

Searching for the Best System Configuration to Fit Your BW Needs?

The New BW Enhanced Mixed Load Benchmark Makes It Easy to Find

by Dr. Thomas Becker and Tobias Kutning, SAP AG



Dr. Thomas Becker

(th.becker@sap.com) joined SAP in 2001 and is now a Quality Manager responsible for SAP NetWeaver Business Warehouse and SAP HANA performance in the In-Memory Platform Performance and Quality Engineering organization.



Tobias Kutning (tobias.kutning@sap.com)

joined SAP in 2001, worked as a database support consultant until 2007, and joined SAP IT in 2008. Since 2010, he has worked on the Performance & Scalability Team, responsible for SAP Standard Application Benchmark Product Management.

SAP customers searching for information about how to enhance the performance and scalability of their SAP NetWeaver Business Warehouse (SAP NetWeaver BW) implementations have always turned to results of SAP's mixed load (BI-MXL) benchmark for help. The BI-MXL benchmark, like all SAP standard application benchmarks, was designed to represent relevant, real-life scenarios involving various SAP business applications to help customers find the most appropriate hardware configurations to support their needs.

However, there are now many changing requirements related to improving decision making, like the demand for instant availability of the latest business data. This also means that the aspects customers are looking at when considering which hardware configuration is the right fit for their SAP NetWeaver BW implementation has begun to evolve as well. In response, SAP decided to develop a new set of metrics and a new benchmark: the BW enhanced mixed load (BW-EML) benchmark.¹

3 Business Requirements from SAP NetWeaver BW Customers

These days, when it comes to systems that support decision making, customers are looking for ones that can enable:

- **Near real-time reporting.** To make informed decisions in a fast-paced world, the ability to get instant information from analytical business applications is crucial. Not only do companies require quick access to information,

they also need this information to include up-to-the-minute details. Smart meter analytics and trade promotion management are just two examples of business processes that rely on near real-time reporting.

- **Ad hoc reporting.** Data volumes in enterprise data warehouses have grown significantly over the past few years due to the increased complexity of business data models and the level of detail captured in data. The sheer volume of this data and the demand for unpredictable navigation patterns make it impossible to use standard techniques like pre-aggregation to speed up query response times. Modern analytical applications must allow users to navigate instantly through these huge amounts of data by providing extensive slicing-and-dicing functionality.
- **TCO reduction.** Since data warehouses can contain hundreds of terabytes of data, it is crucial to minimize data redundancy, while at the same time maintaining layered data models. With SAP NetWeaver BW 7.30, it is possible to run reports directly on DataStore Objects (DSOs), which helps reduce TCO by saving precious storage space. DSOs are the core building element of a layered scalable architecture (LSA). Since reports can now analyze data in DSOs as fast as in multidimensional InfoCube data structures, InfoCubes have become completely obsolete in many reporting scenarios.

The BW-EML Benchmark: Built with BW Users in Mind

The BW-EML benchmark was developed especially with the database requirements of SAP

¹ This new benchmark can be performed with any database that is supported by SAP NetWeaver 7.30.

NetWeaver BW customers in mind. **Figure 1** compares the main features of the new BW-EML benchmark with those of its predecessor, the mixed load (BI-MXL) benchmark. Both benchmarks simulate a mix of multi-user reporting workload and the loading of delta data into the database simultaneously with user queries. Let's drill down further into the details of the benchmark.

Data Model

To ensure that the database being benchmarked can efficiently use both InfoCubes and DSOs for reporting, the BW-EML benchmark's data model consists of three InfoCubes and seven DSOs, each of which contain the data from one specific year. The three InfoCubes contain the same data (from the last three years) as the corresponding DSOs. Both types of objects consist of the same set of fields.

The InfoCubes come with a full set of 16 dimensions, which comprise a total of 63 characteristics, with cardinalities of up to one million different values and one complex hierarchy. To simulate typical customer data models, the InfoCube is made up of 30 different key figures, including those that require exception aggregation. In the data model of the DSOs, the high cardinality characteristics are defined as key members, while other characteristics are modeled as part of the data members.

Data Volumes

To test hardware configurations of various sizes, the BW-EML benchmark can be executed with different data volumes. The smallest configuration defined in the benchmark rules starts with an initial load of 500 million records (50 million records for each InfoCube and DSO). The records are loaded from ASCII flat files with a total record length of 873 bytes each. Larger volume configurations of the BW-EML benchmark include initial load volumes of one billion, two billion, or even more records.

In each of the mentioned configurations, the total number of records that need to be loaded, in addition to the initial load during the benchmark run, is one thousandth of the initial amount of records. The high load phase of the benchmark must run for at least one hour. During this time, the delta data must be loaded in intervals of five minutes. Additionally, the same number of records must be loaded in each InfoCube and DSO.

	Mixed load (BI-MXL) benchmark	Enhanced mixed load (BW-EML) benchmark
Loading of delta requests	<ul style="list-style-type: none"> Every 20 minutes 	<ul style="list-style-type: none"> Every 5 minutes
Ad hoc reporting capabilities	<ul style="list-style-type: none"> Predefined navigation paths Only drills down with the same characteristics set in display 	<ul style="list-style-type: none"> Randomized navigation paths Changes displayed characteristics
Reduction of TCO	<ul style="list-style-type: none"> All benchmark queries are defined in InfoCubes 	<ul style="list-style-type: none"> Uses DataStore Objects (DSOs) and InfoCubes for reporting

Query Model

For the BW-EML benchmark, eight reports have been defined on two MultiProviders — one for the three InfoCubes, and another for the seven DSOs. The respective reports on both MultiProviders are identical. This leads to two sets of four reports each. The four reports are categorized as follows:

- Report Q001: Customer-based reporting
- Report Q002: Material-based reporting
- Report Q003: Sales area-based reporting
- Report Q004: Price-comparison reporting

The reports select data for one particular year, randomly picking the InfoCube or DSO that contains the data. Further navigation steps are performed within each report, each of which results in an individual query and database access point.

Although the first three reports share similar navigation patterns, the filter and drill-down operations are randomized to address the demand for ad hoc types of queries. To make sure that the benchmark accesses different partitions of data, random values for filter parameters are used. Additionally, a random choice of characteristics for drill downs or other slice-and-dice operations ensures that a huge number of different characteristic combinations are covered in a multi-user reporting scenario.

To guarantee a high degree of reproducibility of reporting results, characteristics are grouped

FIGURE 1 ▲ A comparison between the BI-MXL benchmark and the new BW-EML benchmark

SAP standard application benchmarks are designed to represent relevant, real-life scenarios involving various SAP business applications to help customers find the most appropriate hardware configurations to support their needs.

Now that the new BW-EML benchmark race has started, SAP hopes to see more results in the coming months, and can't wait to observe how partners master the new benchmark challenge.

A Glance at the First Certified Benchmark Results

SAP partner HP published the first results of the BW-EML benchmark, choosing SAP HANA 1.0 as the database for the benchmark environment and SUSE Linux Enterprise Server as the operating system. This marks the first time that SAP HANA has been used in an SAP standard application benchmark. HP achieved impressive performance numbers with its benchmark setup, including 65,990 ad hoc navigation steps per hour. The figure below shows further details about the benchmark results.

Hardware partner	HP
Benchmark location/date	May 13, 2012 in Houston, TX, USA
Throughput/hour (ad hoc navigation steps)	65,990
CPU utilization of database server	88%
CPU utilization of application server	28%
Operating system (all servers)	SUSE Linux Enterprise Server 11
Relational database management system	SAP HANA 1.0
Technology platform release	SAP NetWeaver 7.30
Configuration	One database server: HP DL580 G7, 4 processor/40 cores/80 threads, Intel Xeon Processor E7-4870, 2.40 GHz, 64 KB L1 cache and 256 KB L2 cache per core, 30 MB L3 cache per processor, 512 GB main memory One application server (dialog/update/message/enqueue): HP BL680 G7, 4 processor/40 cores/80 threads, Intel Xeon Processor E7-4870, 2.40 GHz, 64 KB L1 cache and 256 KB L2 cache per core, 30 MB L3 cache per processor, 512 GB main memory
Certification number	2012023

▲ Details about the BW-EML benchmark results from HP

by their respective cardinalities, and only characteristics of the same cardinality are considered for a randomized operation.

Multi-User Workload

A script controls the multi-user reporting workload in the BW-EML benchmark. With the script, the number of simulated users can be defined. Each simulated user logs on to the system and then executes all eight reports and ad hoc navigation steps consecutively, resulting in a total of 40 ad hoc navigation steps. After finishing all of the predefined steps, the user logs off and then starts the next loop with a new logon. Simulated users are ramped up at a pace of one user logon per second.

Once all configured users are running, the benchmark control environment automatically starts a process chain that controls the delta

load, which is scheduled every five minutes. After a high load phase of at least one hour, the simulated users are ramped down, and the delta loads finish.

A control program then checks if the correct number of records has been uploaded during the benchmark run and if the uploaded records are visible in the report results. The essential key figure that is reported for a benchmark run is the number of ad hoc navigation steps per hour that the database executes successfully.

It's Your Turn

Now that the new BW-EML benchmark race has started, SAP hopes to see more results in the coming months, and can't wait to observe how partners master the new benchmark challenge. For more information on SAP benchmarking, visit www.sap.com/benchmark. ■

© 2012 SAP AG. 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 AG 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 AG 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.

