

WHITE PAPER

W I N T E R C O R P O R A T I O N

SAP BUSINESS INFORMATION WAREHOUSE

MULTI-TERABYTE EVALUATION & FEASIBILITY TEST

*Testing the Performance & Scalability of
SAP's Business Intelligence Warehouse*

Specialists in the World's Largest Databases



www.wintercorp.com

SPONSORED
RESEARCH
PROGRAM

SAP BUSINESS INFORMATION WAREHOUSE

MULTI-TERABYTE
EVALUATION &
FEASIBILITY TEST

*Testing the Performance & Scalability of
SAP's Business Intelligence Warehouse*

by David J. DiMare and
Richard Winter



WINTER CORPORATION
411 WAVERLEY OAKS ROAD, SUITE 327
WALTHAM, MA 02452
617-695-1800

Executive Summary

Business intelligence (BI) applications contribute strategic value by building integrated data warehouses efficiently and providing analysts with the tools needed to transform the data into informed business decisions. As the business and IT leadership of the enterprise seeks to extend the scope and benefits of BI applications, the issue of *scalability* becomes critical. The key challenge is ensuring that BI applications have the capacity to meet both current and future business requirements. For example, will the BI application provide high performance when the data warehouse environment grows to thousands of users accessing terabytes (TBs) of data? The answer to this question cannot be extrapolated from performance in an existing BI environment because of the unique complexities inherent in creating and managing very large data warehouses. The only way to assess BI scalability is to develop and conduct performance tests on the BI application itself.

To help its customers align technology solutions with business goals, SAP recently completed such an evaluation of its SAP Business Information Warehouse (SAP BW) solution. After conducting a survey of SAP BW customers to establish their current and future BI requirements, SAP designed a series of tests for SAP BW. These tests included building a 5.5TB data warehouse containing 1.1TB of business data, and simulating thousands of users executing many queries against the data. The objectives were to measure SAP BW's ability to build such a large data warehouse, and to deliver rapid query response. This evaluation project was conducted jointly by SAP and partners Sun Microsystems and Oracle.

The evaluation tests were conducted in two phases. The first phase tested the ability of SAP BW to import and load the data and create the data warehouse. The second phase simulated the equivalent of thousands of users issuing queries against the data warehouse; in one test, queries accessed aggregate values, and in another, queries accessed detailed data as well.

Winter Corporation was retained by SAP to conduct an extensive, independent audit of the SAP BW evaluation. Winter Corporation is a consulting firm specializing in large database technology and implementation, with in-depth expertise in database performance and scalability. After completing the audit, Winter Corporation certified that the tests were executed properly and successfully.

As a BI solution, SAP BW provides value in several important areas.

- SAP BW delivers built-in business content models, schemas and ETL (extraction, transformation, and load) capability for the enormous range of business information created and maintained by SAP products. The data warehouse administrator can customize these models to meet specific business objectives. SAP BW can automatically import data from a wide range of sources and create multiple types of data warehouse objects for efficient access. Examples include multidimensional InfoCubes and pre-calculated aggregates, or summaries, of data values.
- SAP BW successfully loaded and built the large-scale data warehouse. The product

executed this extensive and demanding phase without any major incidents, an impressive result. Building such a large data warehouse is often accomplished only with difficulty.

- The query simulation tests showed that SAP BW could handle over 600 concurrent in-flight queries while maintaining very good interactive response time. The latest Winter Corporation *Large Database Research Report* shows that most companies have not yet achieved this level of concurrency in decision-support applications, even during peak periods of activity. This volume of concurrent, in-flight queries is equivalent to an interactive workload generated by thousands of concurrent users, depending on the specific data warehouse environment.

By developing and completing this suite of performance tests, SAP has highlighted its commitment to deliver application scalability and performance for very large BI data warehouse environments with thousands of users. The successful completion of this evaluation is a major milestone in validating the scalability and performance of the SAP Business Information Warehouse.

Table of Contents

1. INTRODUCTION1

2. EVALUATION DESCRIPTION3

2.1. OVERVIEW.....3

2.2. OPERATING ENVIRONMENT.....4

2.3. SAP INFORMATION MODEL.....4

2.4. TEST QUERIES5

3. EVALUATION TESTS AND RESULTS.....7

3.1. DATA WAREHOUSE LOAD TEST.....7

3.2. DATA WAREHOUSE LOAD RESULTS.....7

3.3. QUERY SIMULATION TEST8

3.4. QUERY SIMULATION TEST RESULTS9

3.5. OTHER OBSERVATIONS.....10

4. SUMMARY AND CONCLUSIONS.....12

APPENDIX A. DESCRIPTION OF TEST QUERIES.....13

1. Introduction

Most organizations today recognize the strategic value of business intelligence (BI) applications. BI applications encompass the activities required to collect and analyze data from multiple sources, transform the data into useful information, and distribute the information to the right people at the right time in the right format. Effective BI applications integrate data from many sources, build BI databases quickly and efficiently, and provide tools for analyzing the data. BI implementations typically create data warehouses, data marts, operational data stores, or a combination of these to meet unique business requirements.

A key question is what happens when the volume of data grows to multiple terabytes and the number of users increases to many thousands? Enterprises are finding that to achieve their business goals, they must have access to ever-growing volumes of very detailed data about the business. They must also make this information available to increasingly large user populations, including employees, partners, suppliers, and customers.

System capacity requirements are growing even larger than might be imagined from the volume of business data alone. The amount of actual data-storage capacity required to store business data in a data warehouse format, for example, can be 5 or 6 times the size of the business data itself, according to the Winter Corporation *Large Database Research Report*. This additional space includes indexes, aggregates, and summaries designed to speed access to the data. As a result, one terabyte (1TB) of business data would typically result in a 5-6TB data warehouse. Such large data volumes are increasingly common in many industries, such as banking, insurance, telecommunications, transportation, retail sales, and healthcare.

Just building and loading a large, integrated, and multi-terabyte data warehouse requires scalable, high-performance processing. And once the data warehouse is built, users must be able to quickly search through and analyze massive volumes of data. Techniques such as creating multidimensional data structures and summarizing detailed data into aggregates can improve query performance on very large databases. However, it is difficult or impossible to predict the results for a particular environment, or to identify performance challenges, without conducting tests to simulate hundreds of requests for business information coming into the data warehouse at the same time. As we at Winter Corporation have often cautioned, what works for hundreds of gigabytes of data often will not work for many terabytes of data. Therefore, the leadership of the enterprise needs a BI architecture that can grow as the organization's business needs evolve.

While there are many industry-standard benchmarks for testing the performance and scalability of underlying resources—hardware, operating systems, and relational database management systems (RDBMS)—an organization cannot rely solely on these benchmarks to know what will happen in a real-world BI environment. Other critical factors include the performance and scalability of the application itself, how well the application optimizes its underlying resources, and what the expected user workload looks like.

Because BI applications are a separate layer of software on top of the RDBMS, it is important to develop and conduct performance evaluations at the application level. SAP, as a major business solution provider, clearly understands this. The company's BI solution is mySAP Business Intelligence, which includes SAP Business Information Warehouse (SAP BW) as a core component. SAP BW delivers an enterprise-wide business intelligence platform and a suite of BI tools. It includes an information model; extraction, transformation, and loading (ETL) capabilities for moving source data into SAP BW; and the ability to create data warehouses, operational data stores (ODS), and multidimensional models of data (see Figure 1).

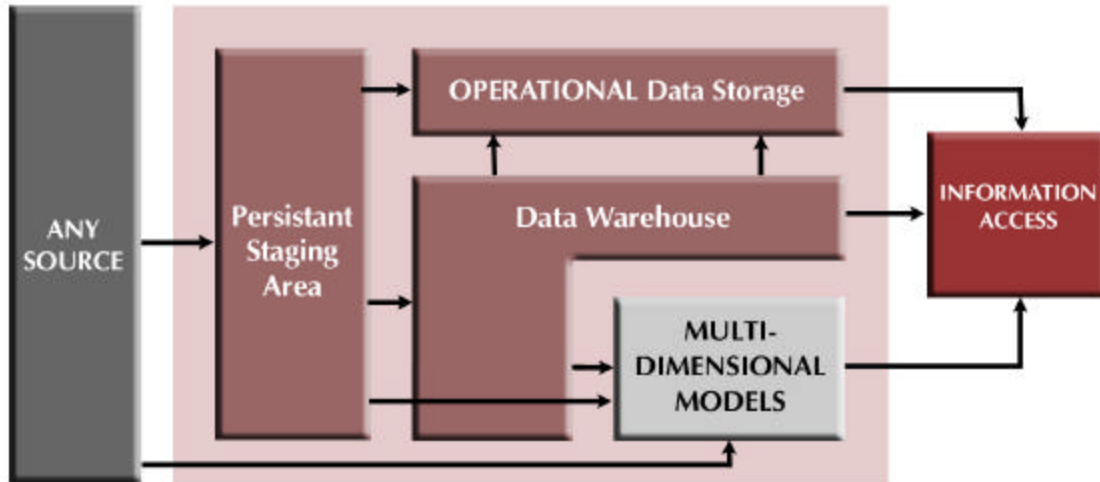


Figure 1. SAP Business Information Warehouse

SAP recently conducted a survey of customers using SAP BW to establish current and future BI requirements. Survey results indicated that the SAP BW user base is experiencing rapid increases in both data volumes and user populations. A growing number of installations have databases approaching 1TB of data, and many are already in the 1TB-2TB range. To address its customers' business and technology objectives, SAP created a test project to assess SAP BW's current scalability, and to identify areas for future enhancements to ensure that SAP BI products will continue to meet user requirements.

The evaluation project was both a feasibility study and a demonstration of SAP BW's performance in an environment that simulated thousands of users accessing a multi-terabyte data warehouse. SAP conducted the project jointly with partners Sun Microsystems and Oracle.

2. Evaluation Description

2.1. Overview

SAP's goals for the evaluation project were the following:

- Demonstrate the extent to which SAP BW can currently scale to handle a large BI environment in terms of both data volume and number of users.
- Use the evaluation as an opportunity to develop and refine strategies for optimizing the data warehouse build-and-load process for terabytes of data.
- Gain valuable experience working with a large test installation and use it to enhance current and future releases of SAP BW.

SAP generated a test data environment that was much larger than the average found in the customer survey. The test data warehouse represented 10 years of sales data for 100,000 customers, 200,000 products, 100 sales organizations, 20 sales divisions, 42 countries, and 8 distribution channels. Once built, the data warehouse totaled 11.7 billion sales records and required 5.5TB of data-storage capacity.

The 5.5TB included 1.1TB of detailed business data plus indexes, aggregates, and temporary data storage facilities used while loading data. To put this test data environment in context, the test represented a volume of business information that is typical of a large data warehouse in a large enterprise today. In addition, the 5.5 ratio of data-storage capacity to detailed business data represents a real-life ratio as described in the section above. According to SAP, this is the largest data warehouse built to date with SAP BW.

The evaluation was run in two phases. The first phase involved loading the BI data warehouse from multiple data sources and then building aggregates and indexes. This process was designed to represent a real-world scenario by creating and populating a range of different data objects, including multidimensional data structures (InfoCubes), ODSs, indexes, and aggregates. (These objects are briefly described in the "SAP Information Model" section below.) The goal here was to see if SAP BW could build the 5.5TB data warehouse without experiencing significant problems.

The second phase simulated thousands of users executing queries against the data. The primary objective was to calculate average response time per query. In the first portion of the test, three types of users issued "light" queries against aggregates. In the second portion of the test, SAP simulated a more mixed workload by including some "heavy" queries against detailed data (InfoCubes) in addition to "light" queries against aggregates.

The evaluation tests were conducted at SAP's offices in Walldorf, Germany, and at Sun Microsystems' Garrison Lab in Palo Alto, California. Winter Corporation conducted an extensive, independent audit of the evaluation and certified that the tests were completed properly and successfully. Please reference the Winter Corporation Audit Letter and the

SAP/Sun Microsystems technical white paper on the project for more details. These documents can be found on SAP's web site, <http://service.sap.com/bw>, in the sub-section titled "Performance." The technical white paper also provides additional information on SAP BW capabilities.

2.2. Operating Environment

Evaluation software included SAP Business Information Warehouse version 3.0B running on the Solaris 2.9 operating system with Oracle 9i version 9.2.0 and Veritas Volume Manager 3.2. The evaluation was run on a Sun Fire 15000 server, a high-end server with 72 900-MHz UltraSPARC III processors and 144GB of main memory. This "central" server functioned as both the data warehouse server and the SAP BW application server.

A second "driver" server was used to drive the query workload. This was a Sun Fire 4800 Server with 12 900-MHz processors and 24GB of main memory. A private network connection was established between the driver server and central server to eliminate interference from network traffic not associated with the evaluation.

Twenty-four Sun StorEdge T3 RAID disk arrays were connected to the central server via fibre channel to support data storage requirements. Each T3 contained nine 73GB disks and was configured as a single RAID device via the hardware RAID controller. Veritas Volume Manager was used to manage logical volumes defined across the disk arrays.

The 12 disk arrays used to support the Oracle database were set up as RAID-0 devices. A RAID-0 configuration provides data striping across disks but has no data redundancy (neither parity or mirroring). I/O performance is enhanced because read and writes are done in parallel across devices.

2.3. SAP Information Model

A key strength of SAP BW is the fact that it delivers built-in business content models, schemas and ETL (extraction, transformation, and load) capability for the enormous range of business information created and maintained by SAP products. The database administrator can extend these models to incorporate new subjects and other information. SAP BW can automatically import data from a wide range of sources and create multiple types of objects for storing the data. Here are some of the objects created as part of the evaluation project.

- *A Persistent Staging Area (PSA)* is a set of database tables into which extracted source data is initially stored for subsequent moves into other SAP BW data objects. Data is loaded into the PSA using SAP BW's ETL capabilities. Data can be extracted from other SAP components, supported RDBMSs, flat files, XML files, and other data sources via the SAP BW open interface (see Figure 2).

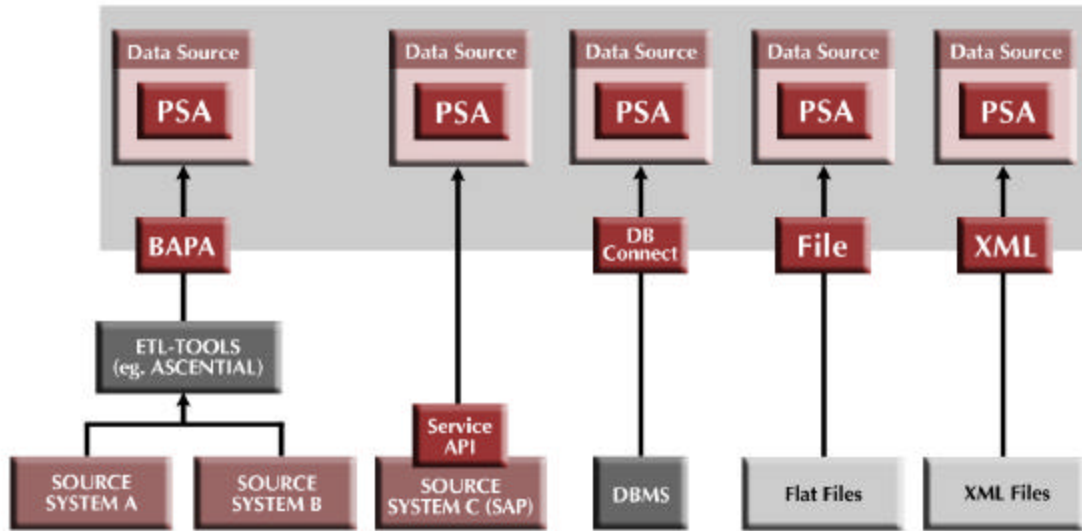


Figure 2. SAP BW supports a wide range of data sources

- An *Operational Data Store (ODS)* describes a consolidated dataset from one or more sources. Data in ODS objects is stored in flat data warehouse tables. ODS data can then be used to populate InfoCubes or other ODS objects. Each of four ODS objects created for the evaluation contained six months of sales data for all distribution channels, representing 584 million records.
- An *InfoCube* represents detailed data stored in a star-schema data warehouse design. A star schema is comprised of one large fact table (e.g., sales orders) and multiple dimension tables (e.g., region, year, and distribution channel). Therefore, a star schema is called a “multidimensional” data structure. With InfoCubes, users can analyze the same information from various business perspectives. The evaluation test created 80 InfoCubes based on 10 years of sales data for 8 distribution channels (one InfoCube for each year and distribution channel = 10 x 8 or 80). Each InfoCube contained about 145 million records.
- An *aggregate* is a summarized view of the data in an InfoCube. Aggregates dramatically improve the response time of queries because calculations are pre-built and stored as summaries. Ten aggregates were created for each InfoCube as described in “Data Warehouse Load Results” below.

2.4. Test Queries

For the query simulation test, SAP used 15 reports provided as business content with SAP BW 3.0B. See Appendix A for a description of these reports. Queries against aggregates represented “light” queries. Queries against detailed InfoCube data represented “heavy” queries.

The input parameters for the queries, including which InfoCube was used to satisfy the

request, were chosen at random. When a parameter was required, a range of values was provided. The low-end value was chosen at random and the high-end value was calculated by adding a range offset to the starting point:

Input Parameter	Range
Customer	10
Division	5
Material	50
Sales Organization	10
Calendar Month	3

This range value determined the number of items that appeared on every report (e.g. sales by customer would have a report showing 10 customers).

The randomization of selection criteria was a positive step in simulating a multi-user business intelligence environment. In future tests, randomizing the range offset as well would provide variable length reports in addition to randomized input criteria, thus increasing the performance challenge.

3. Evaluation Tests and Results

As stated above, the evaluation was divided into two phases: the load of the data warehouse, and the query test simulations.

3.1. Data Warehouse Load Test

The load test phase consisted of the following steps:

- Load master data into the PSA. This included the data for customers, products, sales organizations, sales divisions, countries, and distribution channels.
- Load detailed sales transaction data into the PSA. This data represented monthly sales for 10 years across the 8 distribution channels. There were approximately 400,000 transactions per day for a total of 11.7 billion records.
- Load 80 InfoCubes from the PSA. Each cube represented sales data for a single distribution channel for one year, or approximately 145 million records.
- Load the Operational Data Stores (ODS) from the PSA. Four ODSs were created to store the most recent two years of data in flat database tables rather than in multidimensional InfoCubes. Each ODS contained six months of data. While the ODSs were not accessed during the query tests, it was important to include the ODSs as part of the load test to accurately represent a typical customer BI environment.
- Generate aggregate views of the detailed data. SAP defined ten aggregates per InfoCube. Four were used as interim aggregates to generate six additional aggregates used in the query tests. These six aggregates ranged in size from 12 to 200,000 rows.

The resulting data warehouse occupied 5.5TB of storage, including 1.1TB of data stored in InfoCubes, 0.2TB in aggregates, 0.4TB in ODS objects, and 0.6TB in RDBMS indexes.

3.2. Data Warehouse Load Results

Loading and building a BI data warehouse is a complex process involving many steps. Often, in a large-scale data warehouse, this is accomplished only with difficulty. In contrast, SAP BW executed this extensive and demanding phase without any major incidents. Figure 3 below shows the results measured for each step and the total load/build time.

SAP BW was clearly successful in this phase of the evaluation. The product adds value here with its built-in information model and the ability to automatically import data from many different data sources.

Load/Build Step	Results
<i>Load master/transaction data into PSA</i>	<i>44:50:25 hours</i>
<i>Create PSA indexes and Oracle statistics</i>	<i>35:00:00 hours</i>
<i>Load data from PSA to InfoCubes and ODSs</i>	<i>65:17:19 hours</i>
<i>Create InfoCube indexes and Oracle statistics</i>	<i>09:12:23 hours</i>
<i>Build aggregates</i>	<i>14:14:45 hours</i>
<i>Create indexes and statistics on aggregates</i>	<i>00:52:12 hours</i>
Total Load/Build Time	<i>169:27:04 hours</i>

Figure 3: Summary of Data Warehouse Load Times

3.3. Query Simulation Test

In the query simulation tests, 15 queries in the form of sales reports were submitted from a loop running on the driver server. Using the Web interface of SAP BW, the driver sent the reports as http requests to the Web Application Server on the central server. The driver established multiple parallel loops to simulate submitting hundreds of report requests simultaneously.

As described earlier, the driver was designed to choose report parameters, including the InfoCube used to satisfy the request, at random for each loop. However, the parameters within a single loop remained the same. This simulated a user accessing multiple reports for a single unit of data (e.g. sales for a customer, followed by credits issued to the same customer). Two simulation tests were measured.

Test #1. In the first simulation test, 660 parallel connections processed 5 loops each of all 15 sales reports for a total of 49,500 reports. The queries were tuned to access aggregate views generated from the InfoCubes. Each aggregate ranged in size from 12 to 200,000 rows and the total amount of aggregate data accessed was 7GB. None of the reports ran against detailed data.

Test #2. In the second simulation, 550 connections were established. Of these, 540 processed reports that accessed aggregate views, as in the previous run. In the remaining 10 connections, one of the fifteen reports was modified to obtain results by accessing detailed data stored in an InfoCube.

The 540 connections unchanged from the prior simulation each processed 5 loops of 15 reports. The remaining 10 connections ran 2 loops per connection. This provided a total of 40,800 reports, of which 20 reports accessed detailed InfoCube data. A single InfoCube contained approximately 145 million rows of data.

3.4. Query Simulation Test Results

Test #1. In the first test, the average response time per report was 2.5 seconds. Server utilization was high, averaging 85%.

This simulation showed that SAP BW could handle a large number of concurrent light queries with very good interactive response time. Since the simulation ran with minimal think time between queries, it is likely that all 660 connections were submitting requests simultaneously and generating an estimated equal number of concurrent, in-flight queries. The latest Winter Corporation Large Database Research Report shows that most companies have not yet achieved this level of concurrency in decision-support applications, even during peak periods of activity.

UNIX DSS RESPONDENTS' PROJECTED WORKLOAD

(concurrent, inflight queries) (average)

Four years (1/01 - 12/04)	162
Two years (1/01 - 12/02)	122
1/2001	101

Growth Rate (average)

Four years (1/01 - 12/04)	60%
Two years (1/01 - 12/02)	21%

Figure 4: Expected Average Number of Concurrent, In-flight Queries for Unix-based Decision Support Systems

Source: Winter Corporation Large Database Research Report © 2003

In Figures 4 and 5, Research Report survey participants projected both the average and the maximum (peak) number of concurrent, in-flight queries they expected for Unix-based decision-support systems at the end of 2002 and at the end of 2004. These numbers represent a mix of light, medium, and heavy queries. Test #1 of the SAP BW evaluation demonstrated 660 concurrent queries, most of them light. The frontier identified in the survey is the same number of concurrent queries with a substantial percentage being medium and heavy queries.




Figure 5: Online Workload Frontier for Unix-based Decision Support Systems

Source: Winter Corporation Large Database Research Report © 2003

Test #2. In the second test, the average response time for light queries accessing aggregates increased from 2.5 seconds to 5 seconds. Average response time for reports accessing detailed data was 335 seconds, or 5.6 minutes, significantly longer than the other 14 reports. Therefore, although it was not possible to determine which reports were running at any given point, it is likely that the 10 connections with reports accessing detailed data were all processing heavy I/O reports at the same time. As a result, this put a significant load on the storage subsystem. For this test, an average of 39 I/Os per second per disk was achieved with 85% utilization. There was more I/O wait time and server utilization dropped from 85% to 67%.

This test showed that SAP BW could support a significant number of concurrent connections in a more mixed workload while maintaining an average response time of 5 seconds for light queries. This would be highly satisfactory for most data warehouse applications of this type. A large number of reports were handled while servicing a heavy workload on the storage subsystem. As in the first test, the 550 connections generated an estimated equal number of concurrent queries in-flight. A volume of 550 to 660 concurrent, in-flight queries is equivalent to an interactive workload generated by thousands of concurrent users, depending on the specific data warehouse environment.

3.5. Other Observations

The evaluation also confirmed the following observations for the SAP BW architecture.

- A multi-tier architecture facilitates the distribution of workload over multiple application

servers, balancing capacity among the tiers and increasing query throughput.

- The data model needs to allow for a high degree of parallelism and the ability to easily identify independent tasks that can be run in parallel. For example, parallelism across different objects is more efficient than within one object. In the load test, partitioning data on a logical level and distributing data across multiple InfoCubes was key to reducing data-warehouse load time.
- The creation of suitable aggregates can dramatically improve query performance, but also increases load time. However, the resulting load time in the test was reasonable for such a large data warehouse; building aggregates took about 14 hours, or a little over 8% of the total load time. In addition, SAP's use of hierarchical aggregates effectively reduced the time needed to build the aggregates.
- The pre-generation of keys (that is, loading the dimensions first) speeds up the load of detailed data into the fact tables and can significantly reduce load time.

4. Summary and Conclusions

Completing the evaluation project described in this paper demonstrates SAP's commitment to application scalability and performance in building very large BI databases and providing timely data access to thousands of users. A major goal was to provide important scalability information for customers. In addition, at SAP's request, Winter Corporation conducted an extensive, independent audit of the SAP BW evaluation and certified that the tests were completed properly and successfully. The evaluation also represents a valuable learning experience that will enhance future releases of SAP BW and facilitate the design of future SAP BW evaluation efforts.

A key evaluation finding is the fact that SAP successfully created the 5.5TB data warehouse, with 1.1TB of business data, without any noteworthy problems. There are several observations here. First, this volume of data is typical of a large data warehouse in a large enterprise today. Second, building such a large data warehouse is often accomplished only with difficulty, unlike what was experienced in the SAP BW evaluation. And third, it took SAP only a little over a month to configure and implement the evaluation data warehouse. This demonstrates the value of the integrated SAP BW solution.

A second result of the evaluation was SAP BW's ability to handle over 600 concurrent in-flight queries while maintaining very good interactive response time. The latest Winter Corporation *Large Database Research Report* shows that most companies have not yet achieved this level of concurrency in decision-support applications, even during peak periods of activity. This volume of concurrent, in-flight queries is equivalent to an interactive workload generated by thousands of concurrent users, depending on the specific data warehouse environment.

In the future, we expect SAP to refine its evaluation criteria to test additional SAP BW capabilities, such as executing complex ad hoc queries and the effect of creating and using multidimensional aggregates.

In summary, the results show that SAP BW can support thousands of concurrent users accessing a multi-terabyte data warehouse. This accomplishment is a major milestone for SAP in demonstrating the scalability and performance of the SAP Business Information Warehouse.

Appendix A. Description of Test Queries

Report Title	Input Parameters	Results	Used InfoCube or Aggregate	InfoCube or Aggregate Row Count
Product Analysis	Product Number	Orders by product	Aggregate	200,000
Sales Volume	Customer Number	Sales by customer	Aggregate	100,000
Credit Memos	Customer Number	Credits by customer	Aggregate	100,000
Incoming Orders	Customer Number	Orders by customer	Aggregate	100,000
Customer Returns	Customer Number	Sales & returns by customer	Aggregate	100,000
Customer Statistics	Customer Number	Sales, orders, returns, & credits by customer	Aggregate	100,000
Customer Returns	Month / Year	Orders & returns by calendar month	Aggregate	12
Customer Incoming Orders	Sales Organization & Division	Sales and credit amounts by sales organization & division	Aggregate	2,000
Order Entry Analysis	Sales Organization & Month / Year	Orders by calendar month for a sales organization	(1) Aggregate	12
			(2) InfoCube	145,000,000
Product Profitability Analysis	Product Number	Sales volume by product	Aggregate	200,000
Top Customers	Customer Number	Sales, orders, & returns by customer	Aggregate	100,000
Incoming Order Analysis	None	Orders by calendar month for 1 year	Aggregate	12
Sales Volume/ Cost Analysis	None	Sales by calendar month for 1 year	Aggregate	12
Distribution Channel Analysis	None	Sales by distribution channel	Aggregate	8
Sales Data Overview	None	Sales & orders by division	Aggregate	20

(1) Aggregate used for Simulation Test # 1

(2) InfoCube used for Simulation Test #2

*A leading center of expertise in very large databases,
Winter Corporation provides services in
research, consulting and education.*

*We help users and vendors understand their opportunities;
select their database and data warehouse platforms;
define and measure the value of their strategies, architectures and products;
plan, architect and design their implementations;
and manage their scalability, performance and availability issues.*

Our focus is databases near, at and beyond the VLDB frontier.



WINTER CORPORATION

411 WAVERLEY OAKS ROAD, SUITE 327
WALTHAM, MA 02452
617-695-1800

visit us at www.wintercorp.com