

REAL-TIME DATABASE OBJECT STATISTICS COLLECTION

5

FIELD OF THE INVENTION

The present invention is directed to computer databases. More particularly, the present invention is directed to the collection of computer database object statistics.

10

BACKGROUND OF THE INVENTION

The performance of databases and database management systems ("DBMS"s) is critical to many corporations to meet the need of their customers. Customers desire timely and accurate information, and this information is frequently retrieved by an application program interfacing with a database. Poor performance of a database is usually due to inefficient structured query language ("SQL") by the application program, or inefficient access path selection by the database and DBMS.

15

Most DBMSs, such as the DB2 relational DBMS from IBM Corp., use database object statistics to determine access paths. The access paths may be determined at plan bind time, or at execution time in the case of dynamic SQL. In either case, if the statistics are not up to date the access path selected may not be optimal.

20

Products are available for collecting statistics of database objects. Examples of these prior art products are "RUNSTATS" from IBM Corp., and the "Database Analyzer" from Platinum Technology Inc. However, one drawback with these products is that they must be run in batch mode, which is time consuming. A further drawback is that if the highest accuracy is desired, the database objects must be in a read-only state during statistics collection.

25

Based on the foregoing, there is a need for a system and method for collecting and maintaining database object statistics on a real-time basis without limiting the database objects to read-only status during collection.

SUMMARY OF THE INVENTION

5 One embodiment of the present invention is a system for collecting database object statistics by a database management system. The system selects one or more objects and corresponding one or more object fields and then creates a base set of statistics for the selected object fields. The system further reads active log records and updates the base set of statistics when a change is read in the active log records. Finally, the system begins
10 extracting log records when an end-of-file of the 10 active log records is reached, and updates the updated statistics based on the extracted log records.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block diagram of a system that can implement one embodiment of the present invention.

15 Fig. 2 is a flowchart of the steps performed by the system in accordance with one embodiment of the present invention.

DETAILED DESCRIPTION

Fig. 1 is a block diagram of a system that can implement one embodiment of the present invention. The database object statistics collection system 10 includes a computer 15
20 coupled to a database 20. Computer 15 can be any general purpose computer that includes at least a processor and a memory device (not shown).

Computer 15 executes a DBMS that interacts with database 20. In the embodiment described herein, the DBMS is DB2 from IBM Corp. However, any other DBMS can be used that is capable of responding to the steps described below in conjunction with Fig. 2.

Database 20 is comprised of stored database objects. Examples of database objects include a table, a tablespace, an index, an index space, etc. Various statistics may be collected and monitored for each database object. For example, it may be desired to collect the number of rows included in each table, or the number of active pages in a tablespace.

- 5 These statistics allow the DBMS to generate optimized access paths to database 20.

Each table in database 20 includes a log record. Each log record is a record of changes that were made to its corresponding table, and can be used to rebuild a table in the event of a problem.

- 10 Computer 15 stores in its memory device a set of instructions that, when executed by the processor, allow the DBMS to perform real-time collection of database object statistics in accordance with one embodiment of the present invention. Fig. 2 is a flowchart of the steps performed by system 10 in accordance with one embodiment of the present invention.

- At step 200, at least one database object is selected to collect statistics on, and the statistic fields from the objects are also selected. Examples of database objects include a table and a tablespace. An example of a field for a table is the "number of rows in the table". An example of a field for a tablespace is the "number of active pages in the tablespace." In one embodiment, the fields are columns within a DB2 catalog and are used to optimize access paths to database 20.

- 20 At step 210, a request is initiated to begin collecting statistics for the objects and the corresponding object fields selected in step 200.

At step 220, a base set of statistics are created by reading the entire contents of database 20 and extracting statistics information for the selected objects and fields. Step 220 can be implemented using a known statistics retrieval product such as the "Database Analyzer" from Platinum technology inc.

After the base set of statistics is established, at step 230 the active and archived log records are read. At step 240, the selected fields are updated as changes are read in the log records.

At step 250, it is determined whether the end-of-file of the active log is reached. If
5 the end-of-file has not been reached, step 230 is continued. If the end-of-file of the active log has been reached, system 20 stops reading active log records.

At step 260, a real time extraction of log records is begun by examining log records before they are written to the active log. In one embodiment, the DB2 "Log Exit" function is executed to implement the extraction. Log Exit is a known DB2 application program
10 interface ("API"), but any method available for extracting log records in real time can be used in the present invention.

At step 280, system 10 checks for a gap between the first live log record extracted and the last log record read at step 230. If a gap exists, the gap is processed by rereading the active log from the log record that corresponded to the end-of-file through the first live log
15 record extracted. This insures that no log records were missed between end-of-file and the first extracted log record.

After the gap has been checked and processed, at step 290 the real time extraction of log records is continued.

At step 295, it is determined whether a request is issued by a user to refresh the
20 statistics. If a request is issued, at step 298 the statistic fields are updated and the up-to-date database statistics are made available to the user. Step 290 then continues until an additional refresh request is issued or a request to stop the process is received. In one embodiment, the statistics are refreshed instantly by accessing the desired field and updating the database catalog.

In another embodiment, instead of waiting for a refresh request at step 295, the fields are continuously updated when changes are encountered as the log records are extracted. In this embodiment, the fields, and therefore the statistics, may be continuously displayed to the user. This allows the user to constantly monitor the database statistics.

5 As disclosed, the system and method in accordance with one embodiment of the present invention allows a user to collect and monitor statistics in real-time or on demand, therefore providing highly accurate and timely statistical information to the user.

Several embodiments of the present invention are specifically illustrated and/or described herein. However, it will be appreciated that modifications and variations of the
10 present invention are covered by the above teachings and within the purview of the appended claims without departing from the spirit and intended scope of the invention.

WHAT IS CLAIMED IS:

1. A method of collecting database object statistics by a database management system (DBMS), said method comprising:

(a) selecting one or more objects and corresponding one or more object fields;
5 (b) creating a base set of statistics for the selected object fields;
(c) reading active log records;
(d) updating the base set of statistics when a first change is read in the active log records;

(e) when an end-of-file of the active log records is reached, begin extracting log
10 records; and

(f) updating the updated statistics based on the extracted log records.

2. The method of claim 1, wherein step (f) is executed in response to a user request.

3. The method of claim 1, wherein step (f) is executed in response to a second
15 change read in the extracted log records.

4. The method of claim 1, further comprising:

checking for a gap when a first log record is extracted.

5. The method of claim 4, wherein checking for a gap comprises:

rereading the active log records from the active log record corresponding to the end-
20 of-file through the first extracted log record.

6. The method of claim 1, wherein the DBMS is DB2.



7. The method of claim 6, wherein step (e) comprises executing a log exit function.

8. A system for collecting database object statistics by a database management system (DBMS), said system comprising:

- 5 means for selecting one or more objects and corresponding one or more object fields;
means for creating a base set of statistics for the selected object fields;
means for reading active log records;
means for updating the base set of statistics when a first change is read in the active log records;
10 means for begin extracting log records when an end-of-file of the active log records is reached,; and
means for updating the updated statistics based on the extracted log records.

9. The system of claim 8, wherein said means for updating the updated statistics is executed in response to a user request.

15 10. The system of claim 9, wherein said means for updating the updated statistics is executed in response to a second change read in the extracted log records.

11. The system of claim 8, further comprising:
means for checking for a gap when a first log record is extracted.

12. The system of claim 11, wherein said means for checking for a gap comprises:
20 means for rereading the active log records from the active log record corresponding to the end-of-file through the first extracted log record.

13. The system of claim 8, wherein the DBMS is DB2.



14. The system of claim 13, wherein said means for begin extracting log records comprises executing a log exit function.

