

Query Performance in “BW powered by HANA” (especially in comparison to “BW with BW Accelerator”)

Author

Klaus Nagel: Development Manager, TIP In-Memory Platform BW (SAP AG).

Klaus Nagel joined SAP more than 10 years ago and started as a Developer in the BW team. He now holds the position of Development Manager and leads the BW&HANA Data Management team.

Introduction

With the introduction of the BW Accelerator (BWA) we were able to speed up BW Query execution in most cases dramatically. So if you migrate your system from BW on a traditional RDBMS and an attached BW Accelerator to “BW powered by HANA” you might ask what that means for your Query performance. This paper intends to give some answers and guidelines to this question.

Remark: BW customers not using BW Accelerator today, will most likely experience a huge Query performance boost, based on the “BWA-like” improvements plus the additional HANA capabilities!

Of course, moving to “BW powered by HANA” (BW-HANA) is not just about Query performance, but also enables a ton of other exciting stuff. So just looking at only the Query performance is too one-dimensional, because the real business value comes with things like

- great performance improvements in staging, reducing data latency and enables completely new scenarios,
- new scenarios when combining BW and non-BW data for reporting or staging (see link 4) for more),
- TCO reduction due to: removal of redundant layers (like aggregates/BWA Index and many case InfoCubes), easier and faster modeling, less processes to monitor and schedule (no Rollups, no Changeruns, ...), and of course no more additional BWA required, since it is build-in into HANA,
- new flexibility with BW Workspaces and CompositeProviders,
- and, of course, also the new in-memory planning capabilities,
- ... see link 2) below for the complete list and customer stories.

So, in general, I strongly recommend the move from a BW with a classic RDBMS with or without BWA to BW-HANA. It is not only for the above current benefits, but also it is a move to our strategic and future-proof platform!

But back to my topic: The general statement for Query performance with BW-HANA is: “it is **at least as fast** as with BW Accelerator”. The reason for this is quite simple: the heart of the Query execution engine of HANA is based on BW Accelerator. As with BWA the power of HANA comes from the unique combination of in-memory, columnar data processing technique using highly parallelized query execution optimized on our certified hardware configurations and optimized towards SAP applications (like e.g. BW).

Remark: As you see, it is also hardware dependent – i.e. when comparing BWA with HANA you must ensure that CPU rates and number of cores are comparable.

So, if I say “at least as fast” I mean that there are 3 categories of possible improvements that you will see:

- A) Improvements coming with BW7.30 and its tighter integration of BW and both BWA and HANA by pushing additional calculations to BWA/HANA resp. – see a complete list in our SCN paper. (see link 1) below)
- B) New possibilities due to
 - a. reporting directly on DSOs (reducing the redundant storage in InfoCubes and BWA Index/Aggregates). (see e.g. link 3) below),
 - b. new physical InfoCube layout (HANA-optimized InfoCube without dimension tables) – this is especially helpful for InfoCubes with a “bad” design, i.e. very large dimension tables,
- C) Improvements of the HANA engine capabilities and architecture compared to BWA.

Categorization of Query Runtime

Let’s briefly look at some basics of Query runtime analysis without going into too much detail. You can easily find the Query runtime numbers in the BW Query runtime statistics, e.g. by executing the Query in transaction RSRT with the “Debug” mode “Show Runtime Statistics” or using the BW technical content.

The Query runtime in BW can be split into 3 main parts

- 1) **Client (+Network)** – HANA can do a lot of things, but it doesn’t speed up your frontend rendering times or the time it takes to send the data from the backend to the client. If someone told you so, he probably also told you that you can use HANA to mow the lawn☺. If your Query result are 100s of thousands of cells, the data transfer and rendering will most likely consume most of the execution time – and it will still do so with BW-HANA.

BW-HANA offers a highly sophisticated and powerful OLAP engine – producing millions of cells (for wallpaper print-outs or downloads) is NOT what OLAP stands for.

Whether a Query fits into this category can e.g. be seen in the Query runtime statistics, Event-ID 3200 (“OLAP Data Transfer”), Counter = “# of cells sent to the client”.

- 2) **Analytic Manager** – we have renamed our OLAP engine to the more appropriate “Analytic Manager”. In “BW-HANA” and “BW with BWA” the classic OLAP operations are not only handled by the ABAP application layer, but increasingly more executed by the Calculation Engine of HANA (and partially also by BWA). So actually BW’s OLAP Engine is then a combination of HANA’s Calculation Engine and BW’s Analytic Manager. BW still handles those OLAP operations that are not yet available in HANA and delegates (“manages”) the rest to HANA as it can be executed there much faster (directly at the data) compared to doing it in the application server.

You find the time for the Analytic Manager mainly under the Event-IDs 3000 – 3999 (see exception of 3200 above) of the Query runtime statistics.

- 3) **DB time** – the “DB time” for HANA not only includes the executing of a standard SQL statement on the BW InfoProvider schema, but also those OLAP operations that can be executed by HANA’s Calculation Engine.

You find the DB time under the Query runtime statistics category “Aggregation Layer” and here especially the “DB read time” (also focus on the records read on the DB (“Records, selected”) and transferred from the DB to the application server (“Records, transferred”)).

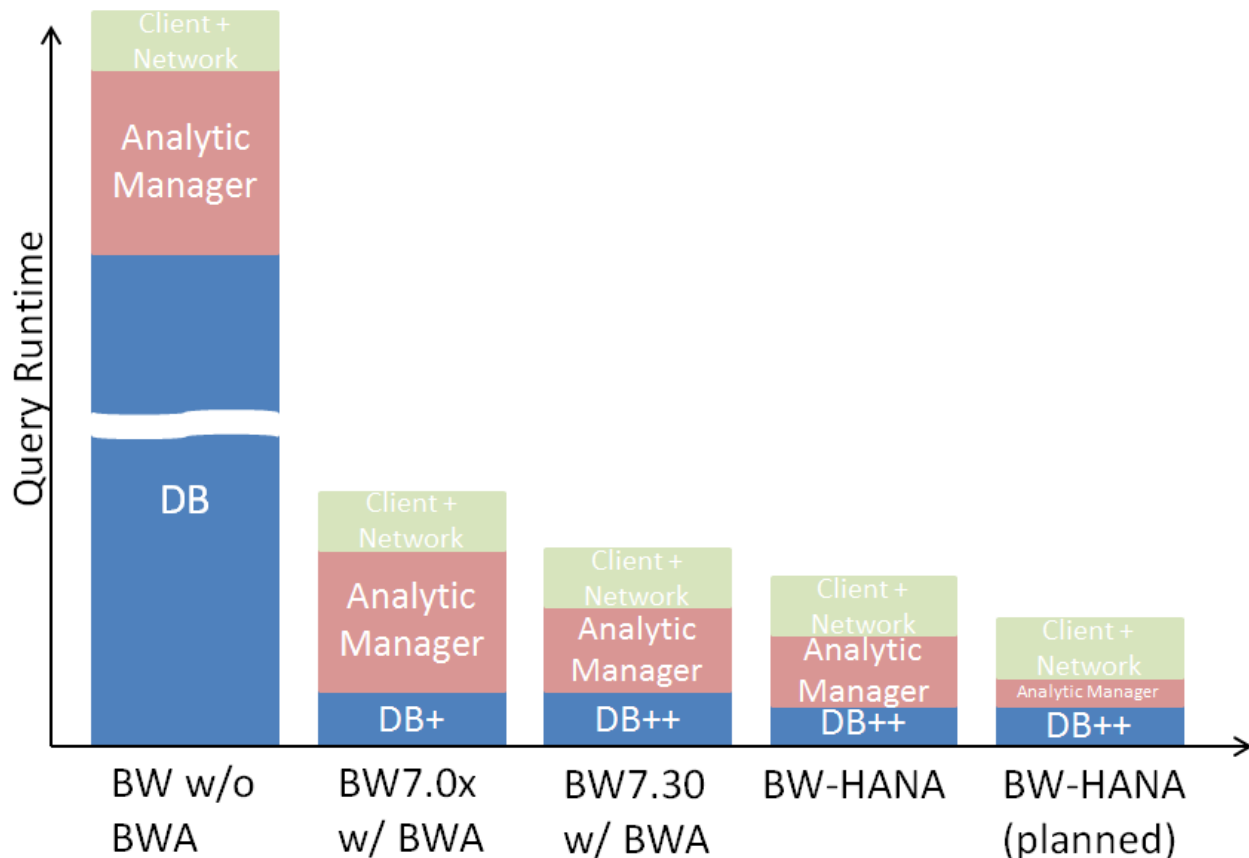


Figure 1: Typical BW Query Runtime distribution and evolution over BW releases.

Query Runtime Improvements

So as mentioned, do not expect any improvements in 1). But what changes with 2) and 3) when switching from BWA to HANA?

Ad part 2) (Analytic Manager) With BW7.00 we pushed aggregation and filtering to BW Accelerator + some hierarchy calculations and special cell-based filter conditions (restricted keyfigures).

With BW7.30 we enabled TopN conditions in Pre-Queries, several cases of exception aggregation, currency conversions and the MultiProvider-UNION in BWA (complete list see link 1) below).

These features are, of course, enabled for BW-HANA as well. But while we stop here for BWA, for HANA this is just the beginning. You will see that we continuously push down additional OLAP calculations to HANA as its Calculation Engine is enhanced further. Currently planned are e.g. Unit conversions and more exception aggregation types and in the long run also formula exception aggregation (there is no concrete timeline available or a sequence of which feature comes first, we will publish those as soon as they are available).

Nevertheless, as of today, there is NO difference between what we push down to HANA and what we already push down to BWA – so you cannot expect performance improvements in this area NOW!

Ad part 3) (DB time) Most Queries perform extremely well on BWA. Since the processing basics have not changed compared to HANA, you will see the same excellent response time with HANA as well (if comparable HW is used).

Remark: Do not forget, that you achieve this performance now directly on the InfoCube data, i.e. no more Rollup into a BWA Index is required. Additionally you can achieve this great performance also on DSOs, as they are also loaded to memory and stored as columnar tables (see more information in e.g. link 3)).

On the other hand, there are some scenarios where BWA had its challenges. For example if the scenario contains very large Masterdata tables (>100 million rows) that had to be split over several blades. If such a large table was joined to a large (and split) fact table and both tables did not have a common partition attribute, a lot of data had to be sent over the inter-blade network.

With HANA and its 512GB (or 1TB) servers (or multiple such nodes in a Scale-Out landscape) there is no need to split even the largest Masterdata tables, making the communication obsolete or much easier to optimize.

Another area where the HANA engine has seen (and will see) improvements is in the way a Query execution plan is generated. Basically HANA has a rule-based optimizer – as BWA. This is a significant difference to classic RDBMS where the cost-based optimizers may work well, but a slight change of your query can result in a different access pattern and a completely different (and potentially bad) runtime. In contrast HANA always offers predictable and, thanks to its advanced techniques, fast runtimes. The rules for HANA's optimizer are continuously improved

and continuously take more application knowhow into account (e.g. of the BW table schemas and Query access patterns). With these improvements being implemented you will see that certain scenarios will be even faster with HANA compared to BWA. An example is the inversion of the JOIN direction for certain accesses and table sizes (this will be shipped shortly).

As HANA's Calculation Engine matures it will also gain an additional performance boost. One example is the BW exception aggregation feature which is pushed down to BWA and HANA (in many cases). As of now the result of the query after the first aggregation type has to be materialized in HANA/BWA as a temporary table and the subsequent accesses (second aggregation type, result rows, ...) then read from this temporary table. This materialization step can be time consuming and for very large intermediate result sets it may even come to memory problems on the smaller BWA blades. For HANA we plan to have this intermediate result only "virtual", i.e. without materialization, which limits memory consumption dramatically and in most cases will be also significantly faster.

So in general, we continuously invest and improve the calculation and aggregation engine of HANA and have certainly not reached the end for a number of scenarios.

So, where will BW-HANA not improve Query performance? (in addition to the category 1) type of Queries) Basically all Queries with OLAP calculations that can not be pushed down to HANA and which have to be performed on a large number of records by the application server. These can be Queries with Virtual Keyfigures/Characteristics (i.e. user exits written in ABAP), or complex formula exception aggregation that, as of now, can not be handled by the HANA calculation engine.

As long as these operations have to be performed on a small dataset, it makes no significant difference whether the application server or HANA does the calculation. But if the dataset is large the data transfer from HANA to the application server is costly and the application server can not perform the calculations with the same speed as HANA due to its different architecture.

As mentioned above, we are continuously decreasing the number of such queries as we enable the push-down of more and more OLAP operations (especially the costly ones). But there are also those operations that, by nature, can not be pushed down, e.g. those with a user exit involved.

Conclusion

Unfortunately the conclusion whether moving from "BW with BWA" to "BW-HANA" itself improves Query performance is not a simple one, but an "it depends". Queries showing good performance with BWA will certainly still perform at least as good with HANA and for the others it needs to be checked case by case in which category they fall using the guidelines above. This might not be fully satisfying, but you, our customers, solve a wide variety of complex business processes with the BW infrastructure and that inhibits simple answers – even with an egg-laying, wooly milk pig like HANA☺.

Nevertheless, as stated already above: Moving to BW-HANA is not only about some Query performance improvements, but about a whole new way of what you can do with your BW system. Therefore: GO FOR IT☺!

Useful Links

- 1) What's New with BW7.30 and BWA7.20 - <http://scn.sap.com/docs/DOC-12041>
- 2) BW-HANA overview - <http://scn.sap.com/docs/DOC-24092>
- 3) Aspect of HANA-optimized InfoCube - <https://www.experiencesaphana.com/docs/DOC-1363>
- 4) Mixed Scenarios in BW-HANA - <https://www.experiencesaphana.com/docs/DOC-1463>

© 2014 SAP SE. All rights reserved.

SAP, R/3, SAP NetWeaver, Duet, PartnerEdge, ByDesign, SAP BusinessObjects Explorer, StreamWork, SAP HANA, and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP SE in Germany and other countries.

Business Objects and the Business Objects logo, BusinessObjects, Crystal Reports, Crystal Decisions, Web Intelligence, Xcelsius, and other Business Objects products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of Business Objects Software Ltd. Business Objects is an SAP company.

Sybase and Adaptive Server, iAnywhere, Sybase 365, SQL Anywhere, and other Sybase products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of Sybase Inc. Sybase is an SAP company.

Crossgate, m@gic EDDY, B2B 360°, and B2B 360° Services are registered trademarks of Crossgate AG in Germany and other countries. Crossgate is an SAP company.

All other product and service names mentioned are the trademarks of their respective companies. Data contained in this document serves informational purposes only. National product specifications may vary.

These materials are subject to change without notice. These materials are provided by SAP SE and its affiliated companies ("SAP Group") for informational purposes only, without representation or warranty of any kind, and SAP Group shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP Group products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.

