

JACKIE'S Forum

Run SQL Script - Part III



Jackie Jansen

This is the final installment in our explorations of the capabilities within the “Run SQL Script” component of iSeries Navigator.


Last month we looked at the SQL Assist function within “Run SQL Script”. There is an additional way to invoke this functionality which is similar to the interactive command key 4 support. If a job is currently active on the system you can go to the Database section of iSeries Navigator, right mouse click and choose the option to view the “Current SQL for a Job”. This allows you to select an active job in the system. You can then ask to see the last SQL statement from that job. If you choose “Edit SQL” at this point the “Run SQL Script” window is opened with your SQL statement inserted. You can now use any of the capabilities accessible via “Run SQL Scripts” on this statement including “Visual Explain”

“Visual Explain” is the graphical depiction of your query. This component was discussed in detail in the November 2003 column. Visual Explain shows you what the query optimizer was doing every step of the way and where the expensive parts of the query occurred. From the “Visual Explain” menu bar in “Run SQL Scripts” you can choose to either “Run and Explain” your SQL statement or simply to “Explain” it. “Explain” basically sets the query time limit to zero. This allows you to review the query optimizer’s access plan without

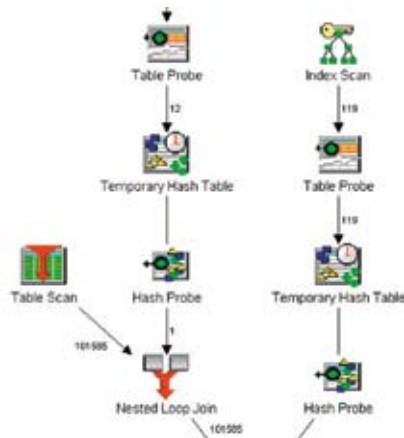
actually running the query. The information about the timing and records retrieved are estimates on which the optimizer is basing its plan. “Run and Explain” will be more accurate, but may also take a lot more time and system resources to execute.

Message ID	Message	Sent
SQL7963	435 rows fetched from cursor CRSR0007.	2/13/06
SQL7968	DESCRIBE of prepared statement STMT0007 completed.	2/13/06
SQL7962	Cursor CRSR0007 opened.	2/13/06
SQL7916	Blocking used for query.	2/13/06
SQL7912	ODP created.	2/13/06
CPI4318	**** Ending debug message for query .	2/13/06
CPI433D	Query options used to build the OS/400 query access plan.	2/13/06
CPI432F	Access path suggestion for file FACT_TABLE.	2/13/06
CPI4329	Arrival sequence access was used for file FACT_TABLE.	2/13/06
CPI4339	Query options retrieved file QAQQINI in library QUSRSYS.	2/13/06
CPI434A	**** Starting optimizer debug message for query .	2/13/06
CPI4339	Query options retrieved file QAQQINI in library QUSRSYS.	2/13/06

“Run SQL Script” has a nice graphical interface to the QAQQINI file. Readers of this column may remember that in March 2004 I wrote about this file and discussed the various query options or settings that you could specify. This file allows a user to specify various query settings in one place. This interactive display makes setting the options much simpler than before. Some of the options in this table refer to system values while yet others come from the “CHGQRYA” command. Examples of functionality controlled via this table would be: if you have an SMP system, i.e. more than one processor, you can specify how aggressively the system will use the available processors.

The Change Query Attributes graphical interface is found under the “Options” menu of “Run SQL Script”. Also under the “Options” tab is the ability to include your SQL debug messages in your job log. Since you can use the “View” tab to display your job log this gives you a very quick and easy way to see what your query is doing. Again, this is much simpler than using a green screen and OS/400 commands to put your job in debug mode, execute your query and then locate the joblog. The debug messages will give you information about the various choices made by the optimizer. Messages include information such as which access paths were considered, which one was chosen and why. Stay tuned for V5R4 and additional database functionality within iSeries Navigator. 

Jackie Jansen is a Senior Consulting IT Specialist. She currently works in the IBM Americas Advanced Technical Support Solutions Centre. Jackie is a frequent speaker at iSeries Technical Conferences and User Group meetings. Contact her at jjansen@ca.ibm.com.



If you are testing your query performance you can go into this table to turn off the caching of your result sets. If you forget this step on V5R3 or later your subsequent runs will always be tremendously faster than your first whether you tune the query or not. As of V5R3 the system will cache query result sets across jobs. This cache is cleared at IPL.