

Re: High CPU in client (Excel, OWC, Proclarity, etc.) accessing Analysis Services

Source:

<http://www.tech-archive.net/Archive/SQL-Server/microsoft.public.sqlserver.olap/2004-03/0062.html>

From: stevo (*intapplications_at_hotmail.com*)

Date: 03/02/04

Date: Tue, 2 Mar 2004 21:22:33 +1000

We have encountered similar problems not limited to Excel, but every client. Based on the AS architecture (as far as I can see) they all must use PTS to access AS. I personally think it is the use of the FILTER function, whose use is dictated by the lack of calc member support by nonemptycrossjoin (yes, even "in-cube" not session based calc members). If you're ever interested, I can send you some MDX to run against Foodmart:Sales that will cause the client to effectively 'crash' – never returning, or taking hours to return is crashing in my book. I think the problems with filter start to occur when there are many dimensions nested on each other on a single axis (worse if many nesting on both axes).

>From what I've seen of Yukon, unless they change it for beta 2 or RTM versions, the nonemptycrossjoin will still not work for calc members, which means that people attempting to 'suppress zeros' (not just empties) will still hit the same performance ceiling.

An interesting approach to "solving" this can be seen in Cognos PPlay – look for the setting to handle zero's from AS cubes. The help file effectively says you can have it fast or you can have it right – you chose. :-)

Enough ranting.... AS is great, but the performance ceilings you hit with what seem to be very simple queries make you scratch your head a bit....

"Kory Skistad" <kskistad@hotmail.com> wrote in message news:c27ce0e1.0402282033.57082255@posting.google.com...

*> I've been dealing with this issue for two years and haven't found a
> satisfactory answer yet. Based on the following postings, this does
> not seem to be an isolated incident:*

>

>

http://groups.google.com/groups?hl=en&lr=lang_en&ie=UTF-8&oe=UTF-8&threadm=%23KtwKPLbDHA.652%40

>

http://groups.google.com/groups?hl=en&lr=lang_en&ie=UTF-8&oe=UTF-8&threadm=01f201c36dfe%24fd440500%2524a001280a%2540phx.gbl%26num%3D2

microsoft.public.sqlserver.olap: Re: High CPU in client (Excel, OWC, Proclarity, etc.) accessing Analysis Services

>
> http://groups.google.com/groups?q=CPU+excel+OLAP&hl=en&lr=lang_en&ie=UTF-8&oe=UTF-8&selm=b08201c
>
> http://groups.google.com/groups?q=CPU+excel+OLAP&hl=en&lr=lang_en&ie=UTF-8&oe=UTF-8&selm=3eac01c
>
> http://groups.google.com/groups?q=CPU+excel+OLAP&hl=en&lr=lang_en&ie=UTF-8&oe=UTF-8&selm=6633def
>
> http://groups.google.com/groups?q=CPU+excel+OLAP&hl=en&lr=lang_en&ie=UTF-8&oe=UTF-8&selm=OORF
>
> http://groups.google.com/groups?hl=en&lr=lang_en&ie=UTF-8&oe=UTF-8&threadm=975fe525.0304211032.6c0fb
>
> http://groups.google.com/groups?hl=en&lr=lang_en&ie=UTF-8&oe=UTF-8&threadm=%23CXuDQc%24CHA.2132
>
> http://groups.google.com/groups?hl=en&lr=lang_en&ie=UTF-8&oe=UTF-8&threadm=c27ce0e1.0209170645.67d93
>
> http://groups.google.com/groups?q=CPU+excel+OLAP&hl=en&lr=lang_en&ie=UTF-8&oe=UTF-8&scoring=d&sel
>
> http://groups.google.com/groups?q=CPU+OWC+OLAP&hl=en&lr=lang_en&ie=UTF-8&oe=UTF-8&selm=uZy%24
>
> http://groups.google.com/groups?hl=en&lr=lang_en&ie=UTF-8&oe=UTF-8&threadm=ueDLI3SfDHA.1060%40TK
> BOLAP%26hl%3Den%26lr%3Dlang_en%26ie%3DUTF-8%26oe%3DUTF-8%26selm%3DueDLI3SfDH
> A.1060%2540TK2MSFTNGP12.phx.gbl%26num%3D7
>
> *I currently have an open case with Microsoft PSS for a week, and they
> haven't been able to provide me with a solution yet.*
>
> *I have a cube based on a single fact table with 2.8 million rows, 8
> dimensions and 12 measures. Originally I also had 8 calculated
> measures. Originally I also had 8 partitions. I created a view in
> the OWC10 with 3 dimensions on the row axis, one on the column axis.
> The largest size of any of these 4 dimensions < 360 members. On the
> page axis I have the remaining dimensions, and I filter two of the
> smaller dimensions. This returns data within 3 minutes. Then I need
> to filter the last remaining dimension which has a single level of
> 3000 members, and I need to exclude only 2 members. This is when the
> CPU = 100% and never, ever, ever, ever returns.*
>
> *I tried this query in OWC, pasted the MDX in the MDX Sample App,
> pasted in Proclarity 5 client, and in Excel XP all with the same
> results... 100% CPU and no returned data.*
>
> *I followed virtually every suggestion in the Analysis Services
> Performance Guide to optimize the query, the cube, the server, etc.
> Tried Execution location=3, Default Isolation=1; Server is dual
> 1.26Gz, 4GB RAM, RAID 5 100GB; partitioned the cube; set the slicer on
> the partition and adjusted the member counts of the slicer dimension;
> used Usage-based optimization based on the query submitted and
> reaggreated using the *EXACT* levels and dimensions needed to resolve
> the query; reinstalled PTSLite from SP3a on the client and already
> have SP3a on the server; ran the query from different PCs; set up
> performance counters (Agg Cache, Query, Processor, Process-msmdsrv,*

Re: High CPU in client (Excel, OWC, Proclarity, etc.) accessing Analysis Services

- > *Physical&Logical Disk, Memory) and recorded activity for 1 hour...*
- >
- > *While the query was grinding, the server was doing absolutely nothing.*
- > *Also the Direct Hit and Filter Hit/sec counters never > 0. Also the*
- > *Indexes Read/sec = 0 always as well. I even submitted several of the*
- > *sample queries from the Foodmart2000 database and none of these 3*
- > *counters shows activity.*
- >
- > *Here is exactly what happens, doesn't matter which client the query is*
- > *submitted from: Query is executed. Query performance counters on*
- > *server show activity and CPUs on server jump. Client machine shows <*
- > *5%. Then 2 minutes later the server drops to 0, and is completely*
- > *idle. Client jumps to 100% and stays there. One hour later client is*
- > *still 100% and server hasn't shown any activity on the Analysis*
- > *Services or other counters.*
- >
- > *PSS called me back and said, "Too many calculated members (8)" So for*
- > *troubleshooting I dropped them all and left only the base measures.*
- > *Same issue. I had 8 partitions originally, so I deleted 7 and*
- > *reprocessed the single, smallest partition. Same issue. I had 4*
- > *dimensions that weren't needed for the query so I dropped them as*
- > *well. No dice. PSS also said my 3000 member dimension was "very*
- > *large" and should have more than 1 level because hierarchies perform*
- > *better. They confirmed that *not* filtering on the remaining*
- > *dimension by excluding a single member allowed the results to return.*
- > *I knew this. I told them this. What have they been doing for a week?*
- > *I asked them about the performance counters, and they said they would*
- > *have to get back to me because they didn't know what direct*
- > *hits/filtered hits/index reads were actually supposed to do... yikes.*
- > *They also said I should look at re-writing my MDX to get better*
- > *efficiency, but while I may be afforded this luxury my users will not*
- > *because they will use tools like Excel, OWC or Proclarity that write*
- > *the MDX for them.*
- >
- > *I can run this query using the same 2.8 million rows within SQL Server*
- > *and filter and exclude the same members to get the desired resultset*
- > *in less than 2 minutes.*
- >
- > *I hope PSS finds an answer because this is not the first query I've*
- > *ran into this issue with, and judging by the numerous postings listed*
- > *above, others have seen this type of behavior as well.*
- >
- > *I suspect PTS is the culprit, but there doesn't seem to be a lot of*
- > *tweaks that we can do with PTS and no easy way to debug, trace, or*
- > *monitor it.*
- >
- > *In my opinion, the benchmark OLAP needs to measure itself by is the*
- > *relational database (SQL Server in my case). For any given query, you*
- > *should be able to do a GROUP BY and aggregate on all the desired*
- > *measures, and then store that in a table. Now perform a simple select*
- > *statement against that aggregated table. If this takes 13 seconds to*

microsoft.public.sqlserver.olap: Re: High CPU in client (Excel, OWC, Proclarity, etc.) accessing Analysis Services

- > *return, it should take roughly 13 seconds to return from an OLAP cube*
- > *that is properly aggregated at the same levels. All of this should be*
- > *done on the server and the results returned to the client. Then PTS*
- > *can cache the necessary results and subsequent queries can be resolved*
- > *from the cached data, or requested additional data from the server*
- > *cache if necessary, and from the server disk if not yet in server*
- > *cache. All this seems straight forward and simple to comprehend, but*
- > *just doesn't seem to be what is happening in actual practice and no*
- > *one I've discussed this with knows exactly what is happening.*
- >
- > *Please, someone help.*