user ID *username* must also exist as a DB2 authorization ID. The user now reverts to using the PUBLIC submission preferences.

Chapter 57. REMOVE SUBMITTER_PROFILE

Deletes a specified submitter profile from the Query Patroller SUBMITTER_PROFILE control tables.

Authorization

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile includes the USER ADMINISTRATION privilege with edit authority.

Required connection

None. This command establishes a database connection.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see Chapter 27, “Query Patroller command line support,” on page 143.

```
▶▶ REMOVE SUBMITTER_PROFILE FOR { USER—'username' | GROUP—'groupname' }
```

Command parameters

USER *username*

Specifies the name of the user whose submitter profile is to be deleted. The user ID *username* must also exist as a DB2 authorization ID.

GROUP *groupname*

Specifies the name of the group whose submitter profile is to be deleted. This group name must also exist as a DB2 authorization ID.

Examples

The following command removes the submitter profile in the SALES database for the group "managers":

```
qp -d sales "REMOVE SUBMITTER_PROFILE FOR GROUP 'MANAGERS'"
```

Chapter 58. RUN HELD_QUERY

Releases a held query. This means that the query runs as soon as resources are available.

Authorization

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile includes the MONITORING privilege with edit authority.

Required connection

None. This command establishes a database connection.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see Chapter 27, “Query Patroller command line support,” on page 143.

►►—RUN HELD_QUERY—*query-id*—————►►

Command parameters

query-id

Specifies the ID of the held query to be run.

Chapter 59. RUN IN BACKGROUND QUERY

Runs a query that was previously submitted. The query may be running or queued. Query Patroller stops the execution of the query and gives control back to the client while the query is resubmitted. The results of the rerun query are stored in a result table.

Authorization

You must be the submitter who submitted the query originally.

Required connection

None. This command establishes a database connection.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see Chapter 27, “Query Patroller command line support,” on page 143.

►►—RUN IN BACKGROUND QUERY—*query-id*—◄◄

Command parameters

query-id

Specifies the ID of the query to be run.

- E** Allows the operator to edit and view Query Patroller query class and system settings.
- V** Allows the operator to view Query Patroller query class and system settings.
- N** Renders the operator unable to view or edit Query Patroller query class and system settings. This is the default value for this parameter.

HISTDATA

Specifies the authority level of the HISTORICAL DATA privilege to be assigned to the operator.

- E** Allows the operator to view and remove historical data.

Note: You must have DBADM authority and SETSESSIONUSER privilege on PUBLIC to generate historical data.

- V** Allows the operator to view historical data.
- N** Renders the operator unable to view or remove historical data. This is the default value for this parameter.

MONITORING

Specifies the authority level of the MONITORING privilege to be assigned to the operator.

- E** Allows the operator to view and manage queries. This includes the ability to monitor, remove, and change the status of managed queries, view query details, and delete result tables.
- V** Allows the operator to view the details of queries managed by Query Patroller.
- N** Renders the operator unable to view or manage queries. This is the default value for this parameter.

USERADMIN

Specifies the authority level of the USER ADMINISTRATION privilege to be assigned to the operator.

- E** Allows the operator to edit and view Query Patroller submitter profiles and submission preferences, and to view operator profiles.
- V** Allows the operator to view Query Patroller operator and submitter profiles and submission preferences.
- N** Renders the operator unable to edit or view Query Patroller operator and submitter profiles and submission preferences. This is the default value for this parameter.

SUSPENDED

Indicates whether or not privileges for this operator profile are to be suspended.

- N** Privileges are not suspended for this operator profile. This is the default value for this parameter.

Y Privileges are suspended for this operator profile.

Examples

The following command updates the operator profile for the HELPDESK group, giving operators with this profile the ability to change the status of held queries run against the PRODUCTION database:

```
qp -d production "UPDATE OPERATOR_PROFILE FOR GROUP 'HELPDESK' USING  
MONITORING 'E'"
```

Chapter 62. UPDATE QUERY_CLASS

Updates the setting details for a specified Query Patroller query class.

Authorization

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile includes CONFIGURATION privilege with edit authority.

Required connection

None. This command establishes a database connection.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see Chapter 27, “Query Patroller command line support,” on page 143.

```
▶▶ UPDATE QUERY_CLASS—query-class-id—USING—————▶▶
|
| [MAX_QUERIES—max-number-of-queries] [MAX_COST—max-query-cost]
|
|
| [DESCRIPTION—'description' ]
|
▶▶
```

Command parameters

query-class-id

Specifies the ID for the query class whose setting details are to be updated.

MAX_QUERIES max-number-of-queries

Specifies the maximum number of queries that can be running simultaneously for this query class. When this threshold of running queries is reached, additional queries for this query class will be queued until resources become available. The value must be greater than or equal to 0 and less than or equal to the value of MAX_TOTAL_QUERIES specified in the Query Patroller system settings.

MAX_COST max-query-cost

The maximum cost for a single query that this query class will accept. The value must be greater than 0 and less than or equal to the value of MAX_TOTAL_COST specified in the Query Patroller System settings. This value must be distinct for every query class defined in the system.

DESCRIPTION description

Optionally specifies a text description for the query class. This parameter is nullable.

Examples

The following command updates query class 8 in the TESTDB database so that it will accept up to 50 queries:

```
qp -d testdb "UPDATE QUERY_CLASS 8 USING MAX_QUERIES 50"
```

Usage notes

1. Query classes can be created, removed, or modified while Query Patroller is started. Creation, changing the maximum query cost, or removal of a query class will take effect immediately unless there are queued or running queries. If there are queued or running queries, including newly submitted queries, the query class changes will take effect when these complete. If you do not want to wait for all queued and running queries to complete, a Query Patroller server restart is required. Updating the maximum number of queries for a query class always takes effect immediately.

Chapter 63. UPDATE SUBMISSION_PREFERENCES

Updates submission preferences for a specified submitter.

Authorization

You must meet one of the following requirements:


- Have DBADM authority
- Be an operator whose profile includes the USER ADMINISTRATION privilege with edit authority
- Be the owner of the profile associated with the submission preferences being updated.

Required connection

None. This command establishes a database connection.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see Chapter 27, “Query Patroller command line support,” on page 143.



Command parameters

USER *username*

Specifies the username of the submitter whose submission preferences are being updated. If no username is specified, the current login ID is used. The user ID *username* must also exist as a DB2 authorization ID.

DEFAULT

Specifies that the existing submission preferences for the specified user be updated with default values for all parameters (for example, `RESULT_EXCEEDED_ACTION` will be 'A', `RESULT_ACCESSIBILITY` will be 'S', and so on).

parameter-keyword value

The following parameters can be set using this command:

SUBMITTER_PROFILE_NAME 'profile-name'

SUBMITTER_PROFILE_TYPE 'profile-type'

Specifies the group profile to use to establish submitter settings, such as thresholds, if the submitter has no profile of type user. If multiple group profiles exist for this submitter and no group profile is specified here, queries will be submitted using the group

profile with the lowest resource thresholds. You must specify 'G' as the profile-type, to indicate a group profile.

RESULT_EXCEEDED_ACTION

Specifies the action to take when the query results to be stored in a result table exceed the maximum result rows specified in the submitter's profile.

'A' Specifies that no results are stored in the result table if the result set is longer than the limit defined in the submitter's profile. This option is the default.

'T' Specifies that a truncated set of results is stored in the result table if the result set is longer than the limit defined in the submitter's profile.

RESULT_ACCESSIBILITY

Specifies whether the result table containing query results will be available to more users than just the submitter.

'O' Specifies that the result table is accessible by the DB2 IDs listed in the value of the OTHER_GRANTEES parameter. The DB2 IDs listed must have access to the database where the query was submitted.

'S' Specifies that access to the result table is restricted to the query submitter. This option is the default.

OTHER_GRANTEES grantees

Specifies the DB2 user IDs or group IDs that can access the result table. Up to 1024 alphanumeric characters are accepted. Multiple IDs must be separated with commas.

RESULT_DESTINATION result-destination-id

Specifies whether the submitting application will wait for query results to return, or will be freed up for further activity.

'A' Specifies that the application that submitted the query will wait for the result set to return while Query Patroller manages the query. When this option is selected, the application that submitted the query may become unresponsive until the result set is returned. This option is the default.

'T' Specifies that the result set is stored in a DB2 table. When the query is submitted, the application that submitted the query becomes free for further processing.

EMAIL_ADDRESSES email-addresses

Specifies the e-mail address or addresses to receive notification regarding queries submitted by this submitter.

Note: This notification only applies if e-mail notification is enabled in the QP_SYSTEM settings.

The value for this parameter may be up to 1024 characters. Multiple e-mail addresses must be separated by a comma.

Examples

A Query Patroller user wants to allow his team members to view the results of his queries against the TEAMDB database. To do this, he updates his submission preferences to make his query results accessible to users "JSMITH" and "AWONG" with the following command:

```
qp -d teamdb "UPDATE SUBMISSION_PREFERENCES FOR USER 'BJONES' USING  
RESULT_ACCESSIBILITY '0' OTHER_GRANTEES 'JSMITH, AWONG'"
```

This command grants users 'JSMITH' and 'AWONG' access to result tables that are created by 'BJONES' after the execution of the command. These users will not have access to result tables created prior to the execution of the UPDATE SUBMISSION_PREFERENCES command.

Chapter 64. UPDATE SUBMITTER_PROFILE

Updates a submitter profile in the SUBMITTER_PROFILE table.

Authorization

You must meet one of the following requirements:

- Have DBADM authority
- Be an operator whose profile includes the USER ADMINISTRATION privilege with edit authority.

Required connection

None. This command establishes a database connection.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see Chapter 27, “Query Patroller command line support,” on page 143.

```
▶▶—UPDATE SUBMITTER_PROFILE—————▶▶
▶—FOR—USER—'username'———USING—DEFAULT————▶▶
      |GROUP—'groupname' |   |profile-parameter—'value' |
```

Command parameters

USER *username*

Specifies the name of the user to be associated with the profile. The user ID *username* must also exist as a DB2 authorization ID.

GROUP *groupname*

Specifies the name of the group to be associated with the profile. The group specified must also exist as a DB2 authorization ID.

DEFAULT

Specifies that the submitter profile should be updated with the default values for all parameters. For any individual parameter that has a default value, the parameter can be set to the default by entering the parameter with DEFAULT as the value. For example, to set the MIN_COST_TO_MANAGE parameter for USERA to the default value, enter the following:

```
qp -d sample "UPDATE SUBMITTER_PROFILE FOR USER 'USERA' USING
MIN_COST_TO_MANAGE DEFAULT"
```

profile-parameter

Specifies the parameter values to be assigned to the profile. The following parameters can be set:

PRIORITY *priority*

Specifies the priority level assigned to queries submitted under this profile.

- Values must be an integer between 0 and 999 inclusive

- The default value is 500

MAX_QUERIES_ALLOWED max-queries

Specifies the maximum number of queries that a submitter is able to run simultaneously. Queries submitted after this limit is reached are placed in a queued state until other submitted queries complete. When creating a submitter profile for a group, note that the value set for this parameter applies to each user. For example, if this value were set to 10 for Group A, then each user belonging to Group A has the authority to run 10 queries simultaneously.

- A value of "-1" indicates that users with this profile can have unlimited queries running simultaneously (up to the value of MAX_TOTAL_QUERIES specified in the QP_SYSTEM table.)
- The default value for this parameter is 100.

MAX_RESULT_ROWS max-number-of-result-rows

Specifies the maximum number of result rows that will be stored in a result table for a single query submitted under this profile. Only queries whose results are to be stored in a result table are subject to this limit.

- A value of "-1" indicates that users with this profile can store results with as many rows as necessary to accommodate the complete result set.
- The default value for this parameter is 1,000,000 rows.

MAX_COST_ALLOWED max-query-cost

Specifies the maximum query cost for a submitter under this profile. If the estimated cost of a query submitted under this profile exceeds this value, the query is placed in a held state.

- A value of "-1" indicates that users with this profile can run queries of any size (up to the value of MAX_TOTAL_COST specified in the QP_SYSTEM table.)
- The default value for this parameter is 10,000,000 timerons.

MIN_COST_TO_MANAGE min-query-cost

Specifies the minimum cost of a query to be managed by Query Patroller. Queries whose estimated cost is lower than this value will not be managed by Query Patroller. A query that falls below this minimum cost will still be tracked for historical analysis, provided the value of the QUERIES_TO_SAVE parameter in the QP_SYSTEM table is set to A (all queries). The default value is 15,000 timerons.

ACCOUNT_ID account-id

Specifies an alphanumeric ID to use for account tracking purposes. Up to 128 characters will be accepted. You can use this parameter to sort submitters into logical groupings to track usage costs.

SUSPENDED Y/N

Specifies whether or not a submitter is prohibited from submitting queries. The default value for this parameter is "N".

INTERCEPT Y/N

Specifies that Query Patroller is to intercept or manage queries submitted by this submitter. If queries are not intercepted, Query Patroller will not evaluate the cost of the query or track it for historical analysis. The default value for this parameter is "Y".

Examples

The following example suspends the privileges for the user "jsmith". After the execution of this command, this submitter will not be permitted to submit queries against the SALES database.

```
qp -d sales "UPDATE SUBMITTER_PROFILE FOR USER 'JSMITH' USING SUSPENDED  
'Y'"
```

Chapter 65. UPDATE QP_SYSTEM

Updates the Query Patroller system settings for a particular database. This command updates entries in the QP_SYSTEM control table.

Authorization

You must meet one of the following requirements:


- Have DBADM authority
- Be an operator whose profile includes the CONFIGURATION privilege with edit authority.

Required connection

None. This command establishes a database connection.

Command syntax

Note: For information about entering Query Patroller commands using the command line interface and general syntax for Query Patroller commands, see Chapter 27, “Query Patroller command line support,” on page 143.



Command parameters

DEFAULT

Resets the entries in the QP_SYSTEM control table to the default values.

Setting-keyword value

Specifies the parameter to be updated and the value that is to be assigned to that parameter. You can update the following system setting categories and their associated parameters using this command:

- system thresholds
 - MAX_TOTAL_QUERIES
 - MAX_TOTAL_COST
- held query handling
 - RUN_HELD_QUERIES
 - RUN_HELD_DURATION
- query intercept
 - INTERCEPT_APPLICATIONS
 - INCLUDE_APPLICATIONS
 - EXCLUDE_APPLICATIONS

MAX_TOTAL_COST—*max-total-query-cost*

Parameters Descriptions

MAX_TOTAL_QUERIES *max-number-of-queries*

The maximum number of queries that can run simultaneously for the entire database. When this threshold of running queries is reached, additional queries are placed in a queued state where they wait until sufficient resources become available.

- This parameter may be set to "-1" to allow an unlimited number of queries to run simultaneously.
- The default setting for this parameter is -1 (unlimited).
- This value must not be less than the MAX_QUERIES value for any query classes defined in the QUERY_CLASS control table.

MAX_TOTAL_COST *max-total-query-cost*

Represents the maximum workload cost threshold, in timerons. If the cost of an incoming query causes the aggregate cost of all running queries to exceed this value, then the incoming query is placed in a queued state where it waits until sufficient resources become available.

- If you type a floating point number, the value must be in decimal or exponential notation.
- This parameter may be set to "-1" to allow an unlimited workload cost.
- The default setting for this parameter is -1 (unlimited) .
- This value must not be less than the MAX_COST value for any query classes defined in the QUERY_CLASS control table.

Held query handling settings

These parameters specify whether held queries are to be run on schedule, and how long held queries should be allowed to run.

Command syntax

Note: These parameters are part of the UPDATE QP_SYSTEM command.

Held query handling settings:

RUN_HELD_QUERIES—*'Y'* | *'N'* | RUN_HELD_DURATION—*run-held-queries-duration—time-unit*

Parameter Descriptions

RUN_HELD_QUERIES

- Y** Indicates that held queries are to be run at a scheduled time. No held queries will be run until a schedule is created.

Note: The Query Patroller command line interface cannot be used to create, update, view, or delete schedules. These tasks must be performed through the Query Patroller Center interface.

N Indicates that held queries are not to be run according to a schedule. This is the default value.

RUN_HELD_DURATION run-held-queries-duration time-unit

Specifies the length of time that held queries are to run. Once this specified period of time has passed, no more held queries will be permitted to run. All remaining held queries will be held until the next scheduled start time for running held queries.

time-unit

MINUTES

Held queries are to run for the specified number of minutes.

HOURS

Held queries are to run for the specified number of hours.


- This number must be an integer greater than or equal to 0.
- The default time is 8 hours.

Query interception settings

These parameters specify the applications whose queries are intercepted by Query Patroller.

Command syntax

Note: These parameters are part of the UPDATE QP_SYSTEM command.



Parameter Descriptions

INTERCEPT_APPLICATION

- A** Specifies that queries from all applications are intercepted by Query Patroller. This option is the default.
- I** Specifies that only queries from the applications specified in the INCLUDE_APPLICATIONS parameter are intercepted by Query Patroller.
- E** Specifies that queries from all applications except those specified in the EXCLUDE_APPLICATIONS parameter are intercepted by Query Patroller.

INCLUDE_APPLICATIONS

Specifies the executable filenames of applications that are intercepted by Query Patroller, separated by commas. Up to 1024 characters are accepted. All other applications will not be intercepted.

Note:

1. This setting only applies if the INTERCEPT_APPLICATION parameter is set to "I".
2. The filename values are case-sensitive.

EXCLUDE_APPLICATIONS

Specifies the executable filenames of applications that are not intercepted by Query Patroller, separated by commas. Up to 1024 characters are accepted. All other applications will be intercepted.

Note:

1. This setting only applies if the INTERCEPT_APPLICATION parameter is set to "E".
2. The filename values are case-sensitive.

You can also use the DB2_QP_BYPASS_APPLICATIONS registry variable to specify applications that you do not want intercepted by Query Patroller. The settings specified for the DB2_QP_BYPASS_APPLICATIONS registry variable take precedence over the application interception settings specified using the UPDATE QP_SYSTEM command. Therefore, if you specify using the DB2_QP_BYPASS_APPLICATIONS registry variable that a particular application is not to be intercepted, then it will not be intercepted even if INTERCEPT_APPLICATION is 'A' or INTERCEPT_APPLICATION is 'I' and INCLUDE_APPLICATIONS lists that application.


System maintenance settings

These parameters specify how often queries and result tables are purged from the system.

Command syntax

Note: These parameters are part of the UPDATE QP_SYSTEM command.

System maintenance settings:



Parameters Descriptions

QUERY_PURGE_PERIOD *time-to-keep-queries* *time-unit*

Indicates how long a query will remain in the MANAGE_QUERY_INFO table after the query completes. When this time expires, and when a scheduled purge job is run, the query is removed from the MANAGE_QUERY_INFO control table. This means that the query is no longer available through the Query Patroller Managed Query view, or through the GET QUERY or LIST QUERIES commands.

time-unit

MINUTES

Queries older than the specified number of minutes are removed.

HOURS

Queries older than the specified number of hours are removed.

DAYS Queries older than the specified number of days are removed.

- The value for this parameter must be an integer.
- A value of 0 means that all queries are removed from the `MANAGE_QUERY_INFO` table whenever the scheduled purge job runs.

Note: The Query Patroller command line interface cannot be used to create, update, view, or delete schedules. These tasks must be performed through the Query Patroller Center interface. If no schedule is created, a default schedule is used.

- If this parameter is not set, or is set to `DEFAULT`, then the queries will be kept for one week from the time they are run.
- A value of -1 means that the queries are never deleted.
- Only completed queries will be removed. In particular, queries that are in a held state are not completed, so they are not removed by this command.

RESULT_PURGE_PERIOD time-to-keep-results time-unit

Indicates how long to keep the result table after a query completes. When the next scheduled purge job runs, result tables that are older than this are dropped.

time-unit

MINUTES

Results older than the specified number of minutes are deleted.

HOURS

Results older than the specified number of hours are deleted.

DAYS Results older than the specified number of days are deleted.

- The value for this parameter must be an integer.
- A value of 0 means all result tables are dropped whenever scheduled purge job runs.

Note: The Query Patroller command line interface cannot be used to create, update, view, or delete schedules. These tasks must be performed through the Query Patroller Center interface. If no schedule is created, a default schedule is used.

- If this parameter is not set, or is set to `DEFAULT`, then the result tables will be kept for 604800 seconds (one week) from the time they are created.
- A value of -1 means that the result tables are never dropped.

RESULT_TABLE_SPACE result-table-space-name

Specifies the name of the table space that will hold result tables.

have generated a result table. If a query completes and the results have not been stored in result table, then no e-mail message is sent.

- N Specifies that no e-mail message is sent to notify a submitter when a query completes. This is the default setting.

EMAIL_SERVER

Specifies the host name or IP address of the appropriate SMTP e-mail server. Up to 256 characters are accepted.

SEND_DESIGNATED

- Y Specifies that e-mail is to be sent to the designated address if the e-mail address field is blank in the submitter's Query Submission Preferences dialog.
- N Specifies that no notification messages are sent if there is no submitter e-mail address specified in the submission preferences for the submitter of the query. This is the default setting.

DESIGNATED_EMAIL_ADDRESS designated-email-address

Specifies the e-mail address to send notification to if the results of a completed query are stored in a results table, and there is no e-mail address specified in the query submitter's submission preferences. The value for this parameter can be up to 256 characters.

Part 9. Reference tables

Chapter 66. Control tables

The Query Patroller control tables are created in the target database during the DB2 Query Patroller setup. The control tables contain information DB2 Query Patroller requires to process queries. It is the information in these tables that can be viewed and updated through the Query Patroller Center and the command line interface.

The table schema for the Query Patroller control tables is DB2QP.

There are three types of control tables for Query Patroller:

-
-
-

Operator profiles (OPERATOR_PROFILE) control table

Contains a row for every operator profile that is defined.

Table 14. OPERATOR_PROFILE control table

Column Name	Data Type	Nullable	Description
NAME	VARCHAR(128)	No	Username of operator profile. Together with TYPE column, forms the primary key for the table.
TYPE	CHAR(1)	No	Type of profile: <ul style="list-style-type: none">• 'U'=User• 'G'=Group Default value is 'U'. Together with NAME column, forms the primary key for the table.
SUSPENDED	CHAR(1)	No	Status of operator privileges <ul style="list-style-type: none">• 'Y'= Privileges suspended• 'N'=Privileges not suspended Default value is 'N'.
CONFIGURATION	CHAR(1)	No	CONFIGURATION privilege authority level <ul style="list-style-type: none">• 'E'=Edit privilege• 'V'=View privilege• 'N'=No privilege Default value is 'N'.
USERADMIN	CHAR(1)	No	USER ADMINISTRATION privilege authority level <ul style="list-style-type: none">• 'E'=Edit privilege• 'V'=View privilege• 'N'=No privilege Default value is 'N'.

Table 14. OPERATOR_PROFILE control table (continued)

Column Name	Data Type	Nullable	Description
MONITORING	CHAR(1)	No	MONITORING privilege authority level <ul style="list-style-type: none"> • 'E'=Edit privilege • 'V'=View privilege • 'N'=No privilege Default value is 'N'.
HISTDATA	CHAR(1)	No	HISTORICAL DATA privilege authority level <ul style="list-style-type: none"> • 'E'=Edit privilege • 'V'=View privilege • 'N'=No privilege Default value is 'N'.
RESERVE	BLOB(64K)	Yes	This column is reserved for future use and should not be modified.

Submitter profiles (SUBMITTER_PROFILE) control table

Contains a row for every submitter profile that is defined.

Table 15. SUBMITTER_PROFILE control table

Column Name	Data Type	Nullable	Description
ID	INTEGER	No	System-generated ID number for submitter profile. Primary key for the table.
NAME	VARCHAR(128)	No	Name of user ID associated with the submitter profile
TYPE	CHAR(1)	No	Type of profile: <ul style="list-style-type: none"> • 'U'=User • 'G'=Group Default value is 'U'.
SUSPENDED	CHAR(1)	No	Status of submitter privileges: <ul style="list-style-type: none"> • 'Y'= Privileges suspended • 'N'=Privileges not suspended Default value is 'N'.
INTERCEPT	CHAR(1)	No	Query interception setting: <ul style="list-style-type: none"> • 'Y'= Intercept queries from submitter • 'N'=Do not intercept queries from submitter. Default value is 'Y'.
PRIORITY	SMALLINT	No	Priority level of queries in the queue. Value must be between 0 and 999 inclusive, with 999 as the highest priority, and 0 as the lowest. Default value is 500.
MAX_QUERIES_ALLOWED	INTEGER	No	Maximum queries allowed to run concurrently from this submitter without queuing up. Value of -1= unlimited. Default value is 100.

Table 15. SUBMITTER_PROFILE control table (continued)

Column Name	Data Type	Nullable	Description
MAX_COST_ALLOWED	DOUBLE	No	Maximum cost of query from this submitter without being held. Value of -1= unlimited. Default value is 10 000 000.
MIN_COST_TO_MANAGE	REAL	No	Minimum cost of a query from this submitter in order to be managed. Must be greater than or equal to 0. Default value is 15 000.
MAX_RESULT_ROWS	BIGINT	No	Maximum result rows stored in result table for a single query from this submitter. (Relevant to scheduled queries only.) Value of -1= unlimited. Default value is 1 000 000.
BLOCK_OPTION	CHAR(1)	Yes	Specify what action to take when a query should be queued, but cannot be queued because the DB2 connection concentrator is activated. <ul style="list-style-type: none"> • 'P' = Allow the query to run immediately • 'R' = Reject the query The default value is null, which means that the setting in BLOCK_OPTION column of QP_SYSTEM control table should be used. (Note: this defaults to 'R'.)
ACCOUNT_ID	VARCHAR(128)	Yes	Account ID for submitter profile. Used for chargeback function.
RESERVE	BLOB(64K)	Yes	This column is reserved for future use and should not be modified.

Query submission preferences (SUBMISSION_PREFERENCES) control table

Contains a row for every submission preferences file that is defined.

Table 16. SUBMISSION_PREFERENCES control table

Column Name	Data Type	Nullable	Description
USER_NAME	VARCHAR(128)	No	User ID of owner of preferences. Primary key for table.
PROFILE_ID	INTEGER	Yes	The submitter profile to be used if user does not have their own profile. Foreign key references SUBMITTER_PROFILE table on delete set null.
RESULT_DESTINATION	CHAR(1)	No	Location to return query results: <ul style="list-style-type: none"> • 'A'=return results to application • 'T'=create result table to store results Default value is 'A'.
RESULT_EXCEEDED_ACTION	CHAR(1)	No	Action to take when query results exceed MAX_RESULT_ROWS for submitter ¹ : <ul style="list-style-type: none"> • 'A'=abort • 'T'=return truncated results Default value is 'A'.

Table 16. SUBMISSION_PREFERENCES control table (continued)

Column Name	Data Type	Nullable	Description
RESULT_ACCESSIBILITY	CHAR(1)	No	Accessibility of query result tables: <ul style="list-style-type: none"> • 'S'= submitter only • 'O'= other users in addition to submitter Default value is 'S'.
OTHER_GRANTEES	VARCHAR(1024)	No ²	List of user IDs that can access results of queries from this submitter. <p>Multiple values separated by a comma.</p>
EMAIL_ADDRESSES	VARCHAR(1024)	Yes	E-mail addresses for notification regarding queries from this submitter. Multiple values separated by a comma. <p>Only applies if e-mail notification is enabled in QP_SYSTEM table.</p>
RESERVE	BLOB(64K)	Yes	This column is reserved for future use and should not be modified.

Note:

1. This column is only applicable to scheduled queries, where results are returned to a result table and not to the submitting application.
2. This column is not null if the value for RESULT_ACCESS_TYPE = 'O'.

Managed query information (MANAGE_QUERY_INFO) control table

Contains a row for every query managed by Query Patroller. Entries in this table are deleted after the period of time specified in the QUERY_PURGE_PERIOD column of the QP_SYSTEM table or when deleted manually.

Table 17. MANAGE_QUERY_INFO control table

Column Name	Data Type	Nullable	Description
ID	INTEGER	No	Query ID. DB2 triggers reference TRACK_QUERY_INFO table on delete cascade. <p>Primary key for table.</p>
STATUS	CHAR(1)	No	Status of query ¹ : <ul style="list-style-type: none"> • 'H'=Held • 'Q'=Queued • 'R'=Running • 'L'=Released from Held state
RUN_IN_BACKGROUND	CHAR(1)	No	Does this query runs in background and saves results in a DB2 result table <ul style="list-style-type: none"> • 'Y'=results stored in result table • 'N'=results returned to application The default value is 'N'.
QUERY_CLASS_ID	SMALLINT	Yes	Query class that query is assigned to run in ² .
USER_MAX_COST_ALLOWED	REAL	Yes	Query submitter's MAX_COST_ALLOWED from SUBMITTER_PROFILE table.

Table 17. *MANAGE_QUERY_INFO* control table (continued)

Column Name	Data Type	Nullable	Description
APPLICATION_HANDLE	BIGINT	Yes	ID of the application that the query is submitted from.
MAX_RESULT_ROWS	BIGINT	Yes	Query submitter's MAX_RESULT_ROWS from SUBMITTER_PROFILE table.
TIME_UPDATED	TIMESTAMP	No	Time that this record was last updated. Default value is the current timestamp.
SESSION_AUTH_ID	VARCHAR(128)	No	DB2 session authorization ID.
SESSION_AUTH_TYPE	CHAR(1)	No	ID type of the SESSION_AUTH_ID: <ul style="list-style-type: none"> • 'U'=USER • 'G'=GROUP • (future: 'R'=ROLE)
MESSAGE_RETURNED	VARCHAR(1024)	Yes	DB2 message returned after execution ³ .
RESERVE	BLOB(64K)	Yes	This column is reserved for future use and should not be modified.

Note:

1. This is the status of the query prior to completion. Once the query has completed, the final status is recorded in the COMPLETION_STATUS field of the TRACK_QUERY_INFO table.
2. If the query ran under the default query class, the value for QUERY_CLASS_ID is 0.
3. The MESSAGE_RETURNED field is usually empty if the query completes successfully.

Query result information (RESULT_INFO) control table

Contains a row for every query with results stored in a results table. Entries in this table are deleted after the period of time specified in the RESULT_PURGE_PERIOD column of the QP_SYSTEM table or when deleted manually.

Table 18. *RESULT_INFO* control table

Column Name	Data Type	Nullable	Description
QUERY_ID	INTEGER	No	Query ID. DB2 triggers reference MANAGE_QUERY_INFO table on delete cascade. Primary key for table.
STATUS	CHAR(1)	No	Status of query results: <ul style="list-style-type: none"> • 'D'=Dropped • 'E'=Exists • 'N'=Does not exist • 'P'=Purged • 'T'=Truncated
OWNER	VARCHAR(128)	Yes	User ID that owns results.
RESULT_TABLE	VARCHAR(128)	Yes	Database table containing result set.
RESULT_SELECT	CLOB(2MB)	Yes	The SELECT statement issued to retrieve results.
RESERVE	BLOB(64K)	Yes	This column is reserved for future use and should not be modified.

Historical analysis (QUERY_ANALYSIS) control table

Contains data generated by historical analysis using DB2 Explain data. This table is populated when historical data generation is performed. Records are deleted from this table when the corresponding entry in the TRACK_QUERY_INFO table is deleted.

Note: A single query might have multiple entries in this table depending on the type of Explain data that exists for the query.

Table 19. QUERY_ANALYSIS control table

Column Name	Data Type	Nullable	Description
QUERY_ID	INTEGER	No	Query ID. Foreign key references TRACK_QUERY_INFO table on delete cascade.
STATEMENT_TYPE	CHAR(2)	No	Descriptive label for type of query. <ul style="list-style-type: none">• 'S'=Select• 'D'=Delete• 'DC'=Delete where current of cursor• 'I'=Insert• 'U'=Update• 'UC'=Update where current of cursor
OBJECT_TYPE	CHAR(1)	Yes	Type of data recorded: <ul style="list-style-type: none">• 'C'=Column• 'I'=Index
OPERATOR_TYPE	CHAR(1)	Yes	Reserved for future enhancement.
OBJECT_SCHEMA	VARCHAR(128)	Yes	The schema for the table index. Applicable when OBJECT_TYPE='I'.
OBJECT_NAME	VARCHAR(128)	Yes	The name of the table column or index.
TABLE_SCHEMA	VARCHAR(128)	No	The schema of the table that the column or index belongs to.
TABLE_NAME	VARCHAR(128)	No	Name of the table that the column or index belongs to.
RESERVE	BLOB(64K)	Yes	This column is reserved for future use and should not be modified.

Historical query information (TRACK_QUERY_INFO) control table

Contains a row for every query managed by Query Patroller. For queries intercepted but not managed by Query Patroller, an entry is stored in this table only if the QUERIES_TO_SAVE field in the QP_SYSTEM table has a value of 'A' (track all queries).

Entries in this table are deleted after the period of time specified in the HISTORY_PURGE_PERIOD column of the QP_SYSTEM table or when deleted manually.

Table 20. TRACK_QUERY_INFO control table

Column Name	Data Type	Nullable	Description
ID	INTEGER	No	ID number for the query. Primary key for the table.
TYPE	SMALLINT	No	Type of statement: <ul style="list-style-type: none"> • 0x0001= Select • 0x0002=Modify
COMPLETION_STATUS	CHAR(1)	No	Status that query completed with: <ul style="list-style-type: none"> • 'A'=Aborted • 'C'=Cancelled • 'D'=Done • 'N'=Not completed • 'U'=Unknown • 'J'=Rejected Default value is 'N'.
MANAGED	CHAR(1)	No	Query Patroller managing query:: <ul style="list-style-type: none"> • 'Y'=yes • 'N'=no Default value is 'Y'.
EXPLAIN_RUN	CHAR(1)	No	Status of EXPLAIN run: <ul style="list-style-type: none"> • 'F'=Failed • 'N'=Not run • 'S'=Ran successfully Default value is 'N'.
QUERY_PRIORITY	SMALLINT	Yes	Priority of submitter that submitted query.
STMT_ATTRIBUTES	INTEGER	No	Bitmap with the following bits: <ul style="list-style-type: none"> • 0x00000001 - static SQL • 0x00000002 - host variable/parameter marker/global variable • 0x00000004 - special register • 0x00000008 - DGTT¹ • 0x00000010 - identity/sequence value • 0x00000020 - result set not allowed • 0x00000040 - session variable • 0x00000080 - positioned statement • 0x00000100 - re-opt statement • 0x00000200 - XQuery statement • 0x00000400 - XML type column in query result • 0x00010000 - queuing not allowed Default value is 0.
NESTING_LEVEL	INTEGER	No	Nesting level of the query. Default value is 0.
ROUTINE_ID	INTEGER	Yes	Routine unique identifier.
PARENT_QUERY_ID	INTEGER	Yes	ID of the immediate parent query.

Table 20. TRACK_QUERY_INFO control table (continued)

Column Name	Data Type	Nullable	Description
PACKAGE_SCHEMA	VARCHAR(128)	Yes	
PACKAGE_NAME	VARCHAR(128)	Yes	
PACKAGE_VERSION	VARCHAR(128)	Yes	
SECTION_ENTRY_NUMBER	INTEGER	Yes	
PROFILE_ID	INTEGER	No	Submitter profile used for query.
RESULT_ROWS	BIGINT	Yes	The number of rows in the returned result set ² .
EXECUTION_TIME_SECONDS	BIGINT	Yes	The seconds portion of the query execution time ³ .
EXECUTION_TIME_MILLI_SECONDS	BIGINT	Yes	The milliseconds portion of the query execution time ³ .
SYSTEM_TIME_SECONDS	BIGINT	Yes	The seconds portion of the total system processor time for the query ³ . System time represents the time spent in system calls. In a system with DPF, this reflects only the time spent in the coordinator node.
SYSTEM_TIME_MILLI_SECONDS	BIGINT	Yes	The milliseconds portion of the total system processor time for the query ³ . System time represents the time spent in system calls. In a system with DPF, this reflects only the time spent in the coordinator node.
USER_TIME_SECONDS	BIGINT	Yes	The seconds portion of the total user processor time for the query ³ . User time represents the time spent executing database manager code. In a system with DPF, this reflects only the time spent in the coordinator node.
USER_TIME_MILLI_SECONDS	BIGINT	Yes	The milliseconds portion of the total user processor time for the query ³ . User time represents the time spent executing database manager code. In a system with DPF, this reflects only the time spent in the coordinator node.
ESTIMATED_COST	DOUBLE	Yes	The estimated cost of the query in timerons.
TIME_CREATED	TIMESTAMP	No	Time when the query was submitted. Default value is the current timestamp.
TIME_STARTED	TIMESTAMP	Yes	Query start time.
TIME_COMPLETED	TIMESTAMP	Yes	Query completion time.
TIME_RELEASED	TIMESTAMP	Yes	Time when query was released from Held state
USER_ID	VARCHAR(128)	No	User ID from DB2.
USER_TYPE	CHAR(1)	No	Type of profile: <ul style="list-style-type: none"> • 'U'=User • 'G'=Group

Table 20. TRACK_QUERY_INFO control table (continued)

Column Name	Data Type	Nullable	Description
STMT_AUTH_ID	VARCHAR(128)	No	Statement authorization ID from DB2.
STMT_AUTH_TYPE	CHAR(1)	No	Type of statement authorization: <ul style="list-style-type: none"> • 'U'=User • 'G'=Group
ACCOUNT_ID	VARCHAR(128)	Yes	Chargeback account ID of submitter
APPLICATION	VARCHAR(128)	Yes	Name of submitting application.
APPLICATION_HOST	VARCHAR(255)	Yes	Host name of the machine that submitted the query.
CLIENT_USER_ID	VARCHAR(255)	Yes	The client user ID set by the application using the sqleseti API.
CLIENT_ACCOUNT_ID	VARCHAR(255)	Yes	The client account ID set by the application using the sqleseti API.
CLIENT_APPLICATION	VARCHAR(255)	Yes	The client application name set by the application using the sqleseti API.
CLIENT_WORKSTATION	VARCHAR(255)	Yes	The client workstation name set by the application using the sqleseti API.
REASON_HELD	VARCHAR(255)	Yes	Message that indicates why the query was held.
REASON_QUEUED	VARCHAR(255)	Yes	Message that indicates why the query was queued.
ENVIRONMENT_VALUES	BLOB(64K)	Yes	Compilation environment values.
STATEMENT	BLOB(2MB)	Yes	Text of query ⁴
RESERVE	BLOB(64K)	Yes	This column is reserved for future use and should not be modified.

Note:

1. DGTT= Declared Global Temporary Table.
2. For queries whose results are returned to the client application, this information is only available when the statement monitor switch is set to 'ON'. For queries whose results are returned to a result table, this information is available even when the statement monitor switch is set to 'OFF'.
3. Available only when the timestamp and statement monitor switches are set to 'ON'.
4. Since the data type for this field is BLOB, you will not be able to perform a simple SELECT statement to retrieve the values in the table. In order to query this field, you need to use the db2qp.convertToString function to convert the values for retrieval. For example, to select the statement field from this table, you will enter the following:

```
select db2qp.convertToString(statement) from db2qp.track_query_info
```

Query Patroller system (QP_SYSTEM) control table

The QP_SYSTEM table contains a single entry that stores all of the settings for the database.

Table 21. QP_SYSTEM control table

Column Name	Data Type	Nullable	Description
QUERIES_TO_SAVE	CHAR(1)	No	Queries to be tracked for Historical Analysis: <ul style="list-style-type: none"> • 'M'=Managed queries only • 'A'=All queries Default value is 'M'.
QUERY_PURGE_PERIOD	INTEGER	No	Length of time to keep queries (in seconds). Value of -1 =unlimited. Default value is 604800 (one week).
RESULT_PURGE_PERIOD	INTEGER	No	Length of time to keep query results (in seconds). Value of -1 =unlimited. Default value is 604800 (one week).
HISTORY_PURGE_PERIOD	INTEGER	No	Length of time to keep historical data for queries (in seconds). Value of -1 =unlimited. Default value is 31536000 (365 days).
MAX_TOTAL_QUERIES	INTEGER	No	Maximum number of managed queries that can run concurrently. Value of -1 =unlimited. Default value is -1 (unlimited).
MAX_TOTAL_COST	DOUBLE	No	Maximum total cost of all concurrent managed queries. Value of -1 =unlimited. Default value is -1.
MAX_QUERY_COST	DOUBLE	No	Reserved for future enhancement. Default value is -1.
RESULT_TABLE_SPACE	VARCHAR(128)	Yes	Table space to store result tables. If no value is specified, the default table space is used.
REJECT_HIGH_COST_QUERY	CHAR(1)	No	Reserved for future enhancement. Default value is 'N'.
RUN_HELD_QUERIES	CHAR(1)	No	Run held queries on schedule: <ul style="list-style-type: none"> • 'N'=Queries not run • 'Y'=Queries run. • (future: 'C'=Cancel) Default value is 'N'.
RUN_HELD_DURATION	INTEGER	No	Period of time over which held queries are run, in seconds. Value of -1 =unlimited. Default value is 28800 (8 hours).
EMAIL_ENABLE	CHAR(1)	No	Allow e-mail messages to be sent to submitters regarding query completion or errors. <ul style="list-style-type: none"> • 'N'=E-mail not enabled • 'Y'=E-mail enabled Default value is 'N'. This applies only to queries that return results to result tables.

Table 21. QP_SYSTEM control table (continued)

Column Name	Data Type	Nullable	Description
EMAIL_SERVER	VARCHAR(256)	Yes	The host name or IP address of the SMTP e-mail server.
SEND_DESIGNATED	CHAR(1)	Yes	E-mail is to be sent to the designated address if no e-mail address is specified in submitter's submission preferences: <ul style="list-style-type: none"> • 'Y'=Yes • 'N'=No Default value is 'N'.
DESIGNATED_EMAIL_ADDRESS	VARCHAR(256)	Yes	E-mail address that will receive notification messages in cases where no e-mail address is specified in submitter's submission preferences, and SEND_DESIGNATED='Y'.
INTERCEPT_STATIC_SQL	CHAR(1)	No	Reserved for future enhancements. Default value is 'Y'.
INTERCEPT_APPLICATION	CHAR(1)	No	Intercept applications: <ul style="list-style-type: none"> • 'A'=All applications • 'I'=Intercept only applications listed in INCLUDE_APPLICATIONS • 'E'=Intercept all but the applications listed in EXCLUDE_APPLICATIONS Default value is 'A'.
TIME_HIST_GENERATOR_LAST_RUN	TIMESTAMP	Yes	The date and time that historical data was last generated..
BLOCK_OPTION	CHAR(1)	No	Specify what action to take when a query should be queued, but cannot be queued because DB2 connection concentrator is activated. <ul style="list-style-type: none"> • 'P'=allow the query to run immediately • 'R'=reject the query The default value is 'R'.
INCLUDE_APPLICATIONS	VARCHAR(1024)	Yes	List of applications to intercept, separated by a comma.
EXCLUDE_APPLICATIONS	VARCHAR(1024)	Yes	List of applications not to intercept, separated by a comma.
RESERVE	BLOB(64K)	Yes	This column is reserved for future use and should not be modified.

Query classes (QUERY_CLASS) control table

The QUERY_CLASSES table contains a row for every query class that is defined for the database.

Table 22. QUERY_CLASS control table

Column Name	Data Type	Nullable	Description
ID	SMALLINT	No	Query class ID. Primary key for table.

Table 22. QUERY_CLASS control table (continued)

Column Name	Data Type	Nullable	Description
MAX_QUERIES	INTEGER	No	Maximum number of queries that can run concurrently in the query class. Value of -1 =unlimited.
MAX_COST	DOUBLE	No	Maximum cost of queries running in this class, in timerons ¹ .
DESCRIPTION	VARCHAR(256)	Yes	Text description of query class.
RESERVE	BLOB(64K)	Yes	This column is reserved for future use and should not be modified.

Note:

1. No two query classes defined in the same system can have identical values for this field.

Schedule information (SCHEDULE) control table

The SCHEDULE table contains a row for every schedule that is defined for a database.

Table 23. SCHEDULE control table

Column Name	Data Type	Nullable	Description
ID	INTEGER	No	System-generated ID for schedule. Primary key for table.
COMMAND	VARCHAR(128)	No	Type of schedule: <ul style="list-style-type: none"> • 'H'=Purge historical data • 'Q'=Purge query or result table • 'R'=Run held query
SUSPENDED	CHAR(1)	No	Schedule is suspended: <ul style="list-style-type: none"> • 'N'=No • 'Y'=Yes Default value is 'N'.
START_DATE	TIMESTAMP	No	Start date for schedule.
END_DATE	TIMESTAMP	Yes	End date for schedule.
INTERVAL_UNIT	INTEGER	Yes	Interval unit for schedule
INTERVAL	INTEGER	Yes	Frequency of schedule
INTERVAL_DETAIL	INTEGER	Yes	
NEXT_START_TIME	TIMESTAMP	Yes	Next date and time that scheduled job is to run.
RESERVE	BLOB(64K)	Yes	This column is reserved for future use and should not be modified.

Query information tables

Contains a row for every query with results stored in a results table. Entries in this table are deleted after the period of time specified in the RESULT_PURGE_PERIOD column of the QP_SYSTEM table or when deleted manually.

Table 24. RESULT_INFO control table

Column Name	Data Type	Nullable	Description
QUERY_ID	INTEGER	No	Query ID. DB2 triggers reference MANAGE_QUERY_INFO table on delete cascade. Primary key for table.
STATUS	CHAR(1)	No	Status of query results: <ul style="list-style-type: none"> • 'D'=Dropped • 'E'=Exists • 'N'=Does not exist • 'P'=Purged • 'T'=Truncated
OWNER	VARCHAR(128)	Yes	User ID that owns results.
RESULT_TABLE	VARCHAR(128)	Yes	Database table containing result set.
RESULT_SELECT	CLOB(2MB)	Yes	The SELECT statement issued to retrieve results.
RESERVE	BLOB(64K)	Yes	This column is reserved for future use and should not be modified.

Part 10. Appendixes

Appendix A. 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 25. 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 25. 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 26. 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 26. 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 27. 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 261.

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 B. 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

- ADD OPERATOR_PROFILE command 145
- ADD QUERY_CLASS command 149
- ADD SUBMISSION_PREFERENCES command 151
- ADD SUBMITTER_PROFILE command 153

B

- books
 - printed
 - ordering 264

C

- CANCEL QUERY command 157
- CLP (command line processor)
 - support
 - description 7
 - syntax 143
- collecting historical analysis data 121
- command line processor (CLP)
 - configuring a client to server connection 41
 - support
 - description 7
 - syntax 143
- command line support
 - description 7
 - syntax 143
- commands
 - ADD OPERATOR_PROFILE 145
 - ADD QUERY_CLASS 149
 - ADD SUBMISSION_PREFERENCES 151
 - ADD SUBMITTER_PROFILE 153
 - CANCEL QUERY 157
 - FILE RESULT 161
 - GENERATE HISTORICAL_DATA 159
 - GET OPERATOR_PROFILE 163
 - GET QP_SYSTEM 165
 - GET QUERY 167
 - GET QUERY_CLASS 169
 - GET SUBMISSION_PREFERENCES 171
 - GET SUBMITTER_PROFILE 173
 - LIST OPERATOR_PROFILES 175
 - LIST QUERIES 177
 - LIST QUERY_CLASSES 181
 - LIST SUBMISSION_PREFERENCES 183
 - LIST SUBMITTER_PROFILES 185
 - qpcenter 187
 - qpsetup 189
 - qpstart 193
 - qpstop 195
 - REMOVE OPERATOR_PROFILE 197
 - REMOVE QUERY_CLASS 199
 - REMOVE QUERY_INFO 201
 - REMOVE QUERY_INFO_HISTORY 203
 - REMOVE RESULT 205
 - REMOVE SUBMISSION_PREFERENCES 209
 - REMOVE SUBMITTER_PROFILE 211
 - RUN HELD_QUERY 213
 - RUN IN BACKGROUND QUERY 215

- commands (*continued*)
 - SHOW RESULT 217
 - UPDATE OPERATOR_PROFILE 219
 - UPDATE QP_SYSTEM
 - description 233
 - e-mail notification settings 240
 - held query handling settings 235
 - historical data collection settings 239
 - query interception settings 236
 - system maintenance settings 237
 - system threshold settings 234
 - UPDATE QUERY_CLASS 223
 - UPDATE SUBMISSION_PREFERENCES 225
 - UPDATE SUBMITTER_PROFILE 229
- communications
 - Query Patroller
 - configuring between clients an servers 41
- components
 - Query Patroller 7
- configuration road map 45
- configuring
 - client to server connection
 - for Query Patroller 41
 - query classes
 - description 101
 - steps 103
 - Query Patroller
 - e-mail notification settings 240
 - held query handling settings 235
 - historical data collection 239
 - overview 47
 - query classes, description 101
 - query interception settings 236
 - road map 45
 - submitter profiles, description 62
 - submitter profiles, steps 68
 - system maintenance settings 237
 - system thresholds 234
 - using historical analysis 119
 - system-level settings 53
- connection concentrators
 - using with Query Patroller 135
- control tables
 - Query Patroller
 - component of product 7
 - description 245
- cost
 - of queries 89

D

- data
 - collecting 121
- database configuration parameters
 - dyn_query_mgmt 49
- database objects
 - usage reports 119
- DB2 governor
 - Query Patroller 135
- DB2 Information Center
 - languages 265

- DB2 Information Center (*continued*)
 - updating 266, 267
 - versions 265
 - viewing in different languages 265
- DB2 Setup wizard
 - installing Query Patroller client tools
 - Linux 34
 - Windows 38
 - installing Query Patroller server
 - UNIX 31
 - Windows 37
- DB2_QP_BYPASS_APPLICATIONS registry variable 81
- DB2_QP_BYPASS_COST registry variable 81
- DB2_QP_BYPASS_USERS registry variable 81
- DB2QP schema 7
- db2qp.result_info table 11
- documentation
 - overview 261
 - PDF 261
 - printed 261
 - terms and conditions of use 270
- dropping
 - result tables using Query Patroller 97
- dyn_query_mgmt configuration parameter
 - Query Patroller
 - intercepting queries 48
 - query processing 49

E

- e-mail notification
 - Query Patroller submitters 56

F

- FILE RESULT command 161
- filtering
 - queries
 - for historical analysis 131
 - tables
 - for historical analysis 129

G

- GENERATE HISTORICAL_DATA command 159
- GET OPERATOR_PROFILE command 163
- GET QP_SYSTEM command 165
- GET QUERY command 167
- GET QUERY_CLASS command 169
- GET SUBMISSION_PREFERENCES command 171
- GET SUBMITTER_PROFILE command 173

H

- held queries
 - costs 89
 - execution delays 11
 - handling
 - policies 89
 - settings 235
 - very large query scenario 4
 - scheduling 87, 88
- help
 - configuring language 265
 - SQL statements 265

- historical analysis
 - data collection
 - enabling 57, 117
 - procedure 121
 - QP_SYSTEM table settings 239
 - improving performance (scenario) 6
 - limitations 11
 - overview 111
 - Query Activity graph 101
 - reports 119
 - restrictions 11
- historical data
 - determining when last generated 125
 - generating using Query Patroller 123
- historical queries
 - purging automatically 94, 96
 - removing manually 98
 - viewing details 127

I

- indexes
 - viewing details
 - using Query Patroller 133
- installing
 - clients on UNIX
 - overview 31
 - clients on Windows
 - overview 37
 - using DB2 Setup wizard 38
 - overview 27
 - Query Patroller client tools on Linux
 - using DB2 Setup wizard 34
 - server
 - overview, UNIX 31
 - overview, Windows 37
 - using DB2 Setup wizard, UNIX 31
 - using DB2 Setup wizard, Windows 37
 - verification 32
- interactive mode
 - not available for version 9 11
- intercepting
 - queries
 - enabling 48

L

- LIST OPERATOR_PROFILES command 175
- LIST QUERIES command 177
- LIST QUERY_CLASSES command 181
- LIST SUBMISSION_PREFERENCES command 183
- LIST SUBMITTER_PROFILES command 185
- log files
 - Query Patroller, description 7

M

- MANAGE_QUERY_INFO control table 245
- managed queries
 - filtering
 - Query Patroller 84
 - purges
 - scheduling 95
 - status 78
 - viewing details
 - Query Patroller 81

- managed queries (*continued*)
 - viewing SQL
 - Query Patroller 83
 - workload considerations 89
- managing 94
- Manually removing 98
- Manually removing managed queries 98
- MAX_COST_ALLOWED parameter 109
 - considerations for setting 62
 - default value 62
 - tuning 139
- MAX_QUERIES parameter
 - considerations for setting 101
 - tuning 139
- MAX_QUERIES_ALLOWED parameter 65, 109
 - considerations for setting 62
 - default value 62
 - tuning 139
- MAX_RESULT_ROWS parameter 65
 - considerations for setting 62
 - default value 62
- MAX_TOTAL_COST
 - considerations for setting 53
- MAX_TOTAL_COST parameter 109
- MAX_TOTAL_QUERIES
 - considerations for setting 53
- MAX_TOTAL_QUERIES parameter 109
 - tuning 139
- MIN_COST_TO_MANAGE parameter 65
 - considerations for setting 62
 - default value 62
 - tuning 139

N

- notices 271

O

- OPERATOR_PROFILE control table 245
- operators
 - profiles 59
 - creating 60
 - Query Patroller 59
 - restoring privileges 61
 - suspending privileges 61
- optimizer
 - description of function 89
 - estimated query cost, passed to Query Patroller 49
- ordering DB2 books 264

P

- performance
 - Query Patroller 139
 - reports 119
- post-upgrade tasks
 - DB2 servers
 - enabling Query Patroller 21
- privileges
 - operators
 - restoring 61
 - suspending 61
 - submitters
 - restoring 72
 - suspending 72

- problem determination
 - information available 269
 - query management 139
 - tutorials 269

- profiles
 - operator 60
 - submitter 69

Q

- QP_SYSTEM control table 245
- qpcenter command 187
- qpdiag.log log file 7
- qpsetup command 189
- qpsetup.log log file 7
- qpstart command 193
- qpstop command 195
- qpuser.log log file 7
- queries
 - changing status 80
 - cost 89
 - estimated cost 89
 - filtering for historical analysis
 - Query Patroller 131
 - held
 - policy for handling 89
 - running at a scheduled time 87
 - interception
 - considerations 89
 - enabling 48
 - for processing 49
 - Query Patroller 77
 - maintenance schedules
 - setting 93
 - management
 - considerations 89
 - emergency query scenario 4
 - in query processing 49
 - Query Patroller 77
 - possible problems with delayed execution 11
 - submission preferences
 - Query Patroller 72
 - submitters
 - interception 62
- query classes
 - configuration considerations 101
 - configuring 103
 - creating
 - Query Patroller 104
 - planning considerations 89
 - Query Patroller 106
 - removing
 - Query Patroller 105
 - usage scenario 5
- Query Controller 7
- query interception
 - command settings 236
 - submitter settings 71
- query management policy 89
- query optimizer
 - description of function 89
 - estimated query cost, passed to Query Patroller 49
- Query Patroller
 - changing query status 80
 - command line support 143
 - components 7
 - control tables 245

- Query Patroller (*continued*)
 - description 1
 - e-mail notification of submitters
 - enabling 56
 - enabling
 - after DB2 upgrade 21
 - enabling to intercept queries 48
 - historical analysis interface 113
 - historical data
 - collection enabling 57, 117
 - determining when last generated 125
 - generating 123
 - historical query details
 - viewing 127
 - index details
 - viewing 133
 - installation environment 27
 - installing
 - overview 25
 - limitations 11
 - list of databases
 - updating 56
 - managed queries
 - filtering 84
 - viewing details 81
 - viewing SQL 83
 - operators 59
 - profiles
 - operator 59
 - queries
 - filtering for historical analysis 131
 - interception and management 77
 - submission preferences 72
 - query classes 106
 - creating 104
 - removing 105
 - query processing 49
 - query submission preferences 72
 - query thresholds
 - setting 55
 - restrictions 11
 - result sets 85
 - result tables
 - dropping manually 97
 - viewing 86
 - scenarios
 - handling large queries 4
 - improving performance using historical analysis 6
 - managing queries of different sizes 5
 - managing query submitter needs 3
 - running large emergency queries 4
 - starting 48
 - stopping 49
 - stored procedures 7
 - submitters 65
 - profiles 65
 - system variables 81
 - tables
 - filtering for historical analysis 129
 - thresholds 109
 - upgrading
 - client tools 23
 - server 21
 - using with DB2 governor 135
 - with connection concentrator 135
- Query Patroller Center
 - component of Query Patroller 7

- Query Patroller Center (*continued*)
 - historical analysis
 - enabling collection of data 57, 117
 - Managed Queries folder 101
 - starting 187
- query submission preferences
 - setting
 - another submitter 74
 - individual query 73
- query submitter profile 83
- query thresholds
 - setting 55
- QUERY_ANALYSIS control table 245
- QUERY_CLASS control table 245
- queue priority
 - configuration considerations 62
- queued queries
 - problems with delayed execution 11
- quiesce mode
 - Query Patroller bypassed by all queries 11

R

- REMOVE OPERATOR_PROFILE command 197
- REMOVE QUERY_CLASS command 199
- REMOVE QUERY_INFO command 201
- REMOVE QUERY_INFO_HISTORY command 203
- REMOVE RESULT command 205
- REMOVE RESULT_TABLE_ALIASES command 207
- REMOVE SUBMISSION_PREFERENCES command 209
- REMOVE SUBMITTER_PROFILE commands 211
- Removing managed queries 98
- reports
 - historical analysis
 - description 119
 - on size distribution of queries 101
 - on submitters 62
- resource limits
 - submitter
 - setting 71
- restrictions
 - Query Patroller 11
- result sets 85
- result tables 85
 - dropping manually
 - Query Patroller 97
 - maintenance schedules
 - setting 93
 - orphaned aliases
 - removing 99
 - purges
 - scheduling 95
 - viewing
 - Query Patroller 86
- RESULT_INFO control table 245
- RUN HELD_QUERY command 213
- RUN IN BACKGROUND QUERY command 215
- RUN_HELD_DURATION parameter
 - tuning 139

S

- scenarios
 - handling large queries 4
 - improving performance using historical analysis 6
 - managing queries of different sizes 5

- scenarios (*continued*)
 - managing query submitter needs 3
 - running large emergency queries 4
- SCHEDULE control table 245
- scheduling
 - purging historical queries 96
 - purging managed queries and result tables 95
- servers
 - manual set up 46
- set up Query Patroller server command 189
- SHOW RESULT command 217
- SQL statements
 - displaying help 265
 - list of statements incompatible with Query Patroller function 11
- start Query Patroller Center command 187
- start Query Patroller command 193
- starting
 - Query Patroller 48
 - syntax 193
 - Query Patroller Center 187
- stop Query Patroller command 195
- stopping
 - Query Patroller 49
 - syntax 195
- stored procedures
 - Query Patroller 7
- SUBMISSION_PREFERENCES control table 245
- submitter profile configuration
 - planning 62
- SUBMITTER_PROFILE control table 245
- submitters
 - distinguishing, in a three-tier setup 62
 - distinguishing, in a two-tier setup 62
 - privileges
 - restoring 72
 - suspending 72
 - profiles 65
 - configuring, description 62
 - configuring, steps 68
 - default settings 62
 - usage scenario 3
 - users and groups 69
 - queries
 - interception of 62
 - Query Patroller 65
 - query submission preferences
 - setting individual preferences 73
 - setting preferences for another submitter 74
 - queue priority considerations 62
 - reports on
 - activity 119
 - for configuring submitter profiles 62
 - resource limits
 - setting 71
- system settings
 - Query Patroller maintenance 237

T

- tables
 - filtering for historical analysis
 - Query Patroller 129
- terms and conditions
 - use of publications 270
- testing 56

- three-tier setup
 - distinguishing submitters 62
- thresholds
 - Query Patroller 109
 - in submitter profiles 62
 - system level, settings 234
- timers 89
- TRACK_QUERY_INFO control table 245
- troubleshooting
 - online information 269
 - tutorials 269
- tuning
 - performance
 - Query Patroller 139
- tutorials
 - problem determination 269
 - troubleshooting 269
 - Visual Explain 269
- two-tier setup
 - distinguishing submitters 62

U

- UPDATE OPERATOR_PROFILE command 219
- UPDATE QP_SYSTEM command
 - description 233
 - settings
 - for e-mail notification of submitters 240
 - for held query handling 235
 - for historical data collection 239
 - for query interception 236
 - for system maintenance 237
 - for system thresholds 234
- UPDATE QUERY_CLASS command 223
- UPDATE SUBMISSION_PREFERENCES command 225
- UPDATE SUBMITTER_PROFILE command 229
- updates
 - DB2 Information Center 266, 267
- upgrading
 - Query Patroller client tools 23
 - Query Patroller server 21

V

- variables
 - Query Patroller 81
- verifying
 - server installation 32
- viewing 83
- Visual Explain
 - tutorial 269

W

- workloads
 - analysis using Query Patroller 121
 - management using Query Patroller 1

X

- XQuery language
 - not supported by Query Patroller 11



Printed in USA

SC27-2467-00



Spine information:

IBM DB2 9.7 for Linux, UNIX, and Windows

Query Patroller Administration and User's Guide

