Scalability Requirements for Business Intelligence Platforms

White Paper

Aligning Business and IT to Improve Performance

Ventana Research Headquarters
1900 South Norfolk, Ste. 280
San Mateo, CA 94403
info@ventanaresearch.com
(650) 931-0880

Ventana Research Europe
2nd floor, Berkeley Square House
Berkeley Square, W1J 6BD London
europe@ventanaresearch.com
44 (0) 20 7887 6012

Scalability Requirements for Business Intelligence Platforms

White Paper

Ventana Research

Corporate Headquarters
1900 South Norfolk, Ste. 280
San Mateo, CA 94403
info@ventanaresearch.com
(650) 931-0880
www.ventanaresearch.com

Europe
2nd floor, Berkeley Square House
Berkeley Square, W1J 6BD London
europe@ventanaresearch.com
44 (0) 20 7887 6012

TABLE OF CONTENTS

San Mateo, California

May 2005

Ventana Research authored this paper and performed the associated research for a fee. The purpose of the paper is to frame evaluation criteria for enterprise-scale business intelligence (BI) platforms. Ventana Research worked with various Information Builders clients to write historical accounts of their WebFOCUS deployments as illustrations of highly scaled enterprise BI deployments.

Our explicit purpose was the articulation of the requirements of large enterprise-scale deployments of business intelligence. This research is not intended for use outside of this context and does not imply that organizations are guaranteed success by using only this evaluation framework and the associated case studies.

We certify that Ventana Research wrote and edited this report independently; that the analysis is a faithful representation of our experience and knowledge of requirements for enterprise-scale business intelligence platform deployments; and that the analysis and conclusions we made are our own.

Executive Summary

Business intelligence is now a front-and-center activity in which many enterprises, large and small, are differentiating themselves from their competition. The rules of the game have changed from when BI was in its formative years. Instead of departmental deployments of ad-hoc tools to strategic analysts, BI has a new role, namely providing operational information to an enterprise's front-line workers. Businesses are now using BI to enable operational workers to make better daily decisions, improving their efficiency and effectiveness. In some cases these deployments are to tens or even hundreds of thousands of organization employees. More progressive businesses have even deployed business intelligence applications to their customers, who in some cases number in the millions, to improve their customers' profitability, their net worth and ultimately their relationship with the business. Regardless of whether the deployment is internal or external to the enterprise, the result is, in many ways, the same: improved business profitability and customer relationships.

So what's the big deal? Haven't we been providing reports to our front-line workers and our customers for years? Well, yes we have been doing that. But, in many cases those reports delivered too much data and not enough information at the same time. Users also seem to want never ending changes to their reports. And this issue is only going to get worse. With the advent of Google, RFID, Sarbanes-Oxley, the Patriot Act, worldwide competition, etc. we have magnitudes more data now than ever before. And this ever increasing onslaught of data shows *no…sign…of…letting…up…ever.* Developing a few more reports is not the way organizations will be able to keep up.

Organizations need to realign their perspective about information delivery from reports to business intelligence *systems*. Further, these systems must be deployed upon a business intelligence technology architecture that is capable of:

- Rendering the mountains of data into nuggets of information
- Delivering those information nuggets just in time, not just in case
- Assuring that the information is the latest and best possible
- Providing an enterprise-wide view of customers, products and services

Organizations intending to deploy large-scale BI systems now need to understand the important architectural requirements that will assure the success of their BI system deployments. These requirements point to a fundamental capability, namely that of scalability. But scalability is no longer a simple notion of supporting more users. The scalability required to support enterprise-wide operations is more complex. Of course, the ability to support large numbers of users is an essential facet of scalability. However, enterprises are a mélange of roles, users, business processes, business applications and data sources. Enterprise-wide scalability entails accommodating both the size *and the variation* of an enterprise.

Line-of-business managers often think that business intelligence platform architecture is the domain of IT specialists. Until they need crucial information, and then, overnight, everyone seems to be an expert. Ventana Research suggests that a proactive approach to deploying an enterprise BI system is best. To support that approach, this whitepaper attempts to provide two benefits to line-of-business and IT managers. First, it provides a vision of enterprise-wide BI scalability that reflects real-world issues of deploying enterprise-wide BI applications. Second, it provides a framework to evaluate the scalability potential for BI platforms.

So what are the issues to be addressed? Enterprise scalability requirements for BI platforms are best understood by breaking them down into different categories. Ventana Research believes that the relevant categories are:

User scalability – BI platforms must support large numbers of simultaneous queries with fast response times. The implication is system architectures must be optimized for large numbers of similar queries rather than for small numbers of widely varying ad-hoc queries.

Analytic scalability – Analytic scalability measures the variation of analytic needs that a BI platform can support. When deployed across enterprise departments with different purposes, variation will be significant.

Application scalability – Application scalability is the measure of a BI platform's ability to fit into an IT environment while meeting the specific analytic needs of the business. Here scalability has to do with the breadth and extent of variations the BI platform must support.

Data scalability – Accommodating the growing volume of data is an ongoing issue within organizations. Scaling to accommodate more data requires support for multiplicity of data sources, schema complexity and very large data set size.

System scalability – User, analytic, application and data scalability requirements all drive the need for larger business intelligence systems. Business intelligence systems must leverage hardware infrastructure expansions in a proportionate, linear fashion. Along with response time, throughput and application concurrency, linearity is also an important scalability metric.

Evaluation Framework Introduction

Enterprise-wide scalability for business intelligence is now a strategic topic within IT organizations. Driving this surge in BI popularity is the enterprise need to deliver valuable information out into the farthest reaches of the organization and beyond. Enterprises now realize that by doing so, they can improve efficiency and effectiveness of their business operations, resulting in significantly improved bottom-line profitability. That profitability is possible not because all workers can now participate in strategic decision making. Instead, that profitability comes from improving the daily, even hourly, decision making done within business operations.

The evaluation of an enterprise capable business intelligence platform is complex because BI systems have many moving parts, there are many enterprise implementation specific factors and vendors have an inclination to promote only the requirements that best fit their offering. This white paper prescribes a framework to evaluate BI platform scalability. The intended audience is senior IT and line-of-business executives managing enterprise-wide BI deployments.

Importance of Business Intelligence for Operations

Operational use of business intelligence is not completely new. In fact, small numbers of enterprises have had enterprise-wide deployments of BI to operations for several years. The essential point is that going forward, *not having* enterprise-wide BI deployments to operations will likely leave organizations at a competitive disadvantage. Global 2000 enterprises are now pushing the boundaries on scale with deployments to 100,000's and even millions of users. In all of these cases, the usage mode is operational, namely, to support frequent decisions about operations. These operational deployments of business intelligence platforms are differentiated from small scale deployments for strategic or tactical decision making via several essential attributes: 1) usage by a high number of individuals either internal or external to the organization, 2) frequent use of the information and 3) users that require easy-to-learn, simple-to-use application interfaces.

The 2004 Ventana Research Business Intelligence for Operations study showed that 73% of the organizations surveyed were deploying business intelligence applications for operational use. Further, almost half of those respondents use those applications on a daily, hourly, or more frequent basis. For many organizations, BI technology is an integral part of software infrastructure deployed for business operations. Ventana Research expects the penetration of BI into operations to accelerate as organizations recognize its value.

Applications are the primary form by which business intelligence technology is delivered to operational users. Unlike analysts who use BI as an ad-hoc tool to investigate issues that vary widely, operational users need their business intelligence information to support specific tasks that are repeated continually over time. Because these tasks vary by the role of the operational user, often many different BI-based applications are deployed simultaneously. Ventana Research data shows that half of the organizations that deploy business intelligence for operations have deployed 3 or more applications. Intentions to add more data, more users and more applications are prevalent throughout the industry.

Ventana Research study results indicate organizations that deployed business intelligence applications for operations gained not only increased data access, but also improved efficiency and customer service, and reduced costs. Specific usage modes of these applications follow distinct trends. Comparison of actual performance to planned or forecasted performance was the most frequently cited form of analysis (83%) in a recent Ventana Research study. Gaining a single view of KPIs, ability to uncover threats/opportunities, ability

to perform root-cause analysis and the ability to gain a complete customer or product view are also important usage modes. Factors driving use include businesses' desire to improve profitability through cost cutting and efficiency improvements.

Assessment Framework

Business intelligence has not traditionally been a strategic initiative within organizations. In past years, other initiatives such as ERP application deployment, B2B or B2C e-commerce, Y2K and CRM have had higher priority. During those times, deployment of business intelligence was often handled on a limited basis, either in scale (i.e. departmental sized user communities) or in usage modes (i.e. OLAP only or reporting only, etc.). Investments in BI technology and the resources required to deploy and maintain them could be handled at a lower level within IT organizations.

The advent of enterprise-wide deployments of BI to operations changes the investment scope for BI. BI is now a key CIO-level initiative within many organizations today. Along with increased criticality comes a need by senior IT management to understand the factors required for successful deployment. Cost, time and personnel requirements must all be estimated at the enterprise level. Essential to that estimate is the fundamental ability of the chosen BI platform to scale.

Approaches to Assessing Scalability

Accomplishing the complex task of assessing business intelligence platform scalability should be achieved by combining two approaches, empirical assessment and bottoms-up assessment. Via the empirical approach, architects investigate other similar BI applications and use data from these deployments to help estimate user community size and application infrastructure requirements. Architects can also use vendor or independent benchmarks as proxies for understanding the system requirements for a given user community size.

Empirical data can provide a starting point, but often doesn't provide sufficient information. Augmenting that data with an understanding of the business intelligence platform's architecture can overcome scalability estimation gaps. Specifically, architects should have a detailed bottoms-up understanding of the features available within a BI platform which assist with scalability. The following reference architecture provides a framework by which architects and IT managers can structure their assessment of BI platform scalability features.

Scalable Enterprise-wide BI Platform Requirements

Based upon recent research by Ventana Research into enterprise-wide operational deployment of business intelligence, factors that affect success of those deployments can be divided up into the following categories:

- User scalability
- Analytic scalability
- Application scalability
- Data scalability
- System scalability

For enterprise deployments of business intelligence to be successful, requirements in each of these five categories need to be met.

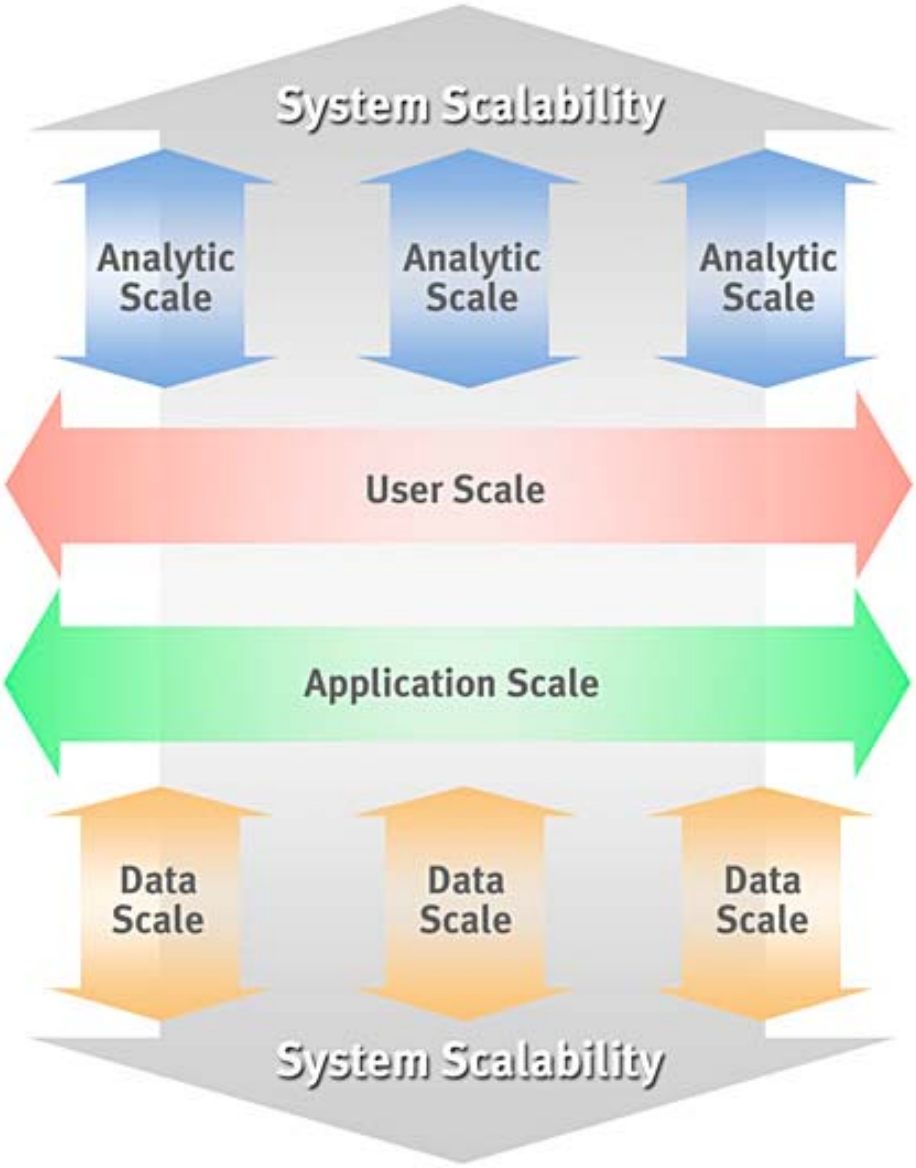

User Scalability

User scalability is perhaps the most essential capability for a scalable system. Driven by the need to deploy information to front-line workers in operations, businesses are now faced with the challenge of deploying business intelligence software to large user communities. As these workers must execute a large number of tasks within their work day, little time is available for conducting ad-hoc analysis. Instead, they need to have rapid-fire access to the specific information needed to complete their tasks. This means that BI platforms must support large numbers of simultaneous queries with fast response times. The implication is system architectures must be optimized for large numbers of similar queries rather than for small numbers of widely varying ad-hoc queries.

Ease of use – BI platforms used in enterprise-wide deployments must be optimized for extreme approachability and ease-of-use, so that front-line workers can interact with little or no training. This is critically important as front-line workers have so little time for system malfunction or UI ambiguity. Similarly, IT departments cannot expect to scale support organizations to provide similar support-to-user personnel ratios found with departmental deployments.

Alerts – Enabling users to interact in the manner in which they most prefer also impacts user scalability. One aspect of this is driving user interaction only when a critical event arises. In sync with users' desire to manage by exception, a system's ability to push notification of and information about critical business events can mean dramatically larger scalability of a system.

Analytic Scalability

Analytic scalability measures the variation of analytic needs that a BI platform can support. When deployed across enterprise departments with different purposes, variation will be significant. However, the added complexity of deploying to large scale operations furthers the extent of analytic scalability required.

Right-time data – Front-line workers live in the here and now. With a short-term time horizon, their data often needs to be only the most recent. BI platforms used to support these users often access data from transactional systems as well as data warehouses. The ability to provide up-to-date insight on customers, products, and other operational processes is an essential function of enterprise-wide deployments of BI platforms.

Data visualization and manipulation – BI platforms that serve front-line workers must also present data so that key insights can immediately be gained. Further, for analyses that require interaction, the controls for those interactions must be unambiguous as to their purpose and reliable in their operation.

Application integration – BI platforms deployed to operations are often developed into 'applications'. These BI applications can also be embedded in other OLTP, ERP, CRM and other transactional applications to provide an integrated operational tool set for the front-line workers they serve.

Customization – Essential to enterprise-wide deployments of BI applications is the optimization and configuration of the deployment to solve business problems specifically identified as critical. All other irrelevant functionality is hidden or stripped away. This does not mean that the applications are not sophisticated, in fact advanced analytics and complex interactions may be hidden underneath a simplified interface.

Application Scalability

Application scalability is the measure of a BI platform's ability to fit into an IT environment while meeting the specific analytic needs of the business. Here scalability has to do with the breadth and extent of variations the BI platform must support. Enterprise-wide deployments of BI platforms will need to support a wider variety of functionality than will departmental deployments.

Business flexibility – Business flexibility has to do with the ability to accommodate the analytic needs of any part of the enterprise. Different business roles require different analytic approaches, some requiring lots of data and others requiring lots of processing.

Application consolidation – Application consolidation is an emerging trend in large organizations driven by IT cost optimization factors along with M&A activity. The outcome is the need to manage multiple applications used by different communities on top of a single BI platform instance. Application concurrency is an important measure for sites where IT infrastructure consolidation is an active initiative.

Multiple data sources – Integrating data from multiple data sources in real-time is also an emerging trend. Information is being integrated to provide a higher level of information context and to give a broader view of organizational performance. Data integration scalability occurs along several dimensions:

- Adaptability of the platform to business analysis need
- Degree to which the data can be intermingled when presented
- Number of data sources from which the data originated

Data Scalability

Accommodating the growing volume of data is an ongoing issue within organizations. Increasingly more detailed data acquisition (even at the sub-transaction level), corporate consolidation and competitive differentiation all are driving the need for broader and deeper data gathering. Scaling to accommodate these factors requires support for multiplicity of data sources, schema complexity and very large data set size.

System Scalability

User, analytic, application and data scalability requirements all drive the need for larger business intelligence systems. Need for additional capacity drives an underlying requirement for hardware system scalability. Business intelligence systems must leverage hardware infrastructure expansions in a proportionate, linear fashion. Along with response time, throughput and application concurrency, linearity is also an important scalability metric. Linear systems minimize coordination overhead amongst servers so the maximum CPU bandwidth is devoted to directly serving users' queries.

Business intelligence software must also be efficient. This efficiency is essential to minimize costs of scalability requirements and to assure the budgeted underlying hardware can support the information delivery and analysis requirements of the enterprise. Throughput is a measure of efficiency that when related to system size provides a measure of scalability.

System Scalability Features

To underpin the specific areas of scalability cited above, business intelligence platforms must support a range of computer science best practices. These include the following:

- Multi-tasking
- Multi-threading
- Data duplication
- Data handling
- Distributed processing
- Process adjacency
- Task prioritization
- Web efficiency
- System support and portability
- Optimized repetition

- Data caching
- Report bursting
- Service oriented architectures
- Native database driver support
- Customization
- Performance tuning and management

While some of these practices are at this point CompSci 101, evaluators of BI platforms will be in a better position to compare alternatives with a detailed understanding of their implementation.

Multi-tasking – Scalable business intelligence software must support large numbers of concurrent users and applications while providing dependable response time to each user. The complication is organizations now intend to deploy BI applications to tens of thousands if not millions of users. Further, the priority of each user's tasks and queries vary. At the same time, numerous tasks must be completed for each user, from data access to report formatting to handling of user commands to security. Handling these large volumes of users and tasks requires efficient designs that minimize the time delay of transitioning from one user or task to another so CPU resources are fully devoted to the essential tasks to be completed. Efficiency of a BI platform should be evaluated for its multi-tasking capabilities.

Multi-threading – An essential architectural component of BI platforms should be multi-threading. Multi-threaded software is faster than previous approaches such as multi-processing because fewer task-related data structures are copied during task switch. Multi-threading can be implemented in specific applications areas or throughout the application. The degree to which multi-threading is implemented will impact how well multiple users, process, tasks and ultimately threads are handled concurrently.

Minimal Data Duplication – As multi-threading minimizes unnecessary copying of data that doesn't change during task switches, so must business intelligence software avoid unnecessary data copying. While data duplication can improve performance when applied for caching purposes, this isn't always the case. Often data duplicated for caching purposes never gets used, so 'the bet' never pays off. In other situations data duplication can ease development complications. In either of these cases, evaluators of BI platforms should critically assess the performance implications for unnecessary data duplication. This is especially important because any persistent data structure that is a duplicate of some other data structure incurs administration overhead, increasing the cost of administration and decreasing the overall system's ease of alteration.

Minimal Data Handling – Business intelligence software involves moving data from the database to the client desktop, i.e. from one place to another. Data must traverse many layers of different software and hardware as it leaves a server's disk and arrives at a user's display. Further, data isn't moved from initial source to final destination in one move. Instead it is moved in a series of shorter, interconnected trips, moving from one memory or disk location to another location thousands if not millions of times. Each of these trips requires preparation and execution, the combination of which can be redundant and inefficient, ultimately degrading scalability potential. Evaluators of BI platforms should understand the cost of data handling to scalability.

Distributed Processing - Highly scalable systems serving large user communities require multi-processor hardware environments. Organizations that intend to deploy business intelligence to user communities larger than 10,000 will likely need not only SMP servers but a cluster of two or more of these servers. In small departmental deployments, the entire BI platform and even the database can be run on one server. In this case, tuning can be a simple process as the system automatically shifts resources to various task bottlenecks.

Larger environments require distribution of server components across servers that are more loosely coupled. In these environments, mapping a pipeline of tasks across multiple CPUs requires a somewhat more manual approach to resource assignment and tuning. The implication is that mapping of tasks to servers and CPUs must be more carefully architected in enterprise-wide environments. Different BI platforms provide different levels of assistance in the form of automation and tuning of task execution and pipelining. Evaluators of BI platforms should understand the degrees of configuration freedom available with each platform and their scalability implications.

Process Adjacency – Organizations intending to deploy to very large enterprises or to extranet environments with user communities of 100,000 or 1+ million may consider deployment of their BI platform on a mainframe computer. Mainframes provide appropriate bandwidth, reliability and administration resources for large deployments. Mainframes also provide sophisticated high-speed internal communications between applications that can be key enabler of scalability for enterprise-wide BI applications. For example, BI platforms running on mainframes can communicate with underlying databases via very high-speed memory-based buffering instead of via Ethernet LANs. These high-speed interfaces can dramatically improve the overall scalability of a BI application for operations by removing a significant communication bottleneck between the BI software and the underlying database.

Task Prioritization – Not every BI user and every BI query within a large enterprise-wide organization will have the same value and priority. As organizations transition from a multiplicity of small departmental deployments to large consolidated enterprise-wide deployments, BI platform architects will need to accommodate a much more diverse set of information needs. With larger deployments come scarce resources are shared more widely which exacerbates the need for appropriate prioritization mechanisms. Evaluators of BI platforms should assess BI platforms for their ability to prioritize users and queries based upon a set of arbitrary business rules.

Web Efficiency – Web deployment of applications is a foregone conclusion in today's IT environment. Desktop environment security, reliability and administration issues have pushed vendors and IT organizations alike to DHTML/JavaScript based user interfaces. Nevertheless, not all DHTML interfaces are equal. Simplified user interfaces often imply a greater need for processing on the BI server and the web server. The implication is that every key stroke and every screen change must be managed by the server. Along with this, latency can significantly impact scalability. Communications between the client's browser and the web server must be infrequent so that many users can share the same network wire. Simple operations should be managed locally on the browser so that the server is not burdened with frequent server queries and networks are not hampered with excessive traffic. Local caching of information on the browser should be provided to minimize network traffic. Evaluators of BI platforms should understand the implications of a thin-client browser-based UI deployment to network traffic and server-side user session state/interactivity management overhead.

System Support and Portability – Another essential capability of a scalable business intelligence system is portability across various system types. Organizations often initially deploy applications at a departmental level on a low cost hardware and operating system platform. This allows initial TCO to be in line with system benefits. As user communities for applications grow, system management sophistication must grow as well. In some instances, this may mean that the organization must transition the application to a different hardware platform and operating system. Deployment ultimately upon a mainframe may even be required. Portability from one operating environment to other larger more sophisticated environments is also an essential capability for scalability.

Optimized Repetition – An essential advantage of a computer is that it can perform a sequence of steps repeatedly very quickly without tiring or getting bored and with equal precision each iteration. As with any software, business intelligence software execution often includes task repetition. Scalable software should be optimized so that repeated steps have the shortest possible duration.

Data Caching – Caching of often used data facilitates performance. In keeping with the notion of optimized repetition, caching, either in memory or on disk, shortens the round trip for repeatedly requested data. Data traverses many systems and sub-systems as it travels from server disk drive to user display. Often this traversal time, which can be roughly equated to query response time, can extend to minutes, hours or even days. At the same time, analysis of data often requires users to visit a sequence of data views, often repeatedly. Scalable BI software minimizes the cost of revisiting data views by providing fast data caches located as near to the user as possible to minimize round-trip query response time. Further, these caches can often be shared between users so one person's interaction can facilitate another person's. Scalable systems have caches to shorten query response times and to allow users reuse of shared data.

Report Bursting – Enterprise-wide deployments of BI platforms serving thousands of users may require the creation, maintenance and distribution of thousands if not millions of reports. Minimizing queries to the database during ad-hoc analysis can greatly increase the performance and scalability for a business intelligence system. Several approaches can be implemented to reduce database query volume. Burst reporting processing is one such approach. For large groups of users where each user receives a report with the same format as their peers, but with different data, report bursting provides a reasonable solution. One large report containing individual reports for all recipient users is run in one 'burst' and then individual reports are accessed either via push or pull by each user. Individual user report access is handled by the report server without incurring database queries. Various approaches to report bursting are touted by different BI vendors. Evaluators of BI platforms should understand the extent to which report bursting is supported.

Service Oriented Architectures – The desire by enterprises for a broader span of insight at any level is driving organizations to relentlessly integrate their application and BI environments. Access to any data from anywhere in the organization is the ultimate goal. This is difficult to accomplish via 'rip-and-replace' because of various cost and productivity impact reasons. So, organizations are now using service-oriented architectural (SOA) approaches to integrate existing IT application assets. This trend will continue as it is the only economically feasible way to integrate an enterprise's application and BI assets end-to-end. Essential to the notion of SOA is reuse of information and application assets throughout the enterprise. Evaluators of BI platforms destined for enterprise-wide deployment should assess candidate platforms for their support of SOA architecture.

Native Database Driver Support – Business intelligence systems that obtain data from underlying relational databases must do so through RDBMS drivers. These drivers either support the full SQL language of the underlying database or some common subset, such as ANSI SQL. Drivers that only support ANSI SQL prevent business intelligence servers from leveraging the unique and potentially useful capabilities of a specific vendor's database. Business intelligence software connecting to databases using native drivers can gain faster query responses because the native drivers allow the underlying database to respond to a query more efficiently.

Customization – Customization to optimize execution is a basic performance tuning tactic. Business intelligence software that is customizable in multiple ways can be tuned to provide optimal performance and scalability. Customization is best enabled when an underlying scripting language is embedded in the business intelligence system. Scripts can be altered in subtle ways to overcome inefficiencies incurred during GUI-based BI object (e.g. reports, queries) development.

Performance Tuning and Management – To assure reliably scalable systems, IT organizations must be able to stress test business intelligence applications and diagnose performance bottlenecks. With time to deployment a critical qualifier of prospective business intelligence projects, the ability to quickly and easily perform stress testing and performance bottleneck diagnosis is essential. Tools should be made available to administrators which diagnose and provide a global view of the entire end-to-end system.

Governors are key safe guards that constrain access to business intelligence applications deployed to large user communities with varying information needs. Governors assure that queries and other actions taken by a user do not impact the experience and performance for other users. Governors should allow administrators to regulate query durations, rows returned and number of concurrent requests per user.

BI Platform Enterprise Scalability Case Studies

Below are a set of case studies that illustrate enterprise scalability of business intelligence platforms. Each of the case studies provides a corporate summary, an application summary and a discussion of the application's scale and the contributing capabilities to that scale.

These case studies illustrate the need for various forms of scalability:

- User scale
- Analytic scale
- Application scale
- Data scale
- System scale

Leading US Bank

Company Description – A leading US bank, with assets in the hundreds of billion of dollars, and one of the top 10 financial services holding companies in the United States. The company operates over 2,000 banking offices and over 4,000 ATMs, and provides a comprehensive line of banking, brokerage, insurance, investment, mortgage, trust and payment services products to consumers, businesses and institutions. It serves over 10 million customers, primarily through full-service branch offices in many states.

Application Background – The bank provides credit cards to a large number of corporations. To better manage their credit card product line the bank created an operational BI application to serve the bank's commercial credit card division. This business unit is driven to increase corporate customer spend using their commercial credit cards. The bank's sales and marketing team uses the application to optimize profitability of the program. This system, accessible by the bank as well as their customers, enables tracking the success of customer programs as well as the profitability of each customer. For corporate customers information like Merchant category (MCC) spend analysis; MCC spend comparison; rebate and fee schedule; and MCC level II and III reports help them decipher not only the success of their commercial card program but also arm them with information when negotiating contract rates with suppliers and vendors. Additionally, it provides controls on maverick spending. The

information gained via ad-hoc analysis allows the commercial card operations and sales teams to work with each customer to identify additional expenditures that can be transacted with commercial cards, thereby increasing spend.

Implementation Specifics – The application, rolled out in the late 90's, currently reaches 2 million users. At any one time there are between 2,000 and 4,000 active report requests by the applications servers. The platform upon which this application is based uses Information Builder's WebFOCUS as its BI server and DB2 as its database. WebFOCUS was chosen because of its track record of success with large scale reporting and ad-hoc analysis environments.

The bank's application development team used a holistic approach to designing this very large scale application. Addressing the need to support millions of users, the application runs in a clustered mainframe processor environment based on 20 IBM P-series (P650's and p690's) processors. Each of these processors supports between 200 and 400 active user report queries at any one time. Numerous performance advantages accrue from this runtime environment.

The IBM processors run as a cluster to load balance and to provide fail-over. Response times for users are guaranteed by creating multiple separate reporting processes (each with pre-allocated memory to speed start-up) on a single multi-processor server, each with only a percentage of the overall workload. By separating users amongst multiple processes, users have to compete with each other less for resources such as memory, CPU and I/O.

WebFOCUS is unique as a BI platform in its ability to communicate with mainframe databases. Because the WebFOCUS BI server processes run in the same cluster as the database, communications between them does not need to traverse an Ethernet wire. Queries submitted by the WebFOCUS processes are submitted using a memory-based interface via DB2's Call Attach API. By communicating with DB2 in this manner, query bottlenecks between WebFOCUS and DB2 are eliminated. Because the database runs on its own SMP processor, DSS queries submitted by WebFOCUS do not impede performance for transactional queries executing simultaneously.

Not all queries from WebFOCUS are directed to the co-located database. Other queries are submitted to other databases using Tibco messaging queues. These queries allow users to get the latest, most comprehensive views of their credit card spend.

Administrators use WebFOCUS's resource analyzer and resource governor to constantly monitor the performance of the application. Tuning is applied as necessary when bottlenecks are uncovered. Tuning is simplified with WebFOCUS because the application executes as a FOCUS application underneath the UI. This procedural language allows fine-grained sequence control that is necessary for the level of scalability needed. Bank developers find FOCUS easy to work with. For example, when requesting the creation of an intermediate data store such as a table, it can direct the application to not persist the table on disk, but only keep it in memory to maximize table read/write response times. When used during run-time analysis, the analyzer tells an administrator what query ran and for how long, within what I/O to which tables, columns and aggregates. Results can be used to delete unused tables and other objects while assuring that the most used objects are cached in memory to minimize access times.

Using WebFOCUS to create libraries of commonly accessed report elements, bank developers have produced libraries of standard report headers and footers. By using the FOCUS scripting language, the developer pre-caches these standard libraries in memory to minimize data reads as reports are being generated.

Ford Motor Company

Company Description

Ford Motor Company is a global company with two core businesses: Automotive and Financial Services. The Automotive business consists of the design, development, manufacture, sale and service of cars, trucks and service parts. In 2003, the Company organized its Automotive business as two primary segments, Americas and International. The Americas segment includes primarily the sale of Ford, Lincoln and Mercury brand vehicles and related service parts in North America and Ford-brand vehicles and related service parts in South America. The International segment includes the sale of Ford-brand vehicles and related service parts outside of North and South America and the sale of Premier Automotive Group brand vehicles and related service parts throughout the world. The Financial Services business includes the operations of Ford Motor Credit Company, a provider of vehicle-related financing, leasing and insurance, and The Hertz Corporation, which operates a car and equipment rental business.

Application Background

Ford Motor Company uses its dealers to provide repair services to Ford auto owners that request warranty-covered repairs. Dealers are paid by Ford for warranty repair work they do. A key facet to this partnership is that the dealers decide when to allow a warranty repair to occur. To assure that dealers do not unfairly charge Ford for unwarranted repair services, Ford runs a dealer oversight program to contain warranty service costs. The reporting and analytics for this program are deployed via WebFOCUS to 15,000 worldwide dealers with each dealer having 1-5 users.

The WebFOCUS application delivers information to dealers about alignment with claiming practices for auto repair areas such as engine and climate control. The reports provide dealer to dealer comparisons via statistical processing to identify abnormal warranty service rates. Elements that are measured include cost per repair, cost per vehicle service, cost per 1000 vehicles. The WebFOCUS application is superior to the prior hard-copy based solution because dealers get information faster, can perform quicker trending and also drill down to actual repairs. When performance statistics were first given to dealers four years ago, dealer only had simple reports printed on 8 x 11 paper with detail columns and numbers. Users can now create personalized reports including stoplighting and perform drill downs.

Implementation Specifics

Dealers have been using the WebFOCUS-based application for three years. Key evaluation points for selecting the business intelligence vendor to support this application were scalability, foreign language support, cost, and fit with the Ford IT infrastructure. The application supports 15 languages and is used by dealers worldwide.

The application runs directly against a Teradata transactional system that contains 500GB of data. The database runs on NCR hardware, the WebFOCUS application runs on windows 2000 server from IBM with 1GB Ram and 70GB RAID array per box.

Essential to assuring performance within the WebFOCUS application was preloading populated reports into the FOCUS hierarchical database. This assured fast referential queries. Concurrent users are limited to 300 per box to assure sufficient CPU time for each user.

Native drivers to Oracle and Teradata are used to assure robust compatibility and access to database specific features that enhance performance. Ongoing stress testing using WebTrends allows the application's administrators to identify and address bottlenecks.

BI Platform Enterprise Scalability Benchmark

In addition to looking at customer case studies as proof of enterprise scalability of a BI platform, it's also useful to analyze lab-based scalability benchmarks. Recently, Information Builders conducted a scalability benchmark with hardware system partner, IBM. This benchmark test intended to prove the scalability and power of WebFOCUS for Linux on the IBM zSeries server. The test pitted 100 concurrent users against 1.4 million rows of data in two files retrieving 3,000 rows of DB2 data per query. The entire system resided on a single IBM zSeries mainframe. Results demonstrated near linear scalability of the WebFOCUS BI platform.

The Test Architecture

The tests were initiated by a testing tool running on a Netfinity Server under the Windows 2003 operating system. The tool issued reporting requests against a WebFOCUS reporting server running on a Linux partition, via a middle tier. The requests accessed data stored on a DB2 database (Version 8) running on a z/OS 1.6 partition. All performance information was collected in a controlled environment.

The middle tier consisted of a WebFOCUS client, several application servers (WebSphere 5.1.1.and Tomcat 4.2), and an Apache web server running on a second Linux partition. The version of Linux used was SUSE Linux 8 SP3. Both Linux partitions ran on z990 hardware with a speed of 450 MIPS and 40GB plus 128GB xStore memory. The Linux partitions were defined within a single zVM zLinux LPAR. Each partition was configured with 2 to 8 CPUs (IFL) and 2GB of memory.

Test Results for 100 Concurrent Users

Number of CPUs	2	4	8
Average Request Processing Time for 61 Rows (Sec.)	1.007	0.503	0.434
Average Request Processing Time for 3000 Rows (Sec.)	6.913	3.427	1.638

Test Results

In the test, HiperSockets accessed DB2 data at a rate of 1GB per second. The architecture was optimized by controlling the ratio of mid-tier (web application) servers to backend reporting servers across images. This was managed, monitored, and changed via z/VM. WebFOCUS' server architecture allows it to leverage the mainframes ability to distribute workloads evenly across its environment. In this case, the WebFOCUS test bed was able to fully utilize all 8 CPUs. Spikes in context switching could signal nonproductive work, wasted CPU, and poor product integration. WebFOCUS showed a low rate of context switching, even under heavy workloads.

Analysis

WebFOCUS is one of the few BI suites that can run on a mainframe. While this is not to be overlooked, it's also important to note that many of the scalability and scale management features described in the above scalability tests are available on other platforms. Clearly, linearity is a good indicator of enterprise scalability and should be an evaluation point for any operational BI initiative. And scalability is not just about execution linearity, but also about configuration and administration.

About Ventana Research

Ventana Research is the leading Performance Management research and advisory services firm. We help companies align business requirements with information technology solutions to improve performance. What makes Ventana Research different from other analyst firms is our focus on Performance Management, the deep experience of our analysts, our research foundation and objective services. We deliver independent, expert workshops, assessments and advisory services on Performance Management and supporting information technologies.

Technology and Business Experts Comprise Analyst/Consultant Team

Our Analyst and Consultant team is made of up of both technology and business professionals with an average of over 20 years each of real-world expertise leveraging information technology to optimize business. Areas covered by our analysts and research include:

Business Research: Business Performance Management, Financial Performance Management, Workforce Performance Management

IT Research: IT Performance Management, Business Intelligence, Business Process Management, Integration Management

Operations Research: Operational Performance Management, Customer Intelligence, Supply Chain Management, Contact Center Management, Sales Management, Innovation Performance Management

Performance Management Focus

We focus strictly on Performance Management, the practice of leveraging assets to better understand, optimize and align business and processes to drive effectiveness throughout the enterprise. We are experts on both technology and business processes, and how to optimize them to drive more value for your company. We examine the underlying requirements and technology for Information Management and where any type of data or information is brought together through Integration Management.

Primary Research on Performance Management Technologies and Best Practices

We regularly conduct primary research on Performance Management technologies and best practices and share this with our clients. Recently completed research programs include: Business Intelligence for Operational Performance, Financial Reporting and Consolidation, Budgeting and Planning, TCO Benchmark Study and Operational Performance Management.

Objective Services and Knowledge Transfer

With research as our foundation combined with knowledge of best practices and methodologies, we deliver customizable services to help you reach clearly defined goals. Our BI services include:

- **Workshop:** Making Business Intelligence Pay
- BI Strategy and Technology Assessment
- TCO for BI Assessment
- BI Product Assessment and Selection
- BI Standardization and Consolidation Assessment

To learn more about Ventana Research workshops, assessments and advisory services, contact clientservices@ventanaresearch.com.