Executive Report

WinterCorp

www.wintercorp.com

SPONSORED RESEARCH PROGRAM

MEASURING PERFORMANCE FOR ENTERPRISE BUSINESS INTELLIGENCE

SAP BusinessObjects XI on IBM Power Systems with Thousands of Concurrent Active Users

THE DRIVE FOR ENTERPRISE BUSINESS INTELLIGENCE

ONCE AN ESOTERIC TOOL used mainly by a few back room specialists, business intelligence (BI) has become a valued tool for employees in nearly every job in the modern enterprise. BI is the accurate, timely and interactive access to the most important information in the company—typically encompassing integrated, correct information on customers, products, prices, inventories, suppliers and many other subjects.

Executives are finding that putting the right facts into the hands of each person in the extended enterprise is a powerful way to enhance service, customer satisfaction and business performance. Employees can make better decisions. Suppliers with access to order and inventory data can do a better job of providing critical items on time. Partners with access to the most up to date pricing and promotions can provide a more effective distribution channel. And customers with ready access to their own order history can more easily serve themselves—and, with higher satisfaction.

Figure 1: Trend Toward Pervasive Use of Business Intelligence

As a result, a growing number of companies are looking for a way to make business intelligence available to thousands, tens of thousands or even hundreds of thousands of users. Executives in these companies are wondering: as I roll out my enterprise business intelligence strategy, will my platform deliver? Just how many online users can I service? Will my BI resources scale to suppliers, partners and customers if need be?

LARGE USER COMMUNITIES

What is required to provide a community of 100,000 or more users with access to business intelligence data?

First, we need a precise vocabulary to analyze this question. The total population of 100,000 users is referred to as the set of "named users."

ABOUT WINTERCORP

WinterCorp is an independent consulting firm that specializes in the performance and scalability of terabyte- and petabyte-scale data management systems throughout their lifecycle.

Since our inception in 1992, we have architected many of the world's largest and most challenging databases in production today. Our consulting services help organizations define business-critical database solutions, select their platforms, engineer their implementations, and manage their growth to optimize business value.

With decades of experience in large-scale database implementations and in-depth knowledge of database products, we deliver unmatched insight into the issues that impede performance and the technologies that enable success.

411 WAVERLEY OAKS ROAD, SUITE 328
WALTHAM, MA 02452
781-642-0300

visit us at
www.wintercorp.com

While usage patterns vary, a general guideline is this: in a large user community, one tenth of the population will actually be connected to the business intelligence system at any one time: we refer to these as the "online users." The rest will be doing some other type of work at that moment: perhaps they will be in a meeting, on the phone, doing something on the factory floor, etc.

Further, of those online users, at any one moment one tenth will actually be interacting with the system—because they have submitted a report, have submitted a query, or have requested some other system service. These are called the "concurrent active" users. The rest will be connected but not otherwise consuming resources because they are studying data that has been displayed for them; thinking about what information to request next; or, perhaps talking to a colleague who has walked into the office.

So, with a total population of 100,000 named users, a general guideline is that one tenth will be online and one tenth of those – or 1,000—will be actually be interacting with the system at any given time, as shown in *Figure* 2.

Figure 2: General Guideline

SAFETY MARGIN

In answering the scalability question, "will my BI platform deliver?", WinterCorp recommends that planners build in a 5x margin of safety on top of the general guideline: so for a total population of 100,000 named users, plan for 5,000 concurrent active users. That way, you have some capacity in reserve to deal with demand peaks: those moments when an unusually large number of people have requested services from the system.

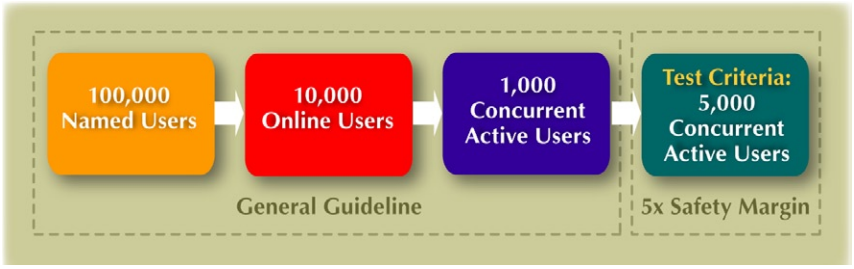

Figure 3: Test Criteria Employs a Five Times Safety Margin

If beyond the planning stage, you should then undertake a requirements study and, if practical, refine your criteria. It is important for planners to recognize that no single rule of thumb is applicable to all deployments.

SAP BUSINESSOBJECTS XI

The business intelligence platform SAP BusinessObjects XI 3.1, introduced in October 2008, is a unified capability with new levels of function and integration designed to support enterprise deployments. Users can access a common view of data through a number of different BI solutions, including SAP BusinessObjects Web Intelligence, SAP BusinessObjects Desktop Intelligence and Crystal Reports.

The common view provides access to a wide variety of data sources, including IBM's DB2, and a wide range of information types.

THE SAP BUSINESSOBJECTS XI / IBM PERFORMANCE STUDY

Realizing that a steadily growing set of enterprises is now planning for, or implementing, business intelligence for large user populations, SAP and IBM recently conducted a performance study. This exercise was designed to measure the ability of the SAP BusinessObjects XI BI platform to service a named user population of 100,000 or more when running on top of IBM DB2 on IBM servers.

TEST CRITERIA

Applying the guideline above with a 5x safety margin built in, the goal of the test was to service 5,000 concurrent active users while delivering a response time of 10 seconds or less at least 90% of the time. The tests were conducted over the last quarter of 2008 at the IBM Benchmark Center in Montpellier, France.

TEST DESIGN PRINCIPLES

In a large population of business intelligence users, it is critical to recognize that needs will vary. A small percentage of users—say, statisticians or financial analysts—will engage in intensive analyses that will run for hours. A much larger group—say, managers or sales associates—will use the system occasionally, just glancing over static reports that provide them such information as the latest customer orders, the latest promotions or their progress toward business objectives.

A key goal in the test design was to create a realistic workload: one in which multiple active users were performing different tasks at the same time; in which the system requests generated by those tasks made significant demands on the resources of the system; and, in which the tasks reasonably modeled real user behavior.

A conscious choice was made to avoid certain simplifications which have sometimes been used in other tests in the industry. As an example, in most real business operations, there are different levels of usage depending on a user's role and skill. Simply testing the ability to pull simple reports would not reflect typical real-world use. As such, the usage patterns incorporated in these tests include that of the "power content consumer," someone who employs sorts, filters and other forms of data manipulation and analysis.

TEST CONFIGURATION

The test configuration was as shown in *Figure 4*. The database management system was DB2 Version 9.5, running on an IBM stack consisting of IBM Power System 595 Servers under AIX, an IBM DS 8300 storage array, and WebSphere Application Server. IBM 595 servers, were employed to perform the tests. An IBM Power 550 was used for administration and monitoring. IBM WebSphere was the application server. The allocation of cores to software components varied with the test and is discussed below.

The injection platform (not shown), used to emulate the behavior of the end users and submit requests to the SAP BusinessObjects platform, consisted of 14 IBM blade servers, each emulating 300 to 600 concurrent active users.

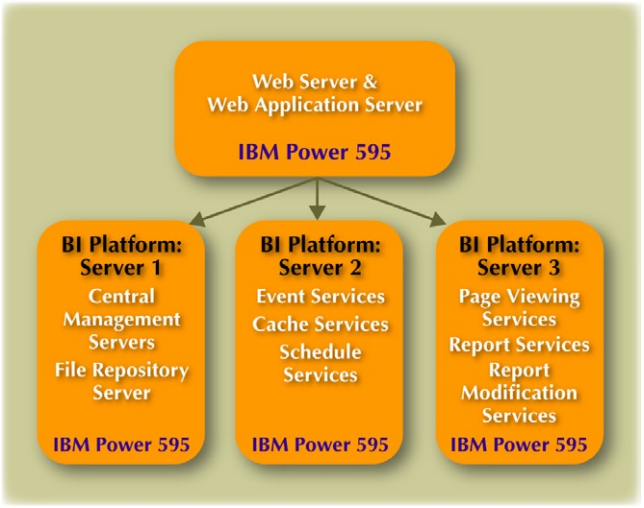

Figure 4: Test Configuration

WORKLOAD

Ten different patterns of end user behavior—referred to as "workflows"—are modeled in the test workload. The patterns include usage of the Crystal Reports (50%), SAP BusinessObjects Web Intelligence (30%) and SAP BusinessObjects Desktop Intelligence (20%). The requests involve a range of queries, from simple to complex. They retrieve as much as 75,000 rows of data.

The patterns also vary with respect to what the user does upon retrieving the data. In simpler patterns, the user pages through the data retrieved, examining each screen of data for a while before going to the next. In the complex patterns, users sort, filter and otherwise manipulate the retrieved data to simulate a more analytical pattern of use.

Those patterns defined as complex comprise 17% of the users in all tests. The presence of these more complex

patterns of usage, in which the data is manipulated and analyzed after retrieval, is important to the realism of the tests, as it models the real life behavior of users who stress the business intelligence platform more heavily. The percentage of users exhibiting each behavior pattern is held constant in the test as the number of concurrent active users is incrementally increased from 1,000 to 5,000 users. Recall that "concurrent active users" is the number of users actually requesting services from the system at one time.

In the judgment of WinterCorp, the tests were designed primarily to focus on the ability of the configuration to handle thousands of concurrent active users performing typical business intelligence reports and queries. They are not testing the scalability of the relational database management system (DB2) with respect to such issues as very large data volumes, highly complex data warehouse queries or high volume data loading.

CORE ALLOCATION

The bar charts and the table in *Figure 5* show the allocation in cores for each of the four tests, and the breakdown of that allocation for each of the four major software components.

Thus, in *Figure 5*, the total number of cores allocated increased from 16 in the 1,000 user test to 64 in the 5,000 user test. In addition, in the 5,000 user test 28 cores were allocated to SAP BusinessObjects Enterprise (BOE); 22 cores were allocated to IBM WebSphere; and, 14 cores

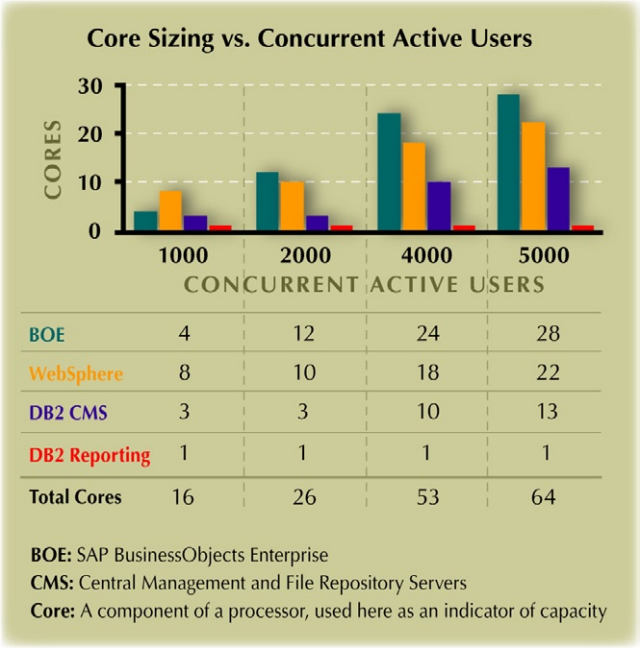

	1000	2000	4000	5000
BOE	4	12	24	28
WebSphere	8	10	18	22
DB2 CMS	3	3	10	13
DB2 Reporting	1	1	1	1
Total Cores	16	26	53	64

BOE: SAP BusinessObjects Enterprise
CMS: Central Management and File Repository Servers
Core: A component of a processor, used here as an indicator of capacity

Figure 5: Processor Allocation, in Cores, from 1,000 to 5,000 Concurrent Active Users

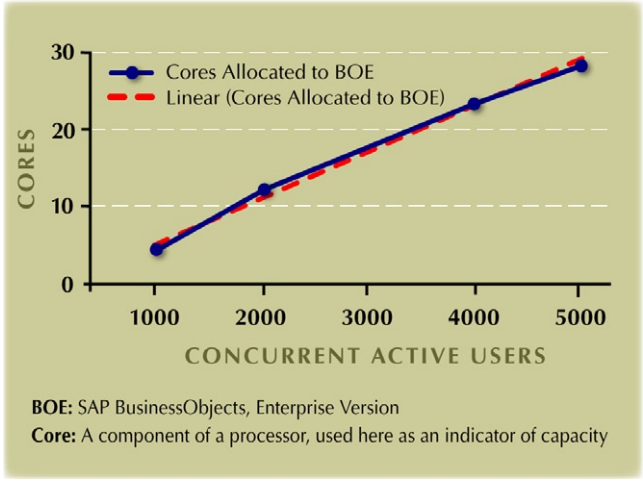

Figure 6: Trend Line of Processor Capacity vs. Concurrent Active Users for SAP BusinessObjects Enterprise (BOE)

were allocated to IBM DB2: one for reporting and 13 for Central Management and the File Repository.

Figure 6 shows the number of cores allocated to SAP BusinessObjects Enterprise (BOE) for each of the four tests vs. the number of concurrent active users in the test, plotted on a graph. These data points are the same as the first line in the table in *Figure 5*. They describe the capacity made available for running SAP BusinessObjects Enterprise itself, separate from any capacity requirements for WebSphere or DB2. It is worthwhile to look at this element separately, since it is both the largest and the fastest growing component of capacity in *Figure 5*.

The blue dots are the actual data points. The thinner, blue line represents the linear function that most closely fits the data. ***Figure 6* indicates that BOE capacity requirements are roughly linear over the span tested, from 1,000 to 5,000 concurrent active users, looking slightly better than linear near the top of the range.**

TEST RESULTS

The tests showed that the SAP BusinessObjects XI platform running on the IBM configuration readily scaled from 1,000 to 5,000 concurrent active users while exceeding the response time target by a significant factor.

Recall that the target was to obtain ten second response or better to at least 90% of system requests. **In fact, the actual average response time was far better, at 2.01 seconds, with 5,000 concurrent active users.**

Figure 7 shows the actual processor utilization for the entire BI platform during the tests, as the number of concurrent active users was increased from 1,000 to 5,000. *Figure 7* thus indicates that the BI platform as a

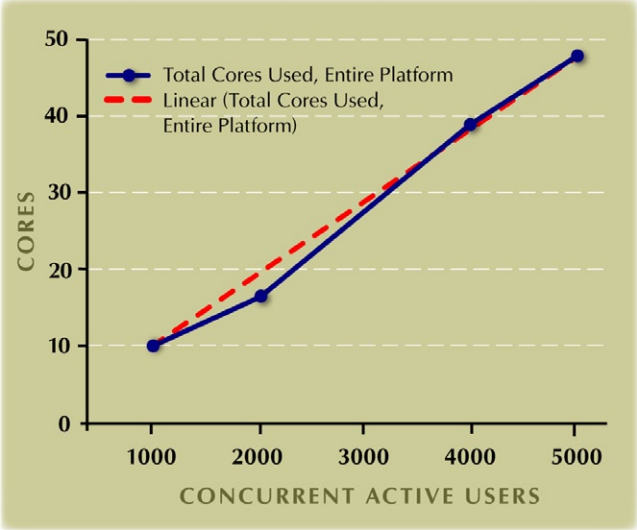

Figure 7: Trend Line of Processor Utilization vs. Concurrent Active Users for the Entire BI Platform

whole scaled in a linear fashion over this range, with respect to the most important measure of its resource consumption: processor cycles. We use the term "entire BI platform" to describe the sum of all the components depicted in *Figure 4*.

Finally, though not part of the formal test protocol, an additional test was run to explore the ability of the test system to scale beyond 5,000 concurrent active users. This additional test demonstrated that the same configuration was actually able to support 7,000 concurrent active users with similarly good response time, as shown in *Figure 8*—which charts both processor consumption in cores and throughput in workflows per second against the number of concurrent active users. Once again, this chart indicates roughly linear scaling over the range tested.

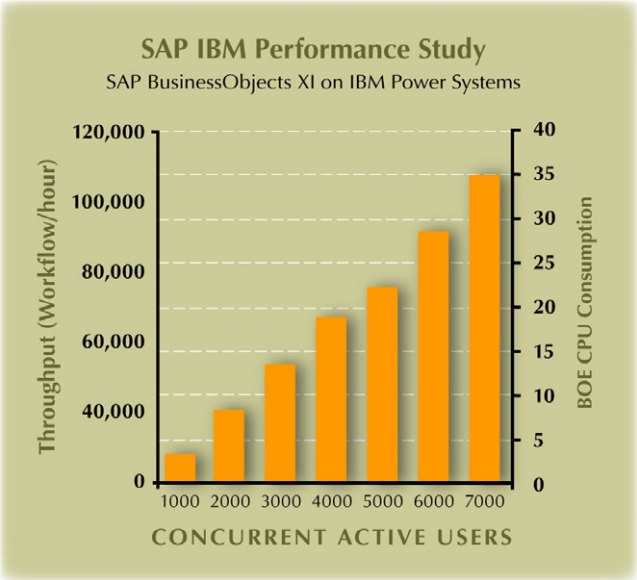

Figure 8: Throughput and Processor Consumption vs. Concurrent Active Users

CONCLUSIONS

In a realistic test of the SAP BusinessObjects XI 3.1 platform, running on an IBM configuration—of database system, operating system and hardware – performance and user scalability was demonstrated consistent with the requirements of an enterprise deployment of business intelligence query and reporting. **Tests involving 5,000 concurrent active users were successful, with an excellent average response time of 2.01 second. These tests are believed to be consistent with supporting a total user population in the range of 100,000 named users. Further, as the concurrent active user population increased from one thousand to five thousand, processor consumption was at or near linear.**

These test results have important implications for customers considering or implementing business intelligence for large user populations. For many companies, this test will provide a basis for moving forward with SAP BusinessObjects XI for deployments in the range of 100,000 named users. The test shows that, for realistically complex patterns of usage, such large populations can be supported by SAP BusinessObjects XI with good performance. Since today's business intelligence implementations often include a broad network of users beyond employees—customers, suppliers, merchandisers, manufacturers, others—such enterprise-scale capacity is an important consideration for larger organizations.

The linearity of the test results also provides an indicator of scalability: that incremental increases in user population will typically result in no more than proportional increases in required system resources, over the range for which the tests were conducted.

Organizations contemplating large scale business intelligence deployments should, whenever practical, conduct their own requirement studies to establish likely usage patterns; likely peak demand levels; and, other parameters that bear on system capacity and performance needs. **WinterCorp recommends that all such users give serious consideration to SAP BusinessObjects XI and IBM on the basis of their demonstrated ability to scale and perform for larger business intelligence user populations.**