



BUSINESS INFORMATION WAREHOUSE

BEx Front end Performance

*Performance Analyses of BEx Analyzer and Web Application
in the Local and Wide Area Networks Environment*

*Document Version 1.1
March 2002*

Contents:

Executive Summary	4
1. Introduction	4
2. Measurements	5
2.1 BEx Analyzer.....	6
2.2 Web Appliation.....	7
3. Example	8
4. Recommendations	9

Executive Summary

Based on typical key figures of different wide area network technologies the data transfer performance of the SAP Business Explorer Analyzer (BEx Analyzer) was compared to an equivalent SAP BW Web Application. The goal of this comparison was a recommendation which of the two reporting tools to use in a wide area network (WAN).

The tests, which are described in more detail below, have shown that in a WAN environment the Web Application can deliver query results by a factor of 4 to 20 faster than the BEx Analyzer. As an example, a query with 210 rows and 7 columns took 1.3 seconds in a local area network (LAN) and 3.6 seconds in a WAN using the Web Application. The BEx Analyzer required about 6 seconds in the LAN and 69 seconds in the WAN for the same query. Therefore, we recommend to use a Web Application as the standard reporting tool for the access of BW data through wide area networks.

Protocol round trips are the decisive factor for the total transfer time of data in a WAN. The Web Application requires exactly one round trip for the transfer of data without graphical elements, whereas the BEx Analyzer has a fixed cost share of 37 round trips and one additional round trip per 200 data cells. As a result, the transfer time for large queries in the BEx Analyzer grows much faster than in a Web Application.

When using graphical elements in a Web Application, attention should be paid to the fact that each graphics (e.g. function buttons, pie charts, etc.) is sent to the front end as GIF image and hence causes an individual round trip. The same applies to frames within a Web page. Therefore, the total number of graphics and frames has a significant impact on the query runtime.

Graphical elements will be cached by the front-end Web browser and therefore need not be transferred a second time when they are reused. If graphical representations of data, such as pie charts, are sent to the front-end, the SAP Web Application Server is in charge of creating the GIF images which then will be transferred to the Web Application. Using the BEx Analyzer, however, the original data is sent to the front-end using several round trips, dependent on the number of data cells, and the graphical representation is produced by Excel.

1. Introduction

By using SAP Business Information Warehouse (SAP BW) as a company-wide Business Intelligence solution, demands are made of reporting tools that are distinguished by using wide area networks (WAN) based on geographical distribution.

SAP BW offers two fundamental tools for access and analysis of data in the warehouse::

- Desktop oriented using the Business Explorer Analyzer (BEx Analyzer)
The BEx Analyzer covers the full range from tabular, formatted reporting to a highly interactive analysis application. Embedded in Microsoft Excel, users can analyze data on different levels of granularity, be it on a detail level or highly aggregated. Interaction can follow predefined navigation paths or can be user-defined, using options like filtering, drilling or slicing and dicing.
- Web oriented using the SAP Web Applications (created by the BEx Web Application Designer)
Web Applications provide the same analysis and interaction options as the desktop oriented reporting. However, they do not require any front end component to be installed on the end user PC. Access to the SAP BW Web Applications is provided through a standard Web browser (within the intranet, extranet or Enterprise Portal).

Important acceptance criteria for reporting tools includes the quality of performance when executing queries. In doing so, it is insignificant to the end user whether he or she has access by using a local network (LAN) or by using a wide area network. At the very least, results are expected within a few seconds when executing "simple" queries.

With regard to network technology, there are two factors that can significantly influence the performance of reporting tools:


- The *band width* of the transfer medium, meaning the maximum amount of data that can be transported per unit of time, and
- The *latency* for a "protocol round trip", meaning the time needed for a message to travel from the sender to the receiver and back again to the sender.

The following table describes the band widths and typical latencies for current network technologies.

Technology	Band Width [Kbit/s]	Latency [ms]
100 Mbit LAN	100,000	< 10
56K Modem	56	250
Terrestrial WAN	128 – 2,048	250
Satellite Connection	64 – 2,048	800

Table 1: Band width and latency of different network technologies

Obviously, latencies are reflected in the total query runtime when using modem and WAN connections and especially when using satellite connections. With satellite connections, the cause for the latency lies mainly with the distance that a signal has to cover on the way from the sender to the satellite and then back to the receiver.



It goes without saying that response times for queries also depend on the amount of data to be transferred. Here we differentiate between

- User data, meaning the values that are displayed in the cells of the results grid, and
- Metadata, meaning data that describes the display of user data (for example, formatting areas and hierarchy nodes).

This document describes the measurement results when determining the protocol round trips, as well as the amount of data transferred with typical query sizes and for different (simulated) network technologies. The measurements are executed, comparing the BEx Analyzer and Web Application reporting tools. Based on the measurement results, this document will conclude with the most important results and a recommendation for using front ends in a wide area network.

2. Measurements

All executed measurements were run on an internal BW test system with release status BW 3.0A (Support Package 1 for Web Applications or Support Package 5 for BEx Analyzer in connection with note 0501839). The SAPGUI front-end version 6.20 PL 7 was used. And the "Low Speed Connection" flag was set when using the BEx Analyzer.

To determine a typical query size, several thousands of queries in about 5 customer systems were analyzed. These analyses yielded an average transfer volume of slightly more than 1000 cells. For our measurements, a typical volume of 1500 cells was specified to simulate, among other things, the transfer of additional metadata.

To simulate the different latencies and band widths, the SAP programs NetPerf and SumRfcTrace as well as the network protocol analysis tool EtherReal (<http://www.etherreal.com>) were used.

2.1 BEx Analyzer


Reporting with the BEx Analyzer is organized into two phases. In the first phase, a registration of the BEx Analyzer occurs in the BW system. During the second phase, data is requested and transferred back to the BEx Analyzer.

In order to measure the transfer costs for the fixed costs of a query, an inquiry is initially used that returns an empty result set (meaning zero cells). In doing so, the part of the data traffic was also measured that is exclusively needed for registering the front end. In order to account for the linearity of the data traffic, additional queries are measured that have 1500 and 10000 result cells, respectively. Target parameters that still needed to be measured included the number of protocol round trips and the number of transferred bytes.

The following table and graphic displays the measurement results for this scenario:

Query	No. of Cells	Round Trips	Bytes
Registration Share	--	33	66,486
Fixed Costs (Incl. Registration)	0	70	100,687
Typical (Incl. Fixed Costs)	1,500	75	113,386
Size (Incl. Fixed Costs)	10,000	84	166,866

Table 2: Key figures of BEx Analyzer



If one considers the total latencies of a query, 37 protocol round trips arise for the fixed costs share alone without registration. While a local network has a total delay time of only about 370 milliseconds through this process, this same delay in a wide area network, depending on the technology, is about 9 to 30 seconds:

Technology	Latency [ms] of 37 Round Trips
100 Mbit LAN	370
56K Modem	9,250
Terrestrial WAN	9,250
Satellite Connection	29,600

Table 3: Latency time induced by fixed costs

With several query executions within one Excel session, metadata that has previously been loaded is re-used. In doing so, the fixed costs share for queries decreases. A quantification of this decrease in costs depends, however, on the degree of reusability and was, therefore, not measured.

2.2 Web Appliation

When reporting with the help of a Web Application, the first authentication information is likewise transferred. Furthermore, the transfer of formatting information (cascading style sheets), Java scripts, and GIF screen information occurs in the first phase. These fixed costs occur due to the start of a session, meaning they occur along with the first query after opening the browser. By using suitable browser-cache settings, cascading style sheets and GIF screen information can be stored longer, so that additional optimization options are available.

Similarly to the BEx Analyzer, three queries with 0, 1500, and 10000 results cells were also measured in this attempt. The target parameters were, once again, the number of necessary protocol round trips as well as the size of HTML information in bytes, which were created or needed to be transferred.


At the time of the attempt, the SAP Basis Software Version 6.10 was available. Part of this Basis is the Web Application Server and the ICMAN. With this software status, an automatic compression of HTML was not yet possible. Through the compression, the generated HTML code is packed before the transfer with the help of the algorithm used in gzip, whereby the actual number of bytes to be transferred substantially decreases. The compression will be available with Basis Version 6.20, which will be first used with BW Version 3.0B. However, in order to currently be able to estimate the effectively transferred amount of data, an HTML file was created in our attempt and was compressed using gzip. The archive file produced by gzip had a similar size to the data to be transferred after the compression.

In the following table, three different data volumes are listed next to the number of http requests. The fourth column contains the number of physically transferred bytes, the fifth column describes the size of the generated HTML file, and the sixth column contains the size of the packed HTML file after compression through gzip.

The http requests can be transferred into a number of protocol round trips according to the following considerations: The first request contains the address of the Web site to be displayed. If an authentication is necessary, then a second round trip is required for this. All additional requests are processed in several, parallel data streams, so that the total number of round trips can be estimated according to the following empirical formula: $\text{Roundtrips} = (\text{http Requests} - 2) / n$, where the number n denotes the degree of parallelism. In our experiments, a degree of parallelism of about four was measured. This standard value is displayed in the table below.

Number of Cells	http - Requests	Round Trips	Transfer [Bytes]	HTML File [Bytes]	Packed HTML File [Bytes]
0 (First Execution)	23	5	156,902	15,671	3,897
0	1	1	16,448	15,671	3,897
1,500	1	1	352,418	351,770	17,942
10,000	1	1	2,274,678	2,273,899	94,034

Table 4: Key figures of Web Application



The above table shows that no noteworthy delays are expected for latencies when using the Web Application after the authentication phase. The duration of the authentication phase strongly depends on the use of formatting information as well as the graphical elements contained in the query. A navigation step within a Web Application query only produces a single protocol round trip. Moreover, the importance of a functioning compression of the HTML stream is noticeable: the larger the amount of transferred data is, the stronger the compression of the data amount.

3. Example

Using a query which returns approximately 1500 cells (about 210 rows with 7 cells per row), typical end-to-end runtimes for the various WAN technologies have been measured, meaning the time between starting the query and the display of its result. For the Web Application, however, we could only measure runtimes based on the transfer of unpacked data. In order to simulate the packed transfer, we used a query which returns approximately the amount of data corresponding to the size of the packed HTML file of the original query. The table below compares the runtimes for our query using the BEx Analyzer and the Web Application. The Web Application did not transfer any graphical elements, since these were already available in the browser cache. Column 3 contains the runtime of the Web Application query based on an unpacked stream of data. Column 5 displays the runtimes of the query simulating the transfer of packed data. The factors mean the speed-up of the Web Application with regard to the BEx Analyzer runtime.

Technology	BEx Analyzer [min:sec]	Web Appl. uncompr. [min:sec]	Factor	Web Appl. compressed [min:sec]	Factor
100 MBit LAN	00:05.8	00:03.0	1.9	00:01.3	4.3
56K Modem	00:34.8	00:45.2	0.8	00:08.8	4.0
Terrestrial WAN	00:25.8	00:04.3	6.1	00:03.0	8.3
Satellite Conn.	01:08.9	00:04.8	14.3	00:03.6	19.2

Table 5: Comparison of runtimes BEx Analyzer vs. Web Application (without graphics)

Looking at the Web Application, it is noticeable that the response time on a modem line is extremely high, due to the low band width of the network. Increasing the band width from 56Kbit/s (modem) to 2Mbit/s (satellite connection) results in a reduction of the response times by a factor 10. As expected, the advantages of the Web Application become more noticeable in terrestrial or satellite based WANs.

Likewise, we have measured response times of the Web application assuming that all graphical elements needed to be transferred, thus simulating an initial execution of the query with an empty browser cache. In total we experienced an additional 15 protocol round trips. The graphics transferred in the query consisted of the standard function buttons of the Web Application.

Technology	BEx Analyzer [min:sec]	Web Appl. uncompr. [min:sec]	Factor	Web Appl. compressed [min:sec]	Factor
100 MBit LAN	00:05.8	00:12.2	0.5	00:11.6	0.5
56K Modem	00:34.8	00:53.9	0.6	00:49.4	0.7
Terrestrial WAN	00:25.8	00:15.8	1.6	00:10.4	2.5
Satellite Conn.	01:08.9	00:36.8	1.9	00:22.5	3.1

Table 6: Comparison of runtimes BEx Analyzer vs. Web Application (with graphics)

A comparison of the transfer times with and without graphics shows the massive impact of those graphics on the overall response time. Basically, each graphical element requires one additional round trip. However, graphical elements are transferred after the data, so that the data is displayed before the graphical elements. Taking into account the transfer of graphical elements, a Web Application still shows faster response times mainly in WAN environments, although its performance comes closer to the one of the BEx Analyzer.

4. Recommendations

The measurements demonstrate that the Web Application in the WAN environment has clear advantages compared with the BEx Analyzer. Depending on the network technology, a Web Application displays query results 4 to 20 times faster than the BEx Analyzer. Therefore, we recommend to use Web Applications for accessing data in SAP BW through a wide area network. Due to its low band width, a modem connection is the slowest access method in any case.

The number of protocol round trips is essential for the transfer time in a WAN environment. The Web Application only requires one single round trip for the complete data transfer, independent of the number of data cells. The BEx Analyzer, however, uses a separate round trip for each 200 cells, and additional 37 round trips for fix costs.

When transferring graphical elements, it is important to notice that each graphical object (e.g. function buttons, graphical representations of data, etc.) requires its own round trip. Graphical elements are cached at the front-end and will not be transferred a second time when reused (e.g. function buttons in standards Web Applications). In general, the number of graphical elements has significant impact on the query runtime.

If a query contains graphical representations of data (e.g. a pie chart), the graphics are generated by the application server in case of the Web Application and transferred to the front-end as GIF image, using one round trip. If the BEx Analyzer is used, however, the application server needs to transfer the data, which requires several round trips, depending on the number of cells. The graphical representation is then generated by Excel on the front-end.

From the above-described measurement results, the following empirical formulas for the expected network traffic for queries is derived. In doing so, the registration of the BEx Analyzer in BW was not included in the calculation. The values are solely regarded as rough clues, since especially in cases involving the Web Application, the data volume is so strongly dependent on the distribution of transferred data and, along with it, the compressibility.

BEx Analyzer

	Fixed Costs	Costs per 1000 Cells
Network Load per Report:	34 KB	7 KB
Round Trips per Report	37	5

Typical Query (1500 Cells, Without Registration in BW): **45 KB / 41 Round Trips**
approx. 6 sec (LAN) / 69 sec. (WAN)

Web Application

	Fixed Costs	Costs per 1000 Cells
Network Load per Report		
Non-Compressed:	16 KB	230 KB
Compressed:	4 KB	10 KB
Round Trips per Report (Acceptance: Screens and Other Static Share Already in the Browser Cache)	1	----

Typical Query (1500 Cells, compressed, without registration): **19 KB / 1 Round Trip**
approx. 1.3 sec (LAN) / 3.6 sec. (WAN)