White Paper
March 2009

Testing the 'Typical BI Day'

Scalability testing of an IBM Cognos 8 BI enterprise deployment

Conducted at the IBM Center for Business Partners

Information Management

Cognos. software

Contents

Abstract

Comprehensive testing in a range of business intelligence activities shows that IBM Cognos® 8 Business Intelligence delivers enterprise level performance—and scalability—to meet the complex needs of global organizations. IBM Cognos 8 Business Intelligence is the first simple, complete and proven platform to address all enterprise business intelligence needs on a single architecture.

Overview

This paper describes scalability testing performed on IBM Cognos 8 Business Intelligence—the first simple, complete, and proven enterprise business intelligence (BI) platform to address all the needs of BI in one product on a single architecture.

The testing was conducted using a range of typical BI activities based on deployment scenarios suggested by customers, analysts, and IBM Cognos professional services personnel. The scenarios were designed to simulate a real world, enterprise BI environment in which users conduct a variety of role-based activities across the full spectrum of reporting and analysis.

The test results show that IBM Cognos 8 BI can deliver enterprise-level performance and scalability and can meet the diverse and complex BI needs of global organizations. Specifically, the results demonstrate:

- **IBM Cognos 8 BI delivers proven, predictable scalability** – IBM Cognos 8 BI scales linearly with increasing user scalability. Results demonstrated better than linear results for horizontal scalability, and proportional scalability. (For definitions, see "About the Testing.")
- **IBM Cognos 8 BI delivers mission-critical performance and scalability for large numbers of users** – Testing was done using the test environment equivalent of 160,000 users performing tasks varying from simple report delivery to complex analysis. This is based on the 10% standard concurrency ratio assumption.

- **IBM Cognos 8 BI delivers exceptional scalability on real-life, enterprise BI deployments** – IBM Cognos 8 BI achieved better-than-linear scalability with the addition of hardware. When hardware was doubled, in a linear world, response time would increase by 50%. In this case, it increased by 65%. This provides confidence that, as user communities grow and infrastructure demands increase, effective strategies can be planned and managed to meet those new demands.

IBM performs scalability tests at the IBM Innovation Center for Business Partners for System x™ and System p™ IBM environments. This testing was performed in a System x environment, including System x servers, IBM DB2®, and IBM WebSphere®. (For environment specifications, see Figure 1: IBM System x Test Environment.) The worldwide IBM Innovation Centers provide IBM Business Partners with easy access to in-depth IBM skills and cross-platform test environments to help design, implement, port, and test applications for an on-demand world.

Business problems

Scalability – Essential for the success of enterprise BI

As customers expand and standardize BI solutions across their organizations, expectations for BI applications increase, especially in the areas of performance and scalability. The testing of IBM Cognos 8 BI described in this paper addresses the key challenges facing customers as they evaluate vendor solutions for deployment in their environments:

- Will the solution scale to meet my needs now and in the future as my deployment grows?
- Will the solution scale for all of the diverse uses and applications across my organization?
- Will the solution scale predictably, so that I can forecast and budget rollouts with no surprises?

> *"When deploying the Cognos 8 BI solution, performance was a key consideration. As we tuned our BI solution, we realized that performance impacts were most often solved with adjustments to database indices, local network configuration, or operating system settings, and were not specifically unique to the Cognos 8 platform. Cognos 8 met our performance expectations for an enterprise wide solution."*
>
> *Associate Director, Business Intelligence Competency Center, large pharmaceutical company*

Too often, vendors claim seemingly impressive numbers for scalability and performance based on tests conducted using scenarios that are too simple and contrived to resemble a real-life enterprise BI environment. Customers gain little value from these claims, and find the test results difficult to apply to their environments and their desired BI deployments.

Unlike those tests, the testing described in this paper is based on business scenarios that reflect real-life customer deployment environments. In addition to considering business needs, IBM took into account important environment considerations that affect the scalability of BI deployments, such as including security as part of the test environment. The tests were carefully designed to provide a basis for customers to evaluate the ability of IBM Cognos 8 BI to scale to address their specific needs.

Factors affecting the scalability of BI deployments

Scalability testing is challenging because of the numerous factors that can influence customer environments. It is difficult to determine a single baseline for testing a solution that can be used to compare results with other vendor solutions or to provide a norm for guidance during deployment.

Also, because the IBM Cognos solution provides a single platform that delivers all BI capabilities within the customer's infrastructure – with any process, in any environment, on any operating system – the testing possibilities are endless, and the uniqueness of the single IBM Cognos platform versus other vendors' multiple disparate tools makes comparison difficult.

Business drivers

The following outlines key considerations to take into account when performing a scalability assessment of a customer's IT environment.

User community and business cycle

These factors consider the characteristics of the user environment. For example:

- How many users access the solution? User community size impacts the potential load on the system at any time.

- Where are the users located? Are they widely dispersed or located in the same location on the same network? A widely dispersed user group impacts the network load.

- Does the business cycle impact the number of users on the system at a given time, such as at month-end or year-end? Business cycle considerations drive peak period usage that needs to be accommodated.

BI solution requirements

These factors focus on how the application affects the performance results. For example:

- Are users generating simple reports or performing complex comparative analysis? More complex queries may require more time to process, depending on data volumes, database indexes, and query optimization.

- Are the requests accessing large data sets? Large data sets may require more processing time, depending on the type of request.

- How often are the requests being made? Special considerations, such as dedicated servers, may be made for those requests that run most often.

IT environment

The complexity of the IT environment affects the ability of the BI solution to scale. The IBM Cognos solution operates within an existing IT environment, and an understanding of the IT infrastructure is necessary when addressing scalability requirements. For example:

- What is the preferred operating system platform? Different platforms may have different optimal configurations upon installation that can affect performance.
- What is the network capacity? Are there windows of time available to schedule more complex requests? Heavy network traffic by other applications may impact BI performance.
- What are the security constraints? Is data encryption required? Are there security parameters that impact the execution of the query? Time for data encryption and decryption needs to be factored into secure deployments.

The solution

Goal of the testing

The testing was designed to achieve specific goals using scenarios that accurately simulate a real-world enterprise environment and that test for scalability as defined by industry best practices.

The goal of the testing was to confirm that the IBM Cognos 8 BI platform is scalable across the three key principles of scalability – user, horizontal, and proportional scalability as defined in the "Definition of Scalability" sidebar. By proving scalability, the testing validates that IBM Cognos 8 BI can address the needs of ever-growing enterprise BI deployments.

To narrow the performance test options, we took a customer-centric approach, applying BI as it is used in a typical organization. The intent was not to test all of the permutations and combinations in hardware and users, nor to develop a comparison between IBM solutions and competitors' solutions. Available competitor test results focus on raw throughput testing of report execution, which is only one aspect of reporting and one capability among the many BI capabilities provided by IBM Cognos 8 BI. Rather than limit ourselves to simple report execution for the purposes of comparison with other products, we focused on providing customers with an understanding of how IBM Cognos 8 BI can perform and scale to meet their needs, in their IT environments, for their user communities, in their enterprise deployments.

Scenario and use cases

In designing the tests, IBM worked with customers, analysts, and professional services personnel to distill the range of typical BI activities into a set of use cases that represent real customer situations. The test scenario represents a "Typical BI Day," simulating multiple users performing activities such as viewing reports, generating reports, and performing analysis. The "Typical BI Day" scenario was developed based on research identifying the typical distribution of BI users within an organization. To accurately reflect market input on typical BI deployments and uses, and to represent the broadest range of typical enterprise needs, IBM defined the following use cases within the "Typical BI Day" scenario:

- **Scale for massive end-user reach** – As BI deployments expand to embrace a broader, more casual end-user community across the organization, BI reaches users that do not necessarily live and interact with BI everyday, yet who need mission critical information at their fingertips. These use cases focused on performance to deliver the viewing of pre-generated content and robust and predictable scalability to enable growing user communities.
- **Access simple reporting on current information** – As needs grow for operational insight and access, users need simple access to reports with the most up-to-date information. The focus of these use cases was on-demand report execution against a mix of data sources, including dimensionally modeled relational data and OLAP data.

Scalability defined

IBM defines scalability following industry best practices for CPU-bound processing that recognizes three aspects of scalability:

User scalability: This involves setting initial performance expectations when hardware is constant and workload increases. For example, for a constant configuration of eight CPUs, as you increase workload by adding users, reports, or units of work, the performance of a scalable solution changes proportionally (e.g., a performance graph shows a linearly increasing trend until the system is saturated).

Horizontal scalability: This involves meeting performance expectations by adding physical resources. For example, with a constant workload provided by 100 concurrent users, as you increase capacity from eight to 16 CPUs, a scalable solution shows proportional decreases in response times and proportional increases in throughput (e.g., with 8 CPUs and 100 concurrent users the response time is 10 seconds, and with 16 CPUs and 100 concurrent users the response time is approximately 5 seconds).

Proportional scalability: This involves maintaining performance expectations as a user community grows. For example, given a constant configuration/workload ratio of 10 concurrent users per CPU, with a scalable solution you can expect to maintain performance by scaling the hardware (e.g., if there are 500 concurrent users on four CPUs with a user base that is expected to grow to 1,000 users, performance can be maintained by doubling the hardware).

- **Address complex and challenging reporting** – Many business needs require interactive reporting, combining multiple sets of data in a single view and providing users with the interactivity to explore through drilling, ranking, and enhancing the content. These use cases focused on interactivity of prompted reports, drilling, and ranking the content and complexity of multi-query reports, WYSIWYG layouts, and master-detail formats.

- **Address the exploration and research needs of information experts** – Most analysts aren't satisfied with preformatted reporting, but require in-depth analysis and the ability to explore and traverse data to derive insight. These use cases focused on exploration tasks in analysis of drilling, rotating, comparative analysis, and addressing the needs of the most demanding analyst community.

The testing approach

To ensure that the testing focused on enterprise needs, we took a three-step approach:

1. Develop the "Typical BI Day" scenario to accurately depict how BI is used within an organization.

2. Identify scalability test cases for the "Typical BI Day."

3. Apply the three scalability principles – user, horizontal, and proportional – to the "Typical BI Day."

1. Develop the "Typical BI Day" scenario

The "Typical BI Day" scenario is based on how users leverage a BI solution in order to accomplish a task, whether it is receiving static output, generating information in a formatted report, or analyzing information on demand.

The scenario is comprised of four business use cases, reflecting the broad range of Enterprise BI needs as previously outlined in "Scenario and Use Cases." These use cases are distributed within the scenario based on industry research[1] and on our own experience of how BI is typically leveraged within an organization.

The business use case distribution is as follows:

- Scale for massive end user reach – 40%
- Access simple reporting on current information – 30%
- Address complex and challenging reporting – 25 %
- Address the exploration and research needs of information experts – 5%

2. Identify scalability test cases for the "Typical BI Day"

Once the "Typical BI Day" scenario was developed, specific scalability test cases were identified. Similar to the variety of processing seen in real-life deployments, multiple test cases were included in each business use case.

Table 1 outlines examples of the test cases that were applied to the "Typical BI Day" scenario. For more details about the example test cases, see Appendix B: Test Case Descriptions.

Table 1. Typical BI Day test cases

Typical BI Day	Solution requirements	Sample test cases
Scale for massive end-user reach: The BI user receives three pre-generated reports, in both PDF and HTML format	Users view generated reports in a variety of formats	Report Viewer (HTML)
Access simple reporting on current information: The BI user generates a report containing the most up-to-date information, possibly from multiple sources	Users run reports that are easy to access with some simple interaction and navigation across different sources	Time Period Report
Address complex and challenging reporting: The BI user needs to view a report, drill through on detail, and rank content	Users require a more interactive experience, such as reports with prompts and filters, in order to assess the same data set in different ways	Multi-fact Master Detail Report
Address the exploration and research needs of information experts: The BI user analyzes data without requiring a preformatted report	Users require the ability to perform data analysis as opposed to viewing a pre-generated report	Analysis

3. Apply scalability principles to the "Typical BI Day"

When testing the "Typical BI Day," the three scalability principles, as described in the scalability definitions already presented, were applied in order to fully understand the impact of scale on an IBM Cognos solution:

a. User scalability – How does the system perform with continuous scale in workload? To address the question of user scalability, testing was performed using a full hardware configuration and continuously adding workload to the system.

Expected result: It is expected that the IBM Cognos 8 BI Platform on System x demonstrates linear scalability and, therefore, performance will decrease proportionally as user workload is added.

b. Horizontal scalability – How does the system perform when there is scaling of hardware capacity? To address the question of horizontal scalability, a consistent number of concurrent users was used while hardware capacity was increased.

Expected result: It is expected that the IBM Cognos 8 BI Platform on System x demonstrates a proportional increase in performance as hardware capacity is added.

c. Proportional scalability – How does the system perform with a proportional increase of concurrent users and hardware capacity?

To address the question of proportional scalability, hardware and concurrent users are added at a proportional rate.

Expected result: It is expected that the IBM Cognos 8 BI platform on System x demonstrates consistent performance as concurrent users and hardware capacity are added.

By analyzing the results of this "Typical BI Day" testing based on the proportional mix of use cases and the three scalability conditions, organizations can gain a better understanding of how the IBM Cognos 8 platform performs as it scales.

User Community terms and sizes

To accurately judge the number of users that can be supported in a real-world environment based on performance in a test environment, it is necessary to distinguish between named, active, and concurrent users.

Named users make up the total population of individuals who can be identified by and potentially use the system. They represent the total user community, and can be active or concurrent at any time. In a real-life BI environment, this is the total number of individuals authorized to use the system. This is the number of most interest to organizations planning a BI implementation, because it tells those organizations how many users they can expect to support in a given environment with the response times reported in a test environment.

Active users are logged on to the system at a given time and can send a processing request at any time. For example, users who are viewing the results returned from a query are active users, although they are not currently stressing the system.

Concurrent users are not only logged on to the system (active), but are sending a request or waiting for a response. They are the only type of user actually stressing the system at any given time. IBM tests using concurrent users who actively stress the system, adding load at all times during the test cycle.

Concurrency ratio: The number of concurrent requests at any moment in time that defines the load on the system. The system should be designed so that at peak loading, the number of simultaneous requests coming in can be satisfied comfortably. A general rule of thumb in estimating workload is that approximately one percent of named users or 10% of active users will equate to the number of concurrent requests the system must manage per second.

Table 2 illustrates the relationship between the number of concurrent users employed in a test environment and the number of users in a real-life environment for which the test results can be reasonably reported. For example, if 250 concurrent users are employed in a test, then the system performance can be reasonably reported to apply to a real-life system with 25,000 named users.

Table 2. Estimating named users based on concurrent users

1 A test involving this many Concurrent users …	2 approximates a real-life environment with this many Active users …	3 and this many Named users
100	1,000	10,000
250	2,500	25,000
500	5,000	50,000

1 Concurrent users are actually stressing the system.
2 Active users are logged on to the system but are not necessarily stressing it.
3 Named users are all the users authorized to use the system – the total number of business end-users in a real-world environment.

Test environment

The testing was conducted at the IBM Innovation Center for Business Partners in Waltham, MA. The worldwide IBM Innovation Centers provide Business Partners with easy access to in-depth IBM skills and cross-platform test environments to help design, implement, port, and test applications for an on-demand world.

Testing was performed using industry-standard load testing software, Mercury LoadRunner.

Figure 1 outlines the IBM System x test environment that was used at the IBM Center for Business Partners.

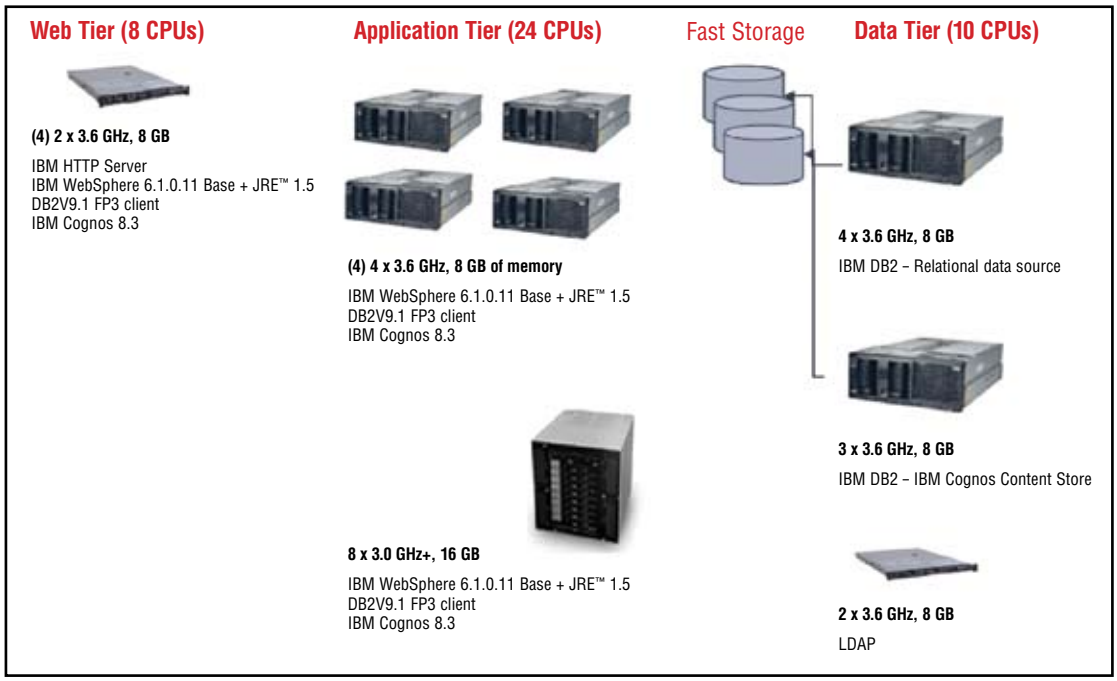

Figure 1: IBM System x test environment

Test results and conclusions

The results described here are based on testing that applied the three scalability principles to the Typical BI Day scenario with its proportional mix of business use cases.

User scalability

Figure 2 shows the total response time trend for the Typical BI Day scenario given an increasing number of concurrent users. Total average transaction time represents the sum of all the averaged transaction times for each of the gestures performed within the tests. A gesture represents a user interaction performed when using IBM Cognos 8 BI such as launching an HTML report.

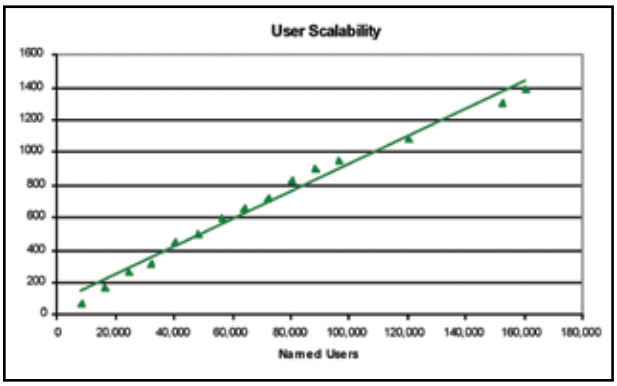

Figure 2: User scalability test results

This test was performed using 16 BI processing CPUs. It is important to note that the optimal configuration will vary, depending on the customer environment.

The results demonstrate that:

- As the number of users increases, there is a proportional increase in response time.
- There is a linear relationship between the expected response time and the number of users.

Conclusion 1: User scalability – Based on the above results, we can conclude that the IBM Cognos 8 platform delivers reliable and predictable user scalability.

Horizontal scalability

Figure 3 shows the response time trend given a constant number of concurrent users with a gradual increase in hardware capacity. In this case, testing was performed with 240 concurrent users.

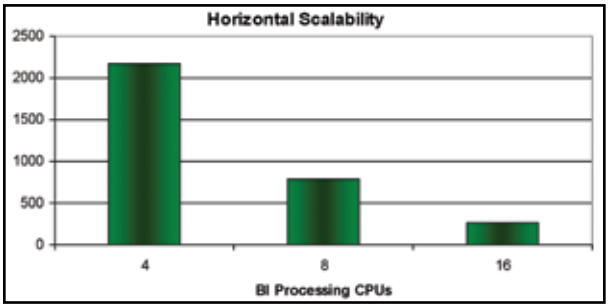

Figure 3: Horizontal scalability test results

The results demonstrate that:

- There is a better-than-linear improvement in response time as hardware capacity is increased.
- Response times improved by 65%, when hardware was doubled. In a perfectly linear world, response would improve by 50% as hardware is doubled.

Typically, with enterprise software, each addition of hardware is used less efficiently than the last. These results demonstrate the IBM Cognos 8 platform scales to meet scalability demands while maintaining the efficient use of the hardware.

Conclusion 2: Horizontal scalability – Based on the above results, we can conclude that the IBM Cognos 8 platform delivers exceptional horizontal scalability.

Proportional scalability

Figure 4 shows the response time trend given a proportional increase in both users and hardware capacity.

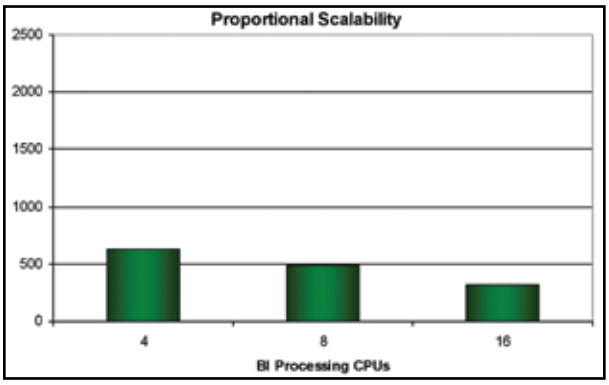

Figure 4: Proportional scalability test results

The results demonstrate that:

- There was improved performance with a proportional increase of both concurrent users and hardware capacity.
- As users and hardware were increased, response times continuously improved at an average rate of 29 percent, demonstrating better than linear results. This means IT can deploy – and grow – the solution as user demands increase.

Conclusion 3: Proportional scalability – From the above results, we can conclude that the IBM Cognos 8 platform delivers exceptional proportional scalability.

Conclusion

Based on the results of the testing, the following conclusions can be made:

IBM Cognos 8 BI delivers proven, predictable scalability
IBM Cognos 8 BI scales linearly with increasing user scalability. Results demonstrated better than linear results for horizontal scalability, and proportional scalability. (For definitions, see "About the Testing.")

IBM Cognos 8 BI delivers mission-critical performance and scalability for large numbers of users
Testing was done using the test environment equivalent of 160,000 users performing tasks varying from simple report delivery to complex analysis. This is based on the 10 percent standard concurrency ratio assumption.

IBM Cognos 8 BI delivers exceptional scalability on real-life, enterprise BI deployments
IBM Cognos 8 BI achieved better-than-linear scalability with the addition of hardware. When hardware was doubled, in a linear world, response time would increase by 50%. In this case, it increased by 65%. This provides confidence that, as user communities grow and infrastructure demands increase, effective strategies can be planned and managed to meet those new demands.

Using these results in your environment

As stated earlier, one of the goals of the testing was for customers to be able to use the results in their own environments, for their own BI deployments.

We have taken a customer-centric approach to our testing to provide results that can be applied to customer environments. The testing results should be considered when scaling the BI solution across an enterprise, because they allow you to:

- **Meet your performance service level agreements (SLAs) –** Predictable performance allows IT to effectively plan deployment hardware requirements to meet the performance SLAs upon implementation.

- **Manage user expectations –** Understanding performance allows the IT department to effectively communicate with users about performance considerations prior to implementation. An effective communication strategy ensures that users' expectations are realistic when a system is deployed.

- **Maximize your IT investment –** A crucial part of maximizing your overall investment is managing the timing of IT infrastructure investment. Last minute purchases to handle unexpected performance results reduce the ability to flexibility manage your IT budget.

Appendix A: Details of the Typical BI Day

The graphs in this appendix depict a sample of the user scalability test results for the business use cases that comprise the Typical BI Day scenario, based on the System x test environment at the IBM labs (see Figure 1). Customer results may vary depending on the specific environment in which a solution is deployed.

In every example shown below, linear scalability was achieved.

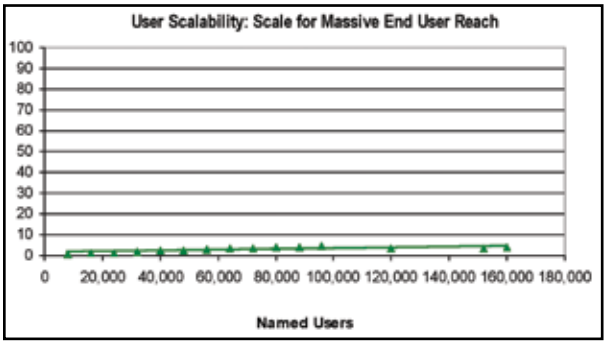

Figure 5: Scaling for massive end-user reach

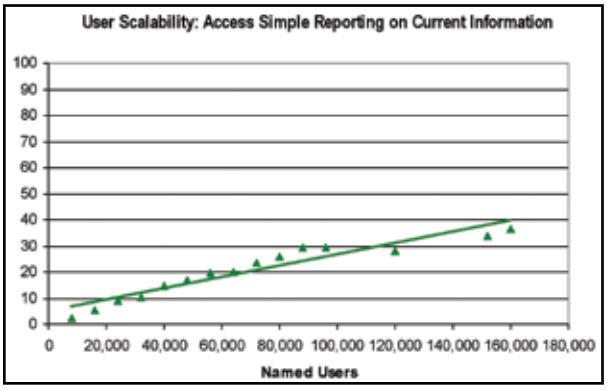

Figure 6: Accessing simple reporting on current information

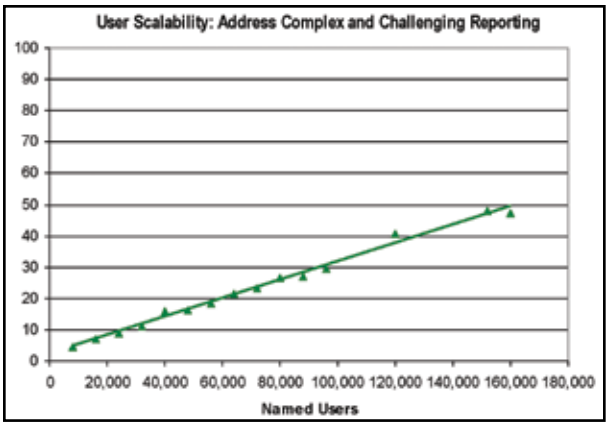

Figure 7: Addressing complex and challenging reporting

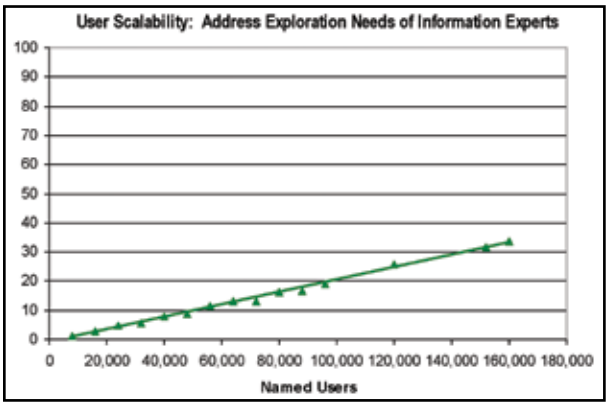

Figure 8: Addressing the exploration and research needs of information experts

Appendix B: Test case descriptions

This appendix provides examples of the reports used as test cases in the Typical BI Day.

Report Viewing (HTML): This report was been pre-generated and contains static data (as at the time of execution).

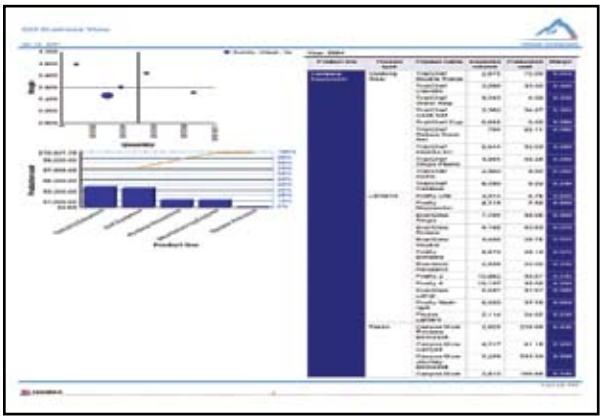

Time Period Report: A report that provides information within the context of a time period. For example: quarterly results, year-over-year comparison

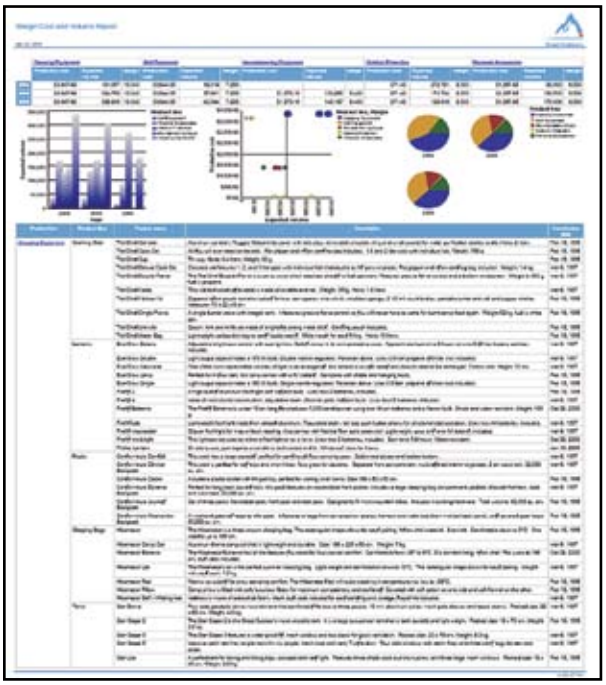

Multi-fact Master Detail: A report that contains both summary level information as well as the associated details. This report combines the use of charts to display summary information and cross-tabs for detailed information.

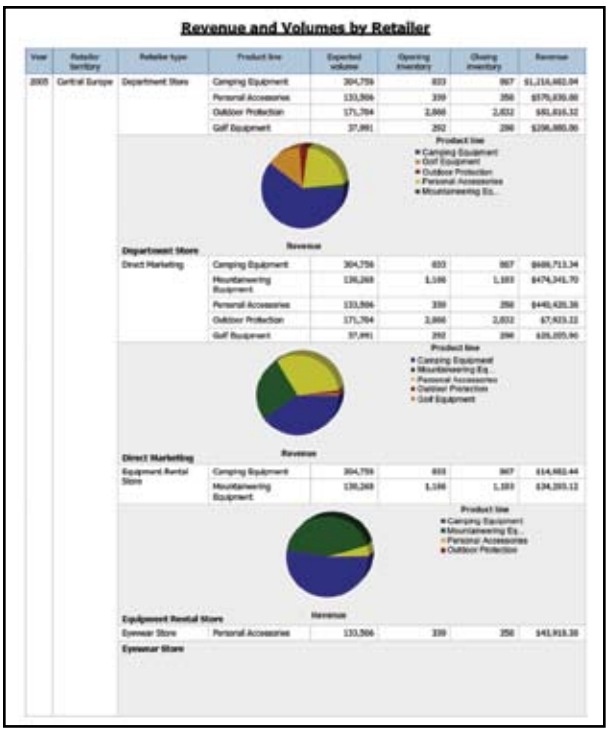

About IBM Cognos BI and Performance Management

IBM Cognos business intelligence (BI) and performance management solutions deliver world-leading enterprise planning, consolidation and BI software, support and services to help companies plan, understand and manage financial and operational performance. IBM Cognos solutions bring together technology, analytical applications, best practices, and a broad network of partners to give customers an open, adaptive and complete performance solution. Over 23,000 customers in more than 135 countries around the world choose IBM Cognos solutions.

For further information or to reach a representative: www.ibm.com/cognos

Request a call

To request a call or to ask a question, go to www.ibm.com/cognos/contactus. An IBM Cognos representative will respond to your enquiry within two business days.

IBM Canada
3755 Riverside Drive
Ottawa, ON, Canada K1G 4K9

Produced in Canada
March 2009

Endnotes

1 Forrester, Profiling the Analytic End-users for Business Intelligence, 2004.

IML14035CAEN