

Scalability

Ensuring Long-Term Value for Business Intelligence and Reporting Investments

A White Paper

Table of Contents

1	Executive Summary
3	What Is Scalability and Why Does It Matter?
3	Historical Precedents
4	Scalability and Standardization
6	Benchmarks and Performance Tests
7	The Truth About Concurrency
7	Developing Reliable Testing Procedures
8	Scalability Testing With WebFOCUS
10	Architectural Considerations
12	The Unique WebFOCUS Architecture
14	WebFOCUS in Action: Large-Scale Deployment Case Studies
16	About Information Builders
17	Appendix A – WebFOCUS Benchmark Test Results
18	Appendix B – Independent Benchmark Comparisons: Cognos Versus Information Builders

Executive Summary

Scalability refers to the ability of a business intelligence or enterprise-reporting environment to leverage new and existing technology resources to support greater numbers of users and increasingly complex activities, ranging from simple page-views to complex analysis exercises. According to research conducted by the Data Warehousing Institute, senior IT executives report that their number one requirement for successfully implementing business intelligence (BI) systems is scalability and availability.¹

Enterprise-wide BI and reporting tools must be able to scale effectively to meet escalating demands from the user base, as well as to minimize hardware and administrative requirements. One key business driver is the escalating growth in the BI market. According to researchers at Gartner Dataquest, the BI market experienced 5.5 percent growth during 2003, and indicators suggest that 2004 will be a year of strong growth in the BI market as well. “As the global economy rebounds, BI is near the top of IT and line-of-business agendas,” Gartner researchers report.²

The fastest market growth is attributed to enterprise business intelligence suites (EBIS), which are designed to support many types of users with ad hoc database query, reporting, and online analytical processing (OLAP) capabilities. For example, Information Builders’ WebFOCUS is a comprehensive and fully integrated enterprise business intelligence suite with a broad range of features to satisfy the diverse needs of an unlimited number of users. As with other leading tools in this market, having a broad range of capabilities has been partially responsible for WebFOCUS’ surge in popularity. However, not all of today’s BI tools will be able to keep up with burgeoning business needs. As demand for BI and reporting solutions increases, organizations are deploying BI apps on an increasingly broader scale, exerting tremendous pressure on the underlying infrastructure.

This white paper offers a comprehensive look at scalability in the BI domain, from historical precedents to current best practices. Scalability is a relative term, so we have included tips on how to use benchmark studies as part of any vendor-selection process. Benchmark testing provides an unbiased evaluation of a product’s capabilities and performance, serving as a useful guide for what you can expect from a given set of BI tools during a variety of performance workloads.

¹ “Getting Real about Real-time BI,” November 2003.

² H. Dresner, B. Hostmann, A. Tiedrich, and F. Buytendijk, “Magic Quadrants for Business Intelligence,” April 19, 2004.

Throughout this paper, we will examine some of the key ways to measure business intelligence scalability, including:

- Total number of users
- Number of concurrent users
- Types of users
- Types and amount of data
- Types and numbers of applications
- Performance over time

Scalability not only involves taking advantage of additional computing, memory, and I/O resources, but using those resources in an efficient, cost-effective way. Most of today's leading BI tools use a server-based architecture to simplify software maintenance, streamline report distribution, and minimize network activity. However, there are many subtle architectural features that are essential to scalability, such as nonpersistence, server multithreading, native data adapters, and automatic report bursting. We'll discuss these and other architectural requirements in the pages that follow.

"User scalability of business intelligence tools and solutions has become paramount," states Gartner analyst Howard Dresner. "Due to various claims by vendors, enterprises must scrutinize offerings to ensure they can deliver on their promises."³

³ H. Dresner, "Scalability and BI: What You Need to Know Now," April 14, 2004.

What Is Scalability and Why Does It Matter?

Gartner defines business intelligence as a user-centered process that includes accessing, exploring, and analyzing data to develop insights and increase understanding, leading to improved decision-making. BI helps business managers to make tactical decisions, as well as to establish, modify, or tune business strategies and processes to gain competitive advantage and achieve corporate goals. Typically this involves an iterative process of accessing and analyzing data – deriving insights, drawing conclusions, and communicating findings to effect positive change within the enterprise.⁴

Well-designed business intelligence and enterprise reporting environments have a way of becoming very popular very fast, which is why many BI installations are characterized by rapid growth. For many organizations, BI is equated to enterprise reporting applications. For these companies, a successful BI deployment is one that can be broadly shared by a large number of users. Once useful information is made available, the demand for it increases exponentially.

Not only does usage proliferate across departments, divisions, and organizational boundaries, but new types of users invariably come on board, making their own unique requests for data and devising new ways to manipulate and analyze that data. Some users are simply interested in running standard reports. Others want to create ad hoc reports or perform statistical analysis. Thus, the ability of the technology environment to scale, grow, adapt, and change is crucial to any BI or reporting initiative.

The corollaries to scalability are threefold: performance, manageability, and reliability. The issue is not merely whether a BI environment can scale to support thousands of users, but whether it can do so while maintaining acceptable levels of reliability and performance – and not pushing system administration costs through the roof.

Historical Precedents

Historically, most BI vendors have provided technology for small groups of highly sophisticated users. Up until about 1996, very few organizations shared information outside their firewalls. The Internet laid the groundwork for extra-enterprise collaboration by standardizing protocols and

⁴ Tiedrich, Alan H., “Business Intelligence Tools: Perspective,” Technology Overview # DPRO-93784, June 19, 2003.

increasing the speed of communication. As a result, today's Web-based BI systems enable organizations to easily communicate with all of their constituents, from employees to customers. The Internet made these broad deployments possible not only by creating standards for collaboration, but also by enabling very light-weight clients consisting of only a Web browser. Today, this "strong server, weak client" model makes large-scale distribution possible.

In practice, today's BI vendors have achieved varying levels of Web maturity. Some products are inhibited by their client/server legacy. Many vendors are only now getting around to creating an efficient, Web-based architecture with proven scalability. We'll discuss this further in the architecture section of this paper.

As a whole, EBIS products have been promoted on a platform of "information democracy," which refers to the spread of BI tools to all types of users, both inside and outside of the enterprise. Gartner believes EBIS suites should be extensible not only to internal users, but also to customers, suppliers, and other extranet users. EBIS products should also support a variety of users by providing extreme ease of use and requiring minimal training.⁵

Extending business intelligence outside of the enterprise doesn't simply mean pushing information in a passive way. Many users want to be able to interact with the information – to obtain it on their terms, when and where they need it. Thus, instead of simply sharing or pushing information out to partners and suppliers, companies are creating self-service applications that give outside users the ability to customize their views of relevant information. These activities place increasingly rigorous demands on the BI infrastructure.

Scalability and Standardization

Until recently, activities such as user reporting and financial analysis were treated in a departmental manner. Each department established its own BI practices, leading to a great variety of reporting products within most companies. Today, large organizations are working to enforce a uniform BI strategy. This consolidation trend is encouraging the adoption of enterprise-level standards for reporting and BI. IT managers at many large companies are

⁵ *ibid.*

insisting on enterprise-class BI tools that can support their growing user populations. Departmental BI applications are no longer sufficient for organizations that need to deploy reliable, secure information to potentially every employee, partner, customer, and prospect.

With BI deployments commonly reaching throughout the enterprise and beyond, it now becomes critical that those applications deliver an acceptable level of performance without forcing organizations to invest in unnecessary server hardware – and the administrative resources that additional servers imply. CIOs want corporate standards to simplify the infrastructure. If a BI tool can't scale, it won't succeed as a standard.

According to researchers at Gartner, most BI vendors boast about their ability to scale to large numbers of users. However, true scalability has been difficult to achieve without expending large sums of money on excessive amounts of equipment and personnel. “BI technology vendors are finally beginning to understand that those vendors with the most-scalable systems stand the greatest chance of winning large enterprise deals,” writes Dresner. “Thus, the newest generation of BI technologies has more-efficient server engines. Through the use of multithreading, intelligent caching, and application servers such as J2EE, many server engines show significant potential as enterprise-class software.”⁶

⁶ Dresner H., “Business Intelligence Technology Directions,” Research Note COM-21-6252, December 5, 2003.

Benchmarks and Performance Tests

A benchmark is an unbiased evaluation of a product's capabilities and performance. For any company planning to purchase an EBIS or deploy an enterprise-reporting environment, these evaluations are essential, since they verify the ability of a given technology to scale under real-world conditions.

On this matter, Dresner is unequivocal. "Before investing in a BI solution, test the software for scalability in a 'real world' environment, using realistic data sets and simulating large numbers of simultaneous users with a mix of workloads (for example, static viewing, casual interactive, or aggressive interactive)," he advises. "Software from Mercury Interactive (LoadRunner) and Compuware (QALoad) can be used to simulate user workloads and give insights into application resource use, which can uncover application inefficiencies. Most testing suite suppliers have labs and will provide application testing services."⁷

When embarking on any benchmarking study, there are a few key questions to ask:

- How do I maximize performance within a given budget?
- How do I minimize costs for a required minimum performance level?
- How do I obtain the best performance/cost ratio within a given budget or given performance requirements?

For example, a major financial services company enlisted help from Information Builders to evaluate a baseline for a new enterprise reporting application. The results of the study were revealing: the application scaled to more than 700 users before memory was depleted on the designated server, and the response times never exceeded one minute throughout the test. While the performance metrics and response times will differ from one configuration to the next, tests such as these help an organization determine whether or not the tools will meet their requirements under maximum loads and stress levels.

Based on experience from this and other tests, Information Builders has devised the following benchmark checklist:

- Use standard tools and applications across all tests
- Include a complete description of the testing environment setup
- Make a chart of all tests and results

⁷ H. Dresner, "Scalability and BI: What You Need to Know Now," April 14, 2004.

- Isolate individual variables
- Verify all results (run your benchmarks to check for any possible variations)

The Truth About Concurrency

When comparing the results of vendor-published benchmark studies, take note of how vendors define their terms, and probe further if necessary. For example, vendors commonly report findings based on *concurrent* users, but different vendors mean different things by concurrency. What are those users doing? Are they merely viewing previously generated HTML pages, or are they submitting report requests that involve queries to the database? Generating reports is a resource-intensive activity that engages the database server, Web server, and application server. Viewing HTML pages is a browser-based activity that involves minimal back-end resources. So if the BI vendor claims it can support hundreds of concurrent users, dig a little deeper: What kinds of users, what kinds of reporting processes, what kinds of concurrent activities?

In this paper, *concurrent* users refers to users who are logged onto the system and actively sending requests or waiting for a response. *Active* users, by contrast, refers to users who are logged on to the system and are simply reviewing the results returned from a previously submitted query.

If you do your own performance testing, you need testing tools that can isolate each variable within the architecture to identify potential bottlenecks. For example, sometimes performance suffers on the back-end because of inefficient database processing. Mercury Interactive's LoadRunner and Compuware's QALoad are both good choices.

Some BI vendors include performance-testing tools that simulate virtual users, with varying levels of sophistication. The WebFOCUS testing tool includes SQL tracers to ensure that each database request is optimized as well.

Developing Reliable Testing Procedures

Reliability is the extent to which an experiment, test, or measuring procedure yields the same results on repeated trials. This comes from several aspects of the procedure:

- The criteria are specific and well defined; each measure is explained and clarified

- The analysis procedure is carried out by carefully trained reviewers
- Each team is providing evidence-based arguments for their judgments

Key questions to ask at the outset of any benchmarking study include the following:

- What testing tools are being used?
- Do these tools support the requisite number of users?
- How many CPUs are supporting the BI environment and what are the processor speeds?
- How much real and virtual memory is available?
- Which Web server and application server are being used, and can you optimize their performance separately?
- Can you test the database without the Web server (pure back-end testing)?

Scalability Testing With WebFOCUS

While most BI vendors boast about their “light-weight,” thin-client software architecture, that’s not good enough for broad external deployments. For example, some vendors use the common gateway interface (CGI) protocol to connect PCs to servers, which requires separate resources for each connected user – often as much as 10MB to 20MB per user. This approach works fine for departmental implementations, and even for some large-scale corporate rollouts. However, when it comes time to deploy a self-service application over the Internet – where user loads can easily number in the tens of thousands – the cost becomes prohibitive.

That’s why WebFOCUS uses a “zero-client” footprint, with no software other than a standard Web browser required on user devices. Not only does this ensure that hardware and administrative costs are kept to a minimum, but it minimizes training costs as well, since users don’t need to contend with a unique client-side software environment.

Ventana Research conducted a six-month study to determine the total cost of ownership (TCO) for scaling various business intelligence environments, as deployed in real-world situations. Information Builders beat all participants in the categories of predefined reporting and parameterized reporting on both Windows and UNIX platforms. The study confirmed that WebFOCUS is both scalable and affordable across a broad range of scenarios, from 1,000 to

10,000 users. Information Builders scored first in three of the 10,000-user scenarios, and second in six of the 10,000-user scenarios.⁸

More recently, Information Builders performed extensive scalability testing of its WebFOCUS enterprise business intelligence application in conjunction with the Unisys ES7000 family of servers.⁹ The tests demonstrated that WebFOCUS can scale to meet the most demanding business intelligence and reporting requirements with exceptional response times – even when tens of thousands of users are involved. Detailed test results are presented in Appendix A, and summarized below.

The UNISYS ES7000 hosted the WebFOCUS software along with an IBM DB2 database containing 6.5 million rows of data. The tests, which measured the activities of both concurrent and active users, were designed to develop reliable, repeatable benchmarks for WebFOCUS. The key goals of the testing were as follows:

- To determine the performance and scalability characteristics of WebFOCUS with an increasing number of users performing typical functions, including report viewing, report execution, batch reporting, and portal navigation
- To test the limits of the number of concurrent users that the system could accommodate
- To develop accurate comparison measurements for competing reporting software and hardware platforms

WebFOCUS, configured on a 32-CPU Unisys ES7000, delivered superior report execution response times – surpassing similar performance benchmark data from Cognos Corporation (see Appendix B).

The same WebFOCUS/Unisys configuration posted average report viewing response times that exceeded the highest previously published thresholds by 83 percent for 900 concurrent users, 70 percent for 1,000 concurrent users, and 64 percent for 1,100 concurrent users.

On a 16-CPU Unisys ES7000, WebFOCUS eclipsed previously published report-execution response time data by 46 percent for 200 concurrent users, 43 percent for 300 concurrent users, and 35 percent for 500 concurrent users.

⁸ “Business Intelligence Total Cost of Ownership Benchmark,” Ventana Research, September 2003.

⁹ “Information Builders’ WebFOCUS and Unisys Enterprise Scalability Benchmarks for the Windows Environment,” Unisys Corporation and Information Builders, April 2004.

Architectural Considerations

There are an endless number of architectural approaches to BI, and each approach could be considered correct if it delivers pertinent information to the business community. However, some architectures, designs, and topologies deliver significantly greater value and more efficiency.

For example, some products require developers to populate local data marts and cubes before users can analyze information. These tools are acceptable for small numbers of users, but become unwieldy as the user base grows, in terms of both the required hardware and the necessary support personnel.

The best BI tools are architected first for Web deployment by using a Java™-based approach with nonpersistent connections. To enable a truly scalable BI environment, the following architectural characteristics are essential:

- Centralized management
- Server-based processing
- Zero-footprint client
- Minimal user training

If you aren't sure a particular product has essential ingredients for scalability, ask the vendor these important questions:

Can it work with data where it lives?

Many BI products are unable to work with existing data. Instead, they require that customers move data around and populate local data marts or cubes. This is tolerable only for deployments of fewer than 100 users. As the number of users increases, such products become difficult to deploy and manage.

Is it Web architected, or just retrofitted?

According to Gartner, many software products have evolved from personal productivity tools to multi-user and server-based offerings. To profit from the shift away from PCs to the Web, many vendors simply ported and retrofitted PC code to servers. These were highly inefficient server-based applications and could not easily scale to support large numbers of simultaneous users. To compensate for this shortcoming, vendors embraced a distributed computing approach, enabling a dispatcher to delegate (and manage) work across multiple servers. Although distributed computing improves scalability, it adds tremendous expense and management

overhead for an enterprise. Gartner emphasizes single-server scalability because, in most instances, enterprises use a single-server approach to deploy a BI solution.¹⁰

The take home point is clear: Windows-based BI products were designed for individual users or small workgroups. They do not scale to support enterprise-wide deployments, and they can't be easily customized to fit specific application needs. Others are Web deployable but process requests in batch mode. This means they deliver static Web pages or documents that users can't interact with, and they don't change when the underlying data changes.

Can it handle large data sets efficiently?

Very few products can work with large data sets efficiently. To make the grade, they must provide native access to data sources (rather than using ODBC or some other generic high-level access mechanism). Additionally, they must perform resource-intensive processing activities on the server rather than the client, such as aggregating data and fielding selection criteria. Without this level of server-based data-handling efficiency, performance will suffer.

Will it query multiple data sources in a single pass, with multiple outputs?

If it won't, system administrators will face network performance and system efficiency issues as the number of users grows.

What exactly will it take to scale the BI environment?

BI products that grew up in the client/server world require many more servers – and consequently many more support personnel. Those products that were initially based on a Web architecture enforce centralized management. These products can support a growing environment with minimal training and maximum user acceptance. As they scale, they require minimal additional hardware.

As Dresner sums up at the end of his informative research note, “Business intelligence vendors are proclaiming their products' ability to scale to large numbers of users. Enterprises should not take these claims at face value, but instead must question vendors and customer references and understand how product scalability is achieved. Even when comfortable that a vendor has a suitable and scalable architecture, enterprises should test the product and simulate real-world situations.”¹¹

¹⁰ H. Dresner, “Scalability and BI: What You Need to Know Now,” April 14, 2004.

¹¹ *ibid.*

The Unique WebFOCUS Architecture

Information Builders' WebFOCUS has a flexible, reliable Web architecture that can scale to meet the needs of a growing user base while delivering consistent, efficient performance at a low cost. WebFOCUS was built entirely by Information Builders. Many other BI products were cobbled together from third-party technology. These BI tools lack the architectural consistency that enables stability and top-notch performance.

Here are a few of the unique characteristics of the WebFOCUS architecture that give it the ability to scale from 5 to 1,000 times as well as its major competitors, depending on the configuration and data sources:

Nonpersistence – The WebFOCUS Reporting Server uses nonpersistent technology to connect users to the server. Users who are logged into the server do not consume resources until they submit requests. Many BI solutions require persistent connections, which means that they must allocate server resources for each connected user even if those users are not submitting queries or generating reports. A nonpersistent connection and a multi-threaded Java-servlet architecture make it possible for WebFOCUS applications to scale to support large numbers of concurrent users without significant hardware investments.

Server multithreading – With most BI technologies, each user consumes server resources – such as CPU, I/O, and memory – because the connections are not multithreaded. As the total number of concurrent users increases, resources become scarce, and servers can quickly become overloaded (particularly in Microsoft Windows NT environments). WebFOCUS is multi-threaded, allowing multiple users to efficiently share the same server resources. This guarantees greater concurrency.

Native data adapters – Relational database vendors have developed their own dialects of SQL. Many of these SQL extensions improve database performance. However, you need to take advantage of the native dialects to reap the performance benefits. Unfortunately, many BI solutions use ODBC to connect to the database. ODBC is a generic middleware translation layer that generates ANSI-standard SQL. While these standard SQL calls can be understood by all relational databases, they cannot take advantage of unique SQL extensions – many of which have been designed to improve database-processing performance. In some cases, additional network traffic is generated when using ODBC. WebFOCUS uses only native database drivers, partly so it can take advantage of database-specific SQL extensions for optimum performance. WebFOCUS even provides native drivers for nonrelational data sources, allowing it to access

data sources that are not accessible by other BI solutions. This represents an intangible form of scalability, since other BI solutions must establish a separate middleware process to move or transform data into a relational source before it can be accessed.

Multi-tier configurations – One of the biggest hindrances to scalability is network traffic. Many Web-based BI solutions must be installed directly on a Web server. In some cases, they only run on Microsoft Windows NT platforms. When the target data source is not stored on the Web server – or not located on a Windows NT platform – data must be moved from the database machine across the network to the BI server before reports can be generated. This is not a big deal where small data sources are concerned; however, it can be a deathblow if you are accessing millions of rows of data, or terabyte-sized data files. WebFOCUS can be distributed across multiple servers, and it supports many different configuration options. For example, the WebFOCUS Report Server can be installed on Windows NT, UNIX, Linux, AS/400, and mainframe computers, enabling developers to place the reporting logic on the same platform as the data. In that instance, WebFOCUS can crunch the data into a report before it is passed across the network, significantly reducing network traffic.

Superior data manipulation technology – WebFOCUS has built-in data-manipulation capabilities that have been developed and refined for almost 30 years. Its native data manipulation capabilities far exceed any SQL dialect. Because of this, it is capable of generating reports in one or two passes through the database. To obtain the same results with a standard SQL-based tool might require ten passes or more.

Load balancing, N-tier processing, and fail-over – WebFOCUS load balancing intelligently routes requests to servers with the most available capacity. N-tier processing allows developers to partition application logic across servers to maximize system performance. WebFOCUS applications are highly portable, from PCs to mainframe systems. Fail-over re-routes users to another available server in the event of an outage.

Zero-footprint client – WebFOCUS applications are accessible to everyone throughout the enterprise without requiring system administrators to install and maintain software on every desktop.

Automatic report bursting – Intelligent bursting eliminates the overhead of creating, distributing, and maintaining multiple copies of the same basic report. For example, a single report can be run once, and then different sections delivered to designated individuals, regions, or offices – all in a single process.

Fully portable application components – WebFOCUS applications are 100 percent portable, enabling administrators to meet burgeoning demands for information as the popularity of an application grows. For example, an application could start out on a Microsoft Windows NT server, then later be scaled to a UNIX system or even mainframe platform, with little or no change to the software.

Integration with other application development efforts – WebFOCUS developers can program links to both J2EE and .Net environments, making it easy to extend their BI assets as part of a Web Services strategy.

Built-in monitoring tools – The WebFOCUS Resource Governor preemptively manages resource consumption to ensure exceptional performance and availability. Sophisticated governing technology helps system administrators predict the impact of a query or report request before it is executed, and to catch resource-intensive operations before production systems are impaired.

WebFOCUS in Action: Large-Scale Deployment Case Studies

Thousands of customers use WebFOCUS for data query, reporting, and analysis activities. Here are a few examples of large implementations, where scalability is essential.

Ford Motor Company

Ford Motor Company created a self-service application to help auto dealers and warranty consultants quickly identify and resolve problems with warranty repair service and claims. Developers replaced a difficult-to-read printed report with a dynamic, Web-based application that uses graphics to instantly identify anomalies and isolate problems. This self-service reporting application not only gives dealers a snapshot of their warranty performance, but also includes powerful tools for examining their businesses. More than 48,000 dealer employees use the system.

Moneris Solutions

Extended applications represent the fastest growing market for business intelligence applications in the years ahead. In some cases, the implications are dramatic, particularly when a self-service application yields a competitive advantage – such as a new way to initiate and solidify relationships with customers. Consider Merchant Direct, a WebFOCUS application that

allows merchants to view debit, Visa, and MasterCard transaction data online. Created by Moneris, a Canadian financial institution, Merchant Direct lets customers generate consolidated reports and obtain customized views of card payment activity. Today, Moneris Solutions serves 25,000 clients, with 1,000 merchants being added every month, using WebFOCUS.

RBC Royal Bank

RBC Royal Bank in Canada is using business intelligence technology to help its customers reconstruct up to six years of banking transactions. Their new self-service application is yielding a dramatic boost in productivity for customer service personnel, who formerly had to reconstruct transaction statements by combining information from legacy information systems with microfiche records. Royal Bank has more than 10,000 users generating about 900,000 reports per month. The response time for creating a report that accesses three years of transactions is about 20 seconds.

Another major bank, which prefers to remain anonymous, has 300,000 WebFOCUS users, and is ramping up to 2 million users.

City of New York

In the public sector, the New York City Department of Health is using WebFOCUS to share health inspection information with citizens via the Internet. The site is immensely popular, logging about 40,000 hits per hour.

About Information Builders

Information Builders was founded 29 years ago with a very simple mission: to help business people access information easily and quickly. We have succeeded. Today, our customers can efficiently put information in the hands of everyone they rely on to support, supply, run, and transact with their businesses. Thousands of customers and partners all over the world – including most of the Fortune 100, all major federal agencies, and many educational institutions – use Information Builders’ software and services to turn data into actionable information that helps their customers, employees, managers, and business partners make better decisions.

We’re an established vendor with mature solutions:

- \$300 million in revenue
- 29 years of experience
- More than 1,750 employees
- Over 12,000 customers
- 47 offices and 26 distributors worldwide

Information Builders sets the standard for enterprise business intelligence and enterprise reporting. Our WebFOCUS product is a fully integrated suite that enables customers to make fast decisions based on complete information. We provide more access to more data sources and run on more platforms than any other BI vendor.

Appendix A – WebFOCUS Benchmark Test Results

This table reveals the output capacity generated by WebFOCUS on a Unisys ES7000 server with 32 processors. The test used a 26-page PDF document as the output. The test configuration produced 1,472 report pages per second over a one-hour period:

Information Builders Production Reporting Results

Number of Reports	Number of Pages	Total Time to Process	Pages per Second	Pages per Hour	Pages per Day
203,815	5,299,200	60 minutes	1472	5,299,200	127,180,800

The table below takes the total volume of output depicted in the table above, and divides those results by the number of processors used in order to determine the number of pages per second per CPU.

Information Builders Production Reporting Results – Pages per CPU per Second

Number of Reports	Number of Pages	Total Time to Process	Pages per Second	Pages per Hour	Pages per Day	Number of CPUs	Pages per CPU per Second
203,815	5,299,200	60 minutes	1472	5,299,200	127,180,800	32	46

Figure 2, Production reporting: pages per second per CPU.

These tests reveal that WebFOCUS scales to readily support more than 1,000 concurrent users involved in report execution activities.

Appendix B – Independent Benchmark Comparisons, Cognos Versus Information Builders

Howard A. Kanter, associate director of the Laboratory for Software Metrics at DePaul University in Chicago, tested the validity and reliability of benchmark data from both Cognos and Information Builders.

Using publicly available information, Kanter compared benchmark test results from both Cognos and Information Builders. His findings are summarized here, in four charts.

The first chart shows that Information Builders' WebFOCUS, running on a 32-CPU server, delivered consistently better response times than Cognos' ReportNet, running on a 44-CPU server:

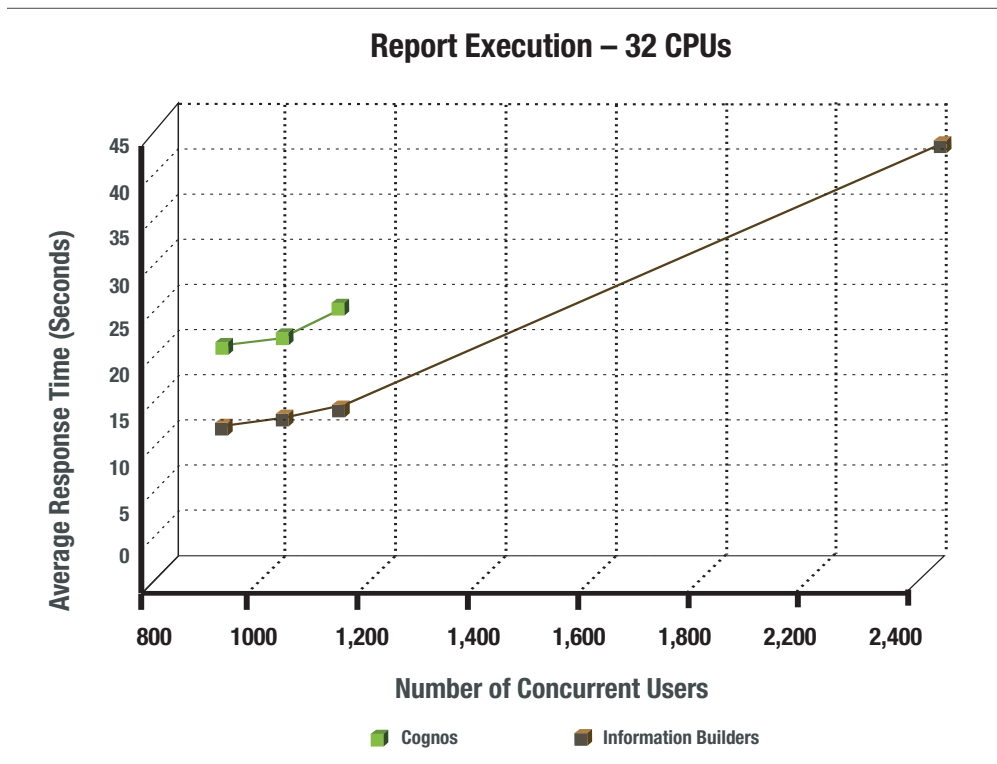


Chart 1, ReportNet versus WebFOCUS: average response time for report execution, as a result of increasing users.

Chart 2 reveals that for a range of 200 through 500 concurrent users, WebFOCUS delivered consistently better response times than ReportNet (both products were tested on servers with 16 CPUs):

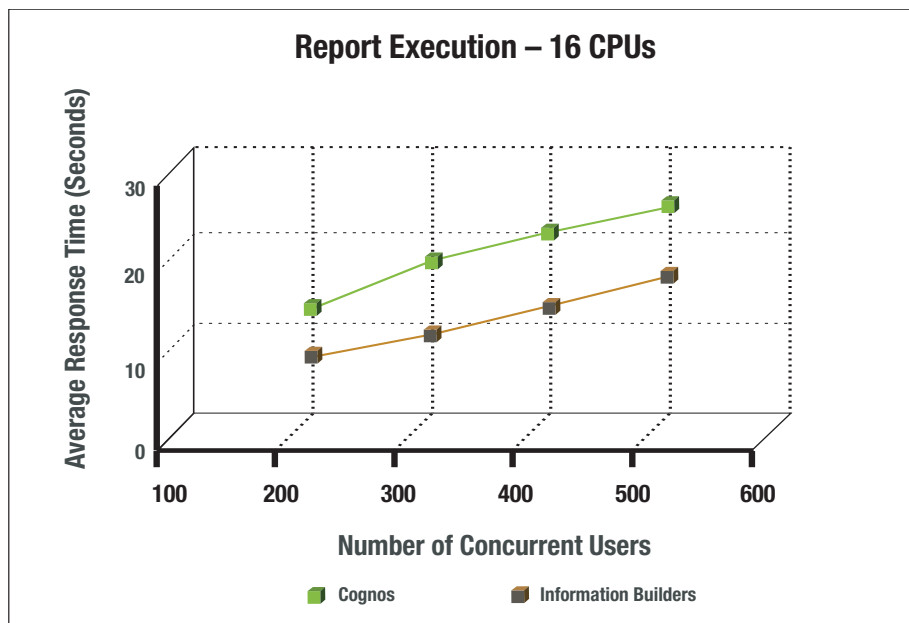


Chart 2, ReportNet versus WebFOCUS: average response time as a result of increasing users.

Chart 3 depicts pages per second for both products. WebFOCUS (running on a 32-CPU server) produced 1,472 pages per second. ReportNet (running on a 36-CPU server) produced only 79 pages per second over a comparable period.

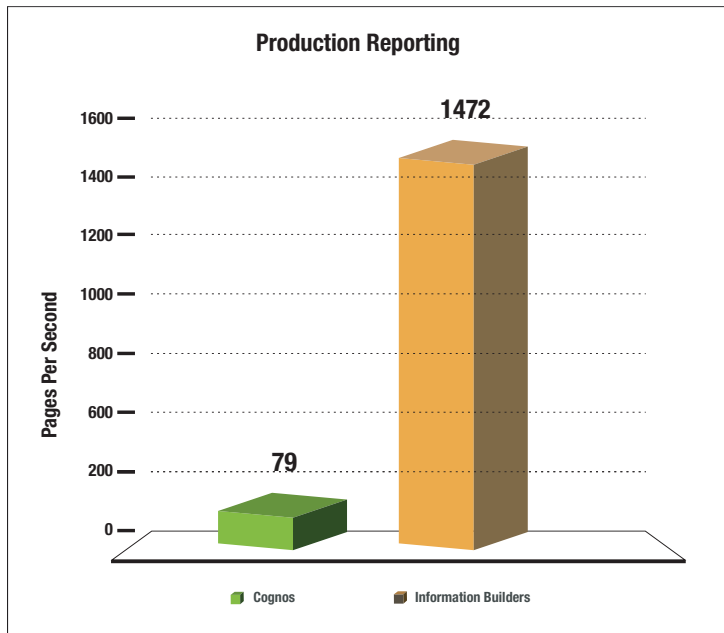


Chart 3, Production Reporting, ReportNet versus WebFOCUS.

Chart 4 demonstrates that, on a per-CPU basis, WebFOCUS is able to produce 45 pages per second, compared to ReportNet, which produces 2 pages per second.

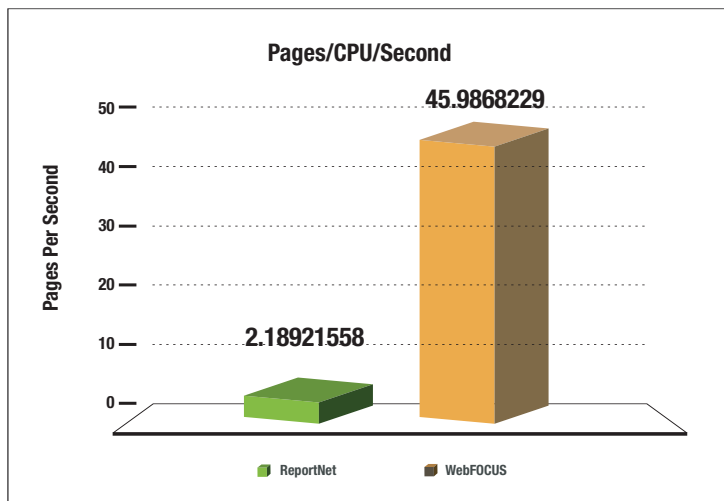


Chart 4, ReportNet versus WebFOCUS, pages per CPU per second.

This independent verification clearly demonstrates that Information Builders' WebFOCUS scales more effectively and produces output pages more efficiently than Cognos' ReportNet. Kanter believes the procedures followed and the results obtained were reasonable and consistent with procedures usually followed in preparing benchmark tests.¹²

¹² "Report on WebFOCUS Scalability, Benchmark Test Results," April 2004.

Sales and Consulting Offices

North America

United States

- **ATLANTA,*** GA (770) 395-9913
- **BALTIMORE,** MD CONSULTING: (703) 247-5565
- **BOSTON,*** MA (781) 224-7660
- **CHARLOTTE,*** NC CONSULTING: (704) 494-2680
- **CHICAGO,*** IL (630) 971-6700
- **CINCINNATI,*** OH (513) 891-2338
- **CLEVELAND,*** OH (216) 520-1333
- **DALLAS,*** TX (972) 490-1300
- **DENVER,*** CO (303) 770-4440
- **DETROIT,*** MI (248) 743-3030
- **FEDERAL SYSTEMS,*** DC (703) 276-9006
- **HARTFORD,** CT (860) 249-7229
- **HOUSTON,*** TX (713) 952-4800
- **LOS ANGELES,*** CA (310) 615-0735
- **METROPOLITAN,*** NY SALES: (212) 736-7928
CONSULTING: (212) 736-4433, EXT. 4443
- **MINNEAPOLIS,*** MN (651) 602-9100
- **NEW JERSEY*** (973) 593-0022
- **ORLANDO,*** FL (407) 804-8000
- **PHILADELPHIA,*** PA (610) 940-0790
- **PITTSBURGH,** PA (412) 494-9699
- **ST. LOUIS,*** MO (636) 519-1411
- **SAN JOSE,*** CA (408) 453-7600
- **SEATTLE,*** WA (206) 624-9055
- **WASHINGTON,*** DC SALES: (703) 276-9006
CONSULTING: (703) 247-5565

Canada

INFORMATION BUILDERS (CANADA) INC.

- **CALGARY** (403) 538-5415
- **MONTREAL*** (514) 630-1134
- **OTTAWA** (613) 233-0865
- **TORONTO*** (416) 364-2760
- **VANCOUVER*** (604) 688-2499
- **VICTORIA** (250) 995-8674

Mexico

INFORMATION BUILDERS MEXICO

- **MEXICO CITY** 85-55-91-71-20-54

Australia

INFORMATION BUILDERS PTY. LTD.

- **MELBOURNE** 61-3-9631-7900
- **SYDNEY** 61-2-8223-0600

Europe

- **BELGIUM** INFORMATION BUILDERS BELGIUM
BRUSSELS 32-2-7430240
- **FRANCE** INFORMATION BUILDERS FRANCE S.A.
PARIS 33-14-507-6600
- **GERMANY** INFORMATION BUILDERS (DEUTSCHLAND)
DUSSELDORF 49-211-522877-0
ESCHBORN 49-6196-77576-0
MUNICH 49-89-35489-0
STUTTGART 49-711-7287288-0
- **NETHERLANDS** INFORMATION BUILDERS
(NETHERLANDS) B.V.
AMSTERDAM 31-20-4563333
- **PORTUGAL** INFORMATION BUILDERS PORTUGAL
LISBON 351-217-230-720
- **SPAIN** INFORMATION BUILDERS IBERICA S.A.
BARCELONA 34-93-344-32-70
BILBAO 34-94-425-72-24
MADRID 34-91-710-22-75
- **SWITZERLAND** INFORMATION BUILDERS
SWITZERLAND AG
WALLISELLEN 41-1-8394949
- **UNITED KINGDOM** INFORMATION BUILDERS
(UK) LTD.
LONDON 44-208-9824700

Representatives

- **AUSTRIA** RAIFFEISEN INFORMATIK CONSULTING GMBH
VIENNA 43-12-1136-3870
- **BRAZIL** INFOBUILD BRAZIL
SÃO PAULO 55-11-3017-5178
- **CHINA** INFOBUILD CHINA, INC.
SHANGHAI 86-21-50805432
- **FINLAND** INFOBUILD OY
ESPOO 358-9-7250-2250
- **GREECE** APPLIED SCIENCE
ATHENS 30-210-699-8225
- **GUATEMALA** IDS DE CENTROAMERICA
GUATEMALA CITY 502-361-0506
- **GULF STATES** ▪ **BAHRAIN** ▪ **KUWAIT** ▪ **OMAN**
▪ **QATAR** ▪ **SAUDI ARABIA** ▪ **YEMEN**
▪ **UNITED ARAB EMIRATES**
AL-GOSAIBI INFORMATION SYSTEMS
973-274-090
- **ISRAEL** NESS A.T. LTD.
TEL AVIV 972-3-5483638
- **ITALY** SELESTA G C APPLICATIONS S.P.A.
GENOVA 39-010-64201-224
MILAN 39-02-2515181
TORINO 39-011-5513-211

- **JAPAN** K.K. ASHISUTO
OSAKA 81-6-6373-7113
TOKYO 81-3-3437-0651
- **KOREA** UNITECH INFOCOM CO. LTD.
SEOUL 82-2-3477-4456
- **MALAYSIA** OPTEGRA SDN BHD
SELANGOR 60-3-80240188
- **NORWAY** ISOLUTIONS AS
STAVANGER 47-51-44-47-44
- **PHILIPPINES**
BEACON FRONTLINE SOLUTIONS, INC.
63-2-750-1972
CORPORATE INFORMATION SOLUTION
63-2-633-1321
- **POLAND** COMPFORT/MERIDIAN POLSKA SP
WARSAW 4822-630-2660
- **SINGAPORE**
AUTOMATIC IDENTIFICATION TECHNOLOGY LTD.
65-6286-2922
LEGATO SOLUTIONS AND SERVICES PTE LTD.
65-684-63150
- **SOUTH AFRICA** INTERNATIONAL COMPUTERS S.A.
(PTY.) LTD.
JOHANNESBURG 27-11-2335911
- **SWEDEN** CYBERNETICS BUSINESS SOLUTIONS AB
SOLNA 46-7539900
- **TAIWAN** GALAXY SOFTWARE SERVICES
TAIPEI 886-22-3897722
- **TURKEY** ISTANBUL
ERDEMISOFT 90-212-257-5555
KEY SOFT LTD. 90-216-428-5933
- **VENEZUELA** INFOSERVICES CONSULTING
CARACAS 58-212-763-1653

Toll-Free Numbers

- **SALES AND INFORMATION**
(800) 969-INFO
- **VAR AND RESELLER INFORMATION**
(800) 969-4636

*Training facilities are located at these branches;
additional locations are available.



Corporate Headquarters Two Penn Plaza, New York, NY 10121-2898 (212) 736-4433 Fax (212) 967-6406

www.informationbuilders.com askinfo@ibi.com

Canadian Headquarters 150 York St., Suite 1000, Toronto, ON M5H 3S5 (416) 364-2760 Fax (416) 364-6552

Copyright © 2004 by Information Builders, Inc. All rights reserved. [30]

All products and product names mentioned in this publication are trademarks or registered trademarks of their respective companies.

For International Inquiries +1(212) 736-4433



Printed in the U.S.A.
on recycled paper