SAP HANA® Performance

Efficient Speed and Scale-Out for Real-Time Business Intelligence



Table of Contents

- 3 Introduction
- 4 The Test Environment

Database Schema

Test Data

System Configuration

Setup

Queries

8 Test Results

Baseline Test

Throughput

Ad Hoc Historical Queries

Results Summary

Real-World Experiences

- 11 Conclusion
- 12 Appendix

Introduction

SAP HANA® appliance software enables organizations to optimize their business operations by analyzing large amounts of data in real time. It runs on inexpensive, commodity hardware and requires no proprietary add-on components. It achieves very high performance without requiring any tuning.

A 100 TB performance test was developed to demonstrate that SAP HANA is extremely efficient and scalable and can very simply deliver breakthrough analytic performance for real-time business intelligence (BI) on a very large database that is representative of the data that businesses use to analyze their operations. A 100 TB¹ data set was generated in the same format as would be extracted from the SAP® ERP application (for example, data records with multiple fields) for analysis in the SAP NetWeaver® Business Warehouse (SAP NetWeaver BW) component.²

This paper will describe the test environment and present and analyze the test results.

SAP HANA appliance software enables organizations to optimize their business operations by analyzing large amounts of data in real time.



The Test Environment

The performance test environment was developed to represent a sales and distribution (SD) BI query environment that supports a broad range of Structured Query Language (SQL) queries and business users.

DATABASE SCHEMA

The star-schema design in Figure 1 shows the SD test data environment and each table's cardinality. No manual tuning structures were used in this design; there were no indexes, materialized views, summary tables, or other redundant structures added for the purposes of achieving faster query performance.

TEST DATA

The test data consists of one large fact table and several smaller dimension tables, as seen in Figure 1. There are 100 billion records in the fact table alone, representing 5 years' worth of SD data. The data is hash partitioned equally across the 16 nodes using "Customer_ID." Within a node, the data is further partitioned into one-month intervals. This results in 60 partitions per node and approximately 104 million records per partition. (See also Figure 2 and the "Loading" section.)

SYSTEM CONFIGURATION

The test system configuration (Figure 2) is a 16-node cluster of IBM X5 servers with 8 TB of total RAM. Each server has:

- 4 CPUs with 10 cores and 2 hyperthreads per core, totaling:
 - 40 cores
 - 80 hyperthreads
- 512 GB of RAM
- 3.3 TB of disk storage

SETUP

These performance tests did not use precaching of results or manual tuning structures of any kind and therefore validated the SAP HANA load-then-query ability. A design that is completely free of tuning structures (internal or external) is important for building a sustainable, real-time BI environment; it not only speeds implementation but also provides ongoing flexibility for ad hoc

querying while eliminating the maintenance cost that tuning structures require.

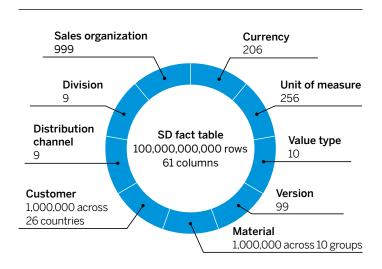
Loading

The loading was done in parallel using the "IMPORT" command in SAP HANA, which is a single SQL statement that names the file to load. Loading is automatically parallelized across all of the nodes and uses the distribution-and-partitioning scheme defined for each table (in this case, hash distribution and monthly range partitions). The load rate was measured at 16 million records per minute, or 1 million records per minute per node. This load performance is sufficient to load 100 million records (representing one business day's activity) in just six minutes.

Compression

Data compression occurs during the data loading process. SAP HANA demonstrated a compression rate 4 of greater than 20 times; the 100 TB SD data set was reduced to a trim 3.78 TB SAP HANA database, consuming only 236 GB of RAM on each node in the cluster (see Figure 2). The large fact table accounts for 85 of the 100 TB in the data set and was compressed to occupy less than half of the available RAM per node.

Figure 1: SD Schema



QUERIES

The query suite totals 20 distinct SQL queries, including 11 base queries plus a number of variations for time interval (month, quarter, and so on). The queries were chosen to represent a mixed workload environment running on data that is representative of its native form (that is, not indexed, tuned, or otherwise denormalized to avoid joins, as would be customary in a conventional database system). They cover a range of BI activities, from departmental to enterprise level, including:

- · General reporting
- Iterative querying (drill downs)
- Ranking
- · Year-over-year analysis

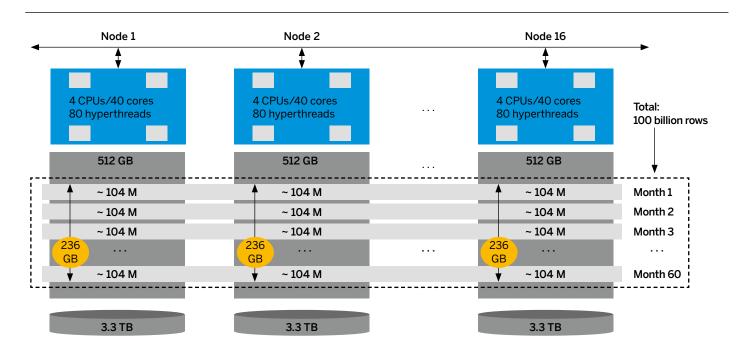
The resulting queries range from moderately complex to very complex, including SQL constructs such as:

- · Multiple joins
- In list
- · Fact-to-fact joins
- Subqueries
- Correlated subqueries (CSQs)
- Union all

The queries are grouped into three general categories of BI usage:

- Reporting Calculations on business performance measures for ranges of materials or customers or both for a time period
- Drill down Iterative user-initiated queries to gather details for a given individual or group of materials or customers or both
- Analytical Periodic deep historical analysis across materials or customers or both

Figure 2: Test Configuration





The table "Query Descriptions" documents each query in terms of its business description, SQL constructs, ranges and qualifiers, and time period variations. If a query is described with multiple time periods, each time period is run as a distinct query; for instance, query R1 is run as four distinct queries, covering

monthly, quarterly, semiannual, and annual time periods, to show its performance relative to changes in data volume (annual data is 12 times the volume of monthly data). Items in bold denote query complexity factors.

Query Descriptions

BI Query	Business Description	SQL Constructs	Ranges and Qualifiers	Time Period(s)
R1	Sales and distribution report by month for: Given range of materials Given range of customers	Star-schema query with multiple joins to dimensions, with group by and order by	Range of materialsRange of customers	Four variations: • 1 month • 3 months • 6 months • 12 months
R2	Sales and distribution report by month for given list of customers	Star-schema query with multiple joins to dimensions, with group by and order by	Range of customers defined using "in list"	1 month
R3	Sales and distribution report by month for: Given range of materials Given customer country	Star-schema query with multiple joins to dimensions, with group by and order by	Range of materialsCustomer country	1 month
R4	Top 100 customers report for: One or several material groups One or several customer countries	Star-schema query with multiple joins to dimensions, with group by and order by	 One or several material groups One or several customer countries 	One month of data for two variations: • 1 country and 1 product group • 3 countries and 2 product groups
R5	Top 100 customers report for: Given material group Given customer country	Star-schema query with multiple joins to dimensions, with group by and order by	Material groupCustomer country	Two variations: • 3 months • 6 months

The majority of queries delivered subsecond response times; and even the most complex query, which involved the entire five-year span of data, ran in under four seconds.



BI Query	Business Description	SQL Constructs	Ranges and Qualifiers	Time Period(s)
D1	Sales and distribution report for: Single customerGiven material group	Star-schema query with multiple joins to dimensions, with group by and order by	Single customer Material group	1 month
D2	Sales and distribution report for: Single customerSingle material	Star-schema query with multiple joins to dimensions, with group by and order by	Single materialSingle customer	1 month
D3	Sales and distribution report for: Single materialGiven customer country	Star-schema query with multiple joins to dimensions, with group by and order by	Single materialCustomer country	1 month
D4	Sales and distribution report for: Single customerGiven range of materials	Star-schema query with multiple joins to dimensions, with group by and order by	Range of materialsSingle customer	1 month
A1	 Year-over-year (YOY) top 100 customers analytical report: Top 100 customers identified for a current time period Business measures calculated for previously identified top 100 customers in a given historical period A final result set calculated that combines current and historical business measures for top 100 customers 	Star-schema query with fact-to-fact join as a correlated subquery that includes multiple joins to dimensions and then group by and order by	 Range of materials Customer country 	Three variations: • 1 month with YOY comparison • 3 months with YOY comparison • 6 months with YOY comparison
A2	 YOY trending report for top 100 customers: Top 100 customers calculated for a given time period of current year Top 100 customers calculated for identical time period of previous year (or multiple years) A final result set calculated that combines business measures for top 100 customers over several years 	Several star-schema subqueries, each with multiple joins, with group by and order by, which are combined by a union all operator	Material group Customer country	Three variations: • 3 months over 2 years • 6 months over 2 years • 3 months over 5 years

Test Results

The test measured both query response time and query throughput per hour. The queries were run first in a single stream, to measure baseline query times, and then in multiple streams of 10, 20, and 25, to measure query throughput (calculated as queries per hour) and query response time in the context of different workloads. The multistream tests randomized the sequence of query submissions across streams and inserted 10 milliseconds of "think" time between each query to simulate a multiuser ad hoc BI querying environment. Individual runtimes are listed in the "Appendix."

BASELINE TEST

As you can see from the baseline results in Figure 3, the majority of queries delivered subsecond response times; and even the most complex query, which involved the entire five-year span of data, ran in under four seconds.

The reporting and drill-down queries (267 milliseconds to 1.041 seconds) demonstrate the excellent ability of SAP HANA to aggregate data. For instance, the longest-running of these, R1-4, ran in

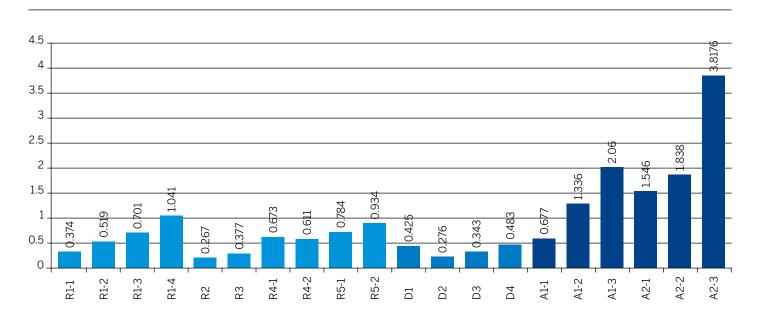
just over a second (1.041) but took only 2.8 times longer to crunch through 12 times more data than its monthly equivalent, R1-1.

The drill-down queries (276 to 483 milliseconds) demonstrate SAP HANA's aggressive support for ad hoc joins and, therefore, to provide unrestricted ability for users to "slice and dice" without having to first involve the technical staff to provide indexes to support it (as would be the case with a conventional database).

The analytic queries (677 milliseconds to 3.8 seconds) efficiently performed sophisticated joins (fact-to-fact joins, subqueries, CSQs, union all) and analysis across a sliding time window (year over year). The queries with one- to six-month date ranges (A1-1 through A2-2) ran in two seconds or less. Query A2-3, which analyzed the entire five-year date range, ran in under four seconds. All analytic query times are well within time frames to support iterative speed-of-thought analysis.

Across the board, the baseline tests show that SAP HANA scaled efficiently (better than linear) as the data volumes increased for a given query.







THROUGHPUT

The throughput tests are summarized in the table "Throughput Tests" and show that, in the face of increasing and mixed BI workloads, SAP HANA scales very well.

Throughput Tests	oughput Tests		
Test Case	Throughput (Queries per Hour)		
1 stream	6,282		
10 streams	36,600		
20 streams	48,770		
25 streams	52,212		

At 25 streams, the average query response time was less than three seconds and is only 2.9 times higher than at baseline (see the "Appendix"), an indicator of SAP HANA's excellent internal efficiencies and ability to manage concurrency and mixed workloads.

A rough estimate of BI user concurrency can be derived by dividing total queries per hour by an estimated average number of queries per user per hour. For instance, the 52,212 queries from the 25 streams per hour divided by 20 (a zesty per-user rate of one query every three minutes) provides a reasonable estimate of 2,610 concurrent BI users across a mixture of reporting, drill-down, and analytic query types.

AD HOC HISTORICAL QUERIES

The performance tests were intended primarily to demonstrate SAP HANA performance in customary real-time BI query environments that focus on analyzing business operations. However, ad hoc historical queries that analyze the entire volume of data that is stored may occasionally need to be run (and they should not impose a severe penalty). In addition to query A2-3 (included in the stream tests), queries R1-1, R4-2, and R5-1 were deemed potentially relevant for a five-year time window and were resubmitted, this time for the entire five-year date range. The table "Ad Hoc Historical Queries" shows that SAP HANA demonstrated better-than-linear scalability running against five times more data with a response time that was three times faster or more.

Overall, these queries all ran in a few seconds and further highlight the ability of SAP HANA to support real-time BI on massive amounts of data without tuning and without massive hardware or proprietary hardware components.

RESULTS SUMMARY

In these performance tests, SAP HANA demonstrated the ability to deliver real-time BI-query performance as workloads increased in terms of capacity (up to 100 TB of raw data), complexity (queries with complex join constructs and significant intermediate results run in less than two seconds), and concurrency (25-stream throughput representing about 2,600 active users).

Ad Hoc Historical Queries

Time Period	R1-1 (in Seconds)	R4-2 (in Seconds)	R5-1 (in Seconds)
1 year	1.1	1.4	1.1
2 years	1.8	2.4	1.8
3 years	2.6	3.4	2.4
4 years	3.3	4.4	3.0
5 years	4.1	5.4	3.6



SAP HANA demonstrated a compression rate of greater than 20 times.

SAP HANA performance is based on comprehensive, well-thought-out design elements, including, for example:

- An in-memory design
- Smart internal data structures (for example, native columnar, advanced compression, and powerful partitioning)
- A clever, cost-based optimizer that can meld these smart data structures into an efficient query plan
- Efficient query execution that smartly executes the query plan to take advantage of internal components (for example, advanced algorithms, multilevel caching optimization in the CPU, and hyperthreading).

Much has been written on the design of SAP HANA, and it is not necessary to rewrite it here. For more information on SAP HANA technical features, please refer to SAP HANA Technical Overview and SAP HANA for Next-Generation Business Applications and Real-Time Analytics.

REAL-WORLD EXPERIENCES

These tests were run in a simulated environment to isolate SAP HANA performance⁷; however, you can expect your own BI query performance to be significantly better on SAP HANA than on your existing conventional database management system – potentially thousands of times better. Most important, SAP customer testimonials confirm SAP HANA performance. Here are a few examples.

"We have seen massive system speed improvements and increased ability to analyze the most detailed levels of customers and products." – Colgate Palmolive

"[R]eplacing our enterprise-wide Oracle data mart and resulting in over 20,000-times speed improvement processing our most complex freight transportation cost calculation . . . our stand-alone mobile applications that were previously running on Oracle are now running on SAP HANA. Our 2,000 local sales representatives can now interact with real-time data instead and have the ability to make on-the-fly promotion decisions to improve sales." – Nongfu Spring

"[O]ur internal technical comparison demonstrated that SAP HANA outperforms traditional disk-based systems by a factor of 408,000." – Mitsui Knowledge Industry Co. Ltd.

"With SAP HANA, we see a tremendous opportunity to dramatically improve our enterprise data warehouse solutions, drastically reducing data latency and improving speed when we can return query results in 45 seconds from SAP NetWeaver BW on SAP HANA versus waiting up to 20 minutes for empty results from SAP NetWeaver BW on a traditional disk-based database." – Shanghai Volkswagen

Conclusion

Real-time BI query environments should be able to support new queries as soon as users formulate them. This speed-of-thought capability is only possible if consistently excellent performance is freely available in the underlying database platform.

These rigorous performance tests and their results presented here confirm that, without tuning and without massive hardware or proprietary hardware components, SAP HANA delivers leading performance and scale-out ability and enables real-time BI for businesses that must support a range of analytic workloads, massive volumes of data, and thousands of concurrent users.

SAP HANA delivers leading performance and scale-out ability and enables real-time BI for businesses that must support a range of analytic workloads, massive volumes of data, and thousands of concurrent users.



FOOTNOTES

- 1. 100 TB = raw data set size before compression
- 2. These tests did not involve SAP NetWeaver BW.
- 3. As defined with the "CREATE TABLE" command
- **4.** Compression rates depend heavily on characteristics of the actual data being compressed. Individual results may vary.
- **5.** A single query stream represents an application connection that supports potentially hundreds of concurrent users.
- **6.** This is done to prevent a query from being submitted by different streams at exactly the same time and to help ensure a practical workload mix (for example, drill downs happen more frequently than quarterly reports).
- 7. These performance tests were designed to isolate SAP HANA system-level performance independent of the variety of applications that may be used in SAP software environments. Therefore, these test results should be used only as a general guideline, and results will vary from implementation to implementation due to variability of platforms, applications, and data.

Appendix

Query Name	Test Case*	Average Runtime (Seconds)	Query Name	Test Case*	Average Runtime (Seconds)
R1-1	1 stream	0.374301625	R3	1 stream	0.3772025
	10 streams	0.7104476375		10 streams	0.687762475
	20 streams	1.11287894375		20 streams	1.03632401875
	25 streams	1.263543115		25 streams	1.28864399
R1-2	1 stream	0.5193995	R4-1	1 stream	0.67338575
	10 streams	0.8852610125		10 streams	1.15797815
	20 streams	1.2852265875		20 streams	1.7928188625
	25 streams	1.55048987		25 streams	2.081943225
R1-3	1 stream	0.70110375	R4-2	1 stream	0.61149025
	10 streams	1.098921025		10 streams	1.0192868125
	20 streams	1.6092578375		20 streams	1.48342026875
	25 streams	1.85689055		25 streams	1.751548475
R1-4	1 stream	1.0411995	R5-1	1 stream	0.784670625
	10 streams	1.5377812		10 streams	1.3914058
	20 streams	2.24272095		20 streams	2.017145075
	25 streams	2.53941444		25 streams	2.37001971
R2	1 stream	0.2671005	R5-2	1 stream	0.9343955
	10 streams	0.4740847375		10 streams	1.6150219
	20 streams	0.765801125		20 streams	2.326104475
	25 streams	0.94418651		25 streams	2.70342093

Query Name	Test Case*	Average Runtime (Seconds)	Query Name	Test Case*	Average Runtime (Seconds)
D1	1 stream	0.4254266875	A1-2	1 stream	1.3362615
	10 streams	0.67727600625		10 streams	2.0676603
	20 streams	0.9883220375		20 streams	3.0319986125
	25 streams	1.1591933925		25 streams	3.62353857
D2	1 stream	0.2762340625	A1-3	1 stream	2.0605925
	10 streams	0.49984061875		10 streams	3.15405455
	20 streams	0.7925614875		20 streams	4.469254975
	25 streams	0.9211584		25 streams	5.01461024
D3	1 stream	0.3439448125	A2-1	1 stream	1.5460765
	10 streams	0.6137608375		10 streams	2.7397760625
	20 streams	0.969213509375		20 streams	3.99316446875
	25 streams	1.1464113225		25 streams	4.90314213157894
D4	1 stream	0.48321675	A2-2	1 stream	1.838777
	10 streams	0.88516171875		10 streams	3.256667375
	20 streams	1.345513390625		20 streams	4.5519178125
	25 streams	1.56276527		25 streams	5.32557834210526
A1-1	1 stream	0.67762625	A2-3	1 stream	3.816785
	10 streams	1.0925558		10 streams	6.371394
	20 streams	1.77672375		20 streams	9.88785475
	25 streams	2.03975081		25 streams	11.9014413333333

^{*10-, 20-,} and 25-stream runs reflect 10 milliseconds of think time between queries.

