

Re: Automatic MOLAP questions

Source:

<http://www.tech-archive.net/Archive/SQL-Server/microsoft.public.sqlserver.olap/2007-10/msg00121.html>

- *From:* "Jeje" <willgart@xxxxxxxxxxxx>
 - *Date:* Wed, 10 Oct 2007 19:04:41 -0400
-

Hi,

1. there is a lock when the polling query is evaluated. so don't do something like "Select max(updatedate) from..." which can cause a large table scan. the automatic notification don't lock anything, SQL Server and SSAS keep a connection alive and when the rows behind the query change SQL Server sent a notification to SSAS to do something. (the same feature is used in ASP.NET to release the cache of a web application when a table change in the source database)
2. the automatic notification is raised when "anything" change in the source table. a polling query execute the query you defined.
3. when you execute a polling query, the result of the query (generally a datetime) and the previous datetime can be used to do an incremental processing (ie loading only the records changed between the 2 dates). SSAS don't try to and can't detect what is new and what is updated. so an incremental process only add content in a cube and can't update the cube. So it's good ONLY if you add data in a cube. if your source database change based on updates, you must do a full process of the partition (incremental is possible but required a lot of works at the DB level; a full process is better and easier)
4. if during a cube process a key is not found SSAS raise an error by default, but you can change the behavior to ignore the error. SSAS will not try to process the dimension. also when an incremental processing of a dimension don't break a cube, a full process of a dimension required that ALL the cubes which use this dimension must be processed too.
5. the performance impact is linked to your parameters. if you do a full process after each update in the source database without any silent interval, this can result in an infinite loop (if you have 1 update by second..). if you setup the automatic mode to do nothing for 1 minute after the last auto process, then during 1 minute there is no process and you have a delay between a DB update and the cube update.
6. no, generally you need to do at least 1 full process, after this it's not required regarding the setup you do. but an incremental process of everything is near impossible in the real life. so a full process is generally required at a regular basis. if you setup everything to do a full process based on a DB trigger, then it's better to do this yourself by running your own XML/A script where you'll have more options (like 1 or multiple transactions, number of process in parallel etc...)

finally...

SSAS don't switch from incremental processing to full processing himself. but you can create different processing mode by partition.

Re: Automatic MOLAP questions

the proactive caching feature is not a solution to simplify the processing of your cubes. it's designed to handle real realtime scenarios and needs.

I recommend to try to use the proactive caching only for your "daily" partition as an incremental processing. and do a full process of the "weekly" partition + "daily" partition every day.

and also, there is a lot of options, so take the time to read the documentation of each option available around the proactive caching.

hope this will helps you a little :)

"Rich" <Rich@xxxxxxxxxxxxxxxxxxxxxxxxxxxx> wrote in message
news:DB1CA9EE-1C25-48B5-8016-8632C27D11BB@xxxxxxxxxxxxxxxxxxxx

Dear All,

We have been doing some testing with Automatic MOLAP and are looking for some detailed information on whether this is the solution we want or not. Its a bit difficult to find details on the inner workings with this.

Our environment is pretty simple: Our database has a star-schema (simple dimensions & facts) and functions as both our OLTP DB and OLAP data warehouse, and both SQL2005 and SSAS run on the same 64-bit, quad-core server. Our Automatic MOLAP is setup to use SQL Notification for polling / event firing.

Some questions I'm hoping someone can help with:

1. Since our data-warehouse is our OLTP, we're wondering if the SQL Notification queries or OLAP processing will put any locks on any objects when its polling or grabbing incremental data?
2. What does SQL Notification look for to identify changes? A timestamp column or new primary key values? The polling queries in Profiler seem to do a TOP lookup by PK.
3. Further to #2 – does it poll for new records only? Or include changing dimensions / UPDATES? Is it possible to do "changed" data with Auto MOLAP incremental processing? (we have a timestamp on all tables to record INSERTs AND UPDATEs)
4. If we choose specific (not all) dimensions and measures to use Auto MOLAP, what happens if a dependancy (FK) in a non-Auto-MOLAP dimension that isn't processed gets referenced in an Auto MOLAP measure? For example, a fact entry with profile_id 100 is added (this measure has Auto MOLAP), profile_id 100 exists in the data warehouse profile dimension table, but the profile dimension has not been processed in OLAP? Is the cube marked for full processing?
5. What is the performance impact in terms of resource usage (memory, CPU,

Re: Automatic MOLAP questions

IO) of doing Auto MOLAP vs just scheduled processing? Is the incremental processing significant? How many temporary OLAP caches are used at any given time if updated are streaming in continually?

6. If all dimensions and measures are set to Auto MOLAP – does this remove the need to fully process?

Any help or references would be greatly appreciated. Thanks and hope you have a great day!

Cheers,
Rich