

Re: Performance Benchmarks?

Source:

<http://www.tech-archive.net/Archive/SQL-Server/microsoft.public.sqlserver.olap/2006-08/msg00132.html>

- *From:* "Jeje" <willgart@xxxxxxxxxxxx>
 - *Date:* Wed, 26 Jul 2006 23:01:09 -0400
-

Hi,

first... 20 seconds its on warm cache or cold cache?
what is your test query?
have you created aggregations?
do you use MOLAP cube or ROLAP?
do you have calculated measures?

second... 40 000 rows is big and unusable for most of the users.
also most of the applications display the records page by page
returning the data is not so long, the delay come from the rendering
process.
you have to train your users to focus on the question and not focus on the
raw data.

there is no formula to anticipate the response time.
there is too many things to consider like the CPU, memory, aggregations,
security, etc...

good luck to convince your users to change their mind :-)

"Greg Hess" <keadrix@xxxxxxxxxxxx> wrote in message
<news:O4xh9cRsGHA.4380@xxxxxxxxxxxxxxxxxxxxxxxxxxxx>

I am looking for performance benchmarks for SSAS 2005. (I tried Google, but couldn't find much there.) We have developed a new cube (the first use of SSAS at my company) and performance on some queries leaves a lot to be desired. There are conflicting views on this. I am starting to think that the performance issues have more to do with the way the query is written (too broad), while another person thinks that the cube is under-performing.

In the problem query we ask for a few measures from three measure groups: sales dollars, shipped quantity, product cost, returned dollars, returned quantity, returned product cost, budgeted sales, budgeted cases. This information is requested for all customers (about 1200) for the past two years by week (104 weeks). In SQL Management Studio this query takes about 20 seconds. The messages tab in the results says it is returning

Re: Performance Benchmarks?

40,000 rows and 8 columns.

If I change the above query to only return data for June 2006 (4 weeks) I get about a 2 second response, and about 2,000 rows. (I left the exact numbers at work.)

Now I'm no expert on SSAS by any means, but I'm guessing that query one will take longer than query two because of the larger volume of data to be returned to the client. My question is: Is the 20 seconds in line because of the amount of data being returned, or does this smell of an issue in the cube? Can one derive a sort of performance indicator using the rows, columns, and time? Perhaps $(\text{rows} * \text{columns}) / \text{time}$?

Thanks,
Greg

PS. It is good to note that this cube only has about 3 months of data in it (June 2005, June and July 2006), the data warehouse behind it was just developed and hasn't been fully populated. The server it is running on isn't lightning fast, 1.6Ghz and 1.2GB RAM. A faster server is in the works, but from what we have been told this server should be able to handle one developer throwing a few queries it's way.