



Myths About Mainframe Business Intelligence

September 6, 2007

Intelligent Solutions Inc.

P.O. Box 4587

Boulder, CO 80306

Voice: 303-444-2411

Fax: 720-293-0060

www.IntelSols.com

Introduction

Business intelligence (BI) has come a long way from the early days of simple reports or isolated groups of statisticians locked in back offices, poring over considerable amounts of numbers and data points. Today, the entire enterprise requires real intelligence in making its decisions. BI has moved from the back office to the entire organization; every part of today's corporation needs some form of decision support.

This pervasiveness of BI has ramifications not only in our BI architectures but also in the technologies we use to support these environments. BI is bigger, faster and must support all employees now. Tremendous stress is put on our BI environments from the terabytes of data being analyzed, response times that mimic operations, sophisticated users who are globally dispersed, simple queries and complex models of behavior. BI is no longer "nice to have"; today, it is mission critical.

As we reviewed the requirements of BI today, we were struck by the synergy that SAS and the IBM mainframe offer to address the requirements of these new business intelligence workloads. There are a number of misconceptions or myths about the mainframe in the market today that are, to put it bluntly, untrue and invalid. This paper addresses these myths and illustrates some of the reasons why today's business intelligence solutions can be successfully deployed on a mainframe today.

The mainframe computer from IBM has historically been the powerhouse for corporate computing, including data warehousing. However, beginning in the 1990s, distributed servers were used increasingly throughout the enterprise to handle departmental and other workloads. The use of distributed systems were particularly attractive for business intelligence workloads as the server architectures matured, and a more robust software environment developed, including improvements in the operating systems, database management systems, and data access and delivery tools.

Over the years, server technology evolved across all the platforms. As distributed systems continue to emulate more and more of the mainframe's unique functions, such as partitioning capabilities, virtualization technologies and workload management controls, the mainframe has also evolved, supporting more of the software vendors and offerings that drive the business intelligence market today. Countless innovations enable the mainframe to provide an extremely reliable, cost-effective, large-scale platform capable of satisfying both core applications and business intelligence needs.

When choosing where to place business intelligence applications, companies should weigh their options, considering not only their requirements for today, but the applications requirements two and three years down the road. If they make their decisions based on misconceptions concerning the relative cost, complexity and capability of mainframe computing environments, then they might overlook the solution that best fits their needs.

This white paper describes 10 misconceptions about the IBM mainframe platform (now called System z) and explains why that platform is a competitive option for business intelligence.

Myth 1: Mainframe total cost of ownership is too high

Deploying any new environment can be expensive, regardless of the platform choice. Often the mainframe is perceived to be expensive and is not considered when selecting a BI platform. However, if you consider only hardware and software costs when making this decision, you are including only a portion of the overall costs of a solution. The *total* cost of ownership (TCO) for any application environment is made up of many components, including labor, hardware, software, electricity, etc. To really evaluate the full price of a solution, you must consider all the costs, including the most expensive one – the staff required to support the system.

In a distributed server environment, costs go up *linearly* with additional workload. That is, adding capacity means adding servers. Each additional server in a distributed environment adds to the human resources needed to manage and maintain the environment. In the mixed workload mainframe environment, the initial hardware costs are higher, but the per-unit cost of incremental capacity actually *decreases* as the total workload grows. With a mainframe environment, incremental capacity can often be added without necessitating staffing increases to manage and maintain the environment.

For those with an existing mainframe environment, many of the initial costs for deploying a new solution have already been paid. The incremental costs associated with adding business intelligence capabilities are therefore much lower than those for a new environment. Creating a data warehouse from data that may already be housed on the mainframe and adding a user tool can give users immediate access to valuable BI capabilities.

Further, the mainframe has recognized the need to target capacity to address specific workloads by offering dedicated utility processors. These provide a high-speed engine that reduces the overall processing costs when data is centralized on the mainframe. The economy of this solution helps break down the walls between transactional data stores and BI, enterprise resource planning (ERP) and customer relationship management (CRM) applications. These specialty processors help minimize the need to maintain duplicate copies of data across a pool of discrete systems while providing high levels of security for critical corporate data. The overall impact is to reduce the need for multiple databases by consolidating applications onto the mainframe, thus leveraging the platform's inherent strengths to manage the concurrent sharing of data by batch, online transaction processing (OLTP) and online analytical processing (OLAP) applications.

These specialty processors, called the System z Integrated Information Processors or zIIP, let customers host data intensive workloads, such as those associated with data warehousing or BI, at a greatly reduced cost. zIIP is well-suited to handle parallel queries (with an emphasis on star schema designs) common in BI workloads by improving the price/performance of parallel processing on the System z.

MYTHS ABOUT MAINFRAME BUSINESS INTELLIGENCE

A sound IT architecture involves separate environments for development, quality assurance and production workloads. In a distributed environment, this is typically accomplished with separate servers, each of which requires a redundant copy of the operating system, database management system, and application software and utilities. Not only is there a cost associated with each of the copies, but they must be synchronized. All too often, problems are discovered because of differences among the production, development and quality assurance environments.

As mentioned, another consideration is the physical movement of the data. This movement is a costly challenge. With approximately 80 percent of the world's critical transactional data already residing on the mainframe, the costs related to moving the data from the centralized server, versus leveraging the single copy of the data, can be substantial. Using a central server such as a mainframe eliminates transporting the data to another platform, thus reducing the audit and control complexities and, more importantly, the security costs. The elimination of an extra server for housing the data warehouse further reduces costs by avoiding duplication of the operating system, database management system and additional copies of other software. Additional savings are generated from decreases in reserve processing power, storage space, and consolidated security and disaster recovery approaches as well as from reduced network costs.

For companies with significant workloads, the mainframe does indeed offer an alternative that provides a lower cost of ownership than a distributed environment. A true comparison requires that you consider all aspects when using price as a criteria for selecting a BI platform.

Myth 2: Advanced data mining and analytic solutions are not available on the mainframe

IDC defines advanced analytics as

software [that] includes data mining and statistical software (previously called technical data analysis). It uses technologies such as neural networks, rule induction, and clustering, among others, to discover relationships in data and make predictions that are hidden, not apparent, or too complex to be extracted using query, reporting, and multidimensional analysis software. This market also includes technical, econometric, and other mathematics-specific software that provide libraries of statistical algorithms and tests for analyzing data.¹

The BI market offers a wide variety of options for delivering business analytics today. A vendor offering a strong portfolio of analytic and business solutions is SAS, which

¹ McDonough, B., and Vesset, D., "Worldwide Business Intelligence Tools 2006 Vendor Share: Excerpt from IDC # 207422," IDC, June 2007.

MYTHS ABOUT MAINFRAME BUSINESS INTELLIGENCE

got its start in the mainframe environment. SAS offers one of the more comprehensive suites of analytic offerings, ranging from predictive to querying capabilities. As a leader in analytics, SAS offers its advanced analytics on a range of platforms, including the mainframe. Examples of how these important analytics are used today include:

- Analyzing customer buying behaviors and determining next-best product offerings.
- Creating models of fraudulent behavior in several industry verticals.
- Predicting market and competitor actions.
- Determining optimal inventory levels and supply chain processes.
- Understanding and predicting cause and effect activities.
- Determining mean time before failure (MTBF) for critical manufacturing components.

SAS clearly offers a differentiating advantage to the platform from other BI vendors by offering strong analytics. The company continues to invest in its mainframe capabilities, not only in its traditional data mining and statistical capabilities, but also for BI analytics, where it remains the only company to do so.

Business intelligence workloads have evolved over time to require significant processing power, memory and I/O bandwidth. The mainframe is unique in the market, as it is designed to provide a balanced system optimized for a mixed workload. That is, the cache size and structure, the internal bandwidth, and the I/O structure and bandwidth are more robust when compared to CPU speeds on other platforms. This leads to an increase in the mainframe's relative capacity to do mixed workloads as opposed to its relative capacity when dedicated to single applications.

It's reassuring to know that the System z has decades of proven processing performance, undisputed reliability and an analytics provider/partner with the sophistication to take advantage of both. The mainframe also supports a number of database functions that have specific features relevant to BI activities, such as sparse indexing for star joins, recursive SQL, materialized query tables, generalized sparse index/memory data caching and dynamic indexing.

Business intelligence software offerings also leverage the latest technologies through an intelligent client or Java client, which allows the server platform to be immaterial. The mainframe remains the workhorse – a role it performs so well – while the business community interfaces through its favorite technology or technique.

The maturing BI environments, with their emphasis on sophisticated analytic capabilities, now require scale, security and performance – all of which are the well-documented advantages of the mainframe platform. For companies that must improve their competitive stances through advanced analytics, this technology is quite appealing. By combining sophisticated BI, statistical and data mining capabilities from

a company like SAS with IBM's proven platform, businesses may get just what they need to reach a new level of competitive competence.

Myth 3: The mainframe is old technology and doesn't offer the same modern capabilities as the newer solutions

Labeling today's mainframe as "old technology" implies that not much has changed in the mainframes since the introduction of distributed servers. Nothing could be further than the truth. IBM has made heavy investments in mainframe research and development, and that has resulted in many new features, some of which are noted here. SAS has ensured that its advanced data mining and statistical capabilities are designed to take advantage of that platform's enhancements as they are introduced to ensure synergy between the software and hardware platform.

System z supports modern BI components such as a true real-time or low-latency Operational Data Store (ODS), high-speed data management technologies and high performance in mixed workloads. In addition, IBM and SAS continue to invest not only in System z but also in analytic and BI capabilities.

SAS has had a strong relationship with IBM since its inception in 1976. The Business Intelligence Architecture in SAS[®]9 provides many innovative features. Some of the more noteworthy are listed here:

The threaded kernel (TK) architecture within SAS[®]9 enables SAS tasks to take advantage of multiprocessors, thereby improving scalability.

The SAS Scalable Performance Data Engine (part of Base SAS) exploits parallelism to improve throughput for data loading, index creation and data retrieval.

SAS applications are designed to migrate to new releases with few, if any, changes in the source code.

The XML engine enables data from an XML document to be imported into a SAS data set and SAS data sets to be exported to an XML document.

SAS[®]9 provides improved graphics support to enable users to produce interactive graphic outputs for the Web. These graphics support zooming, rotation, rubber-banding, drilling and dynamic switching of presentation style.

Various SAS[®]9 features enable batch report delivery to be upgraded to high-quality Web output directly from z/OS.

All of the Foundation SAS servers are supported under z/OS. This includes the SAS Metadata Server, which provides metadata repository services (through XML) to other SAS components, the SAS Workspace Server, which spawns address spaces on behalf of clients in which Foundation SAS work is executed, and the SAS OLAP Server, which enables users to develop and deploy scalable OLAP applications.

MYTHS ABOUT MAINFRAME BUSINESS INTELLIGENCE

IBM has been one of the leaders in bringing packaged solutions (operating system and middleware) into the business intelligence market with a range of warehouse offerings. IBM's range of solutions spans all its server platforms, and is pre-configured and easy for new clients to order, install, upgrade and service. These solutions automate installation, migration and maintenance tasks, making these easy to perform and less costly overall. It is worth noting that significant data warehousing capabilities have been added to DB2 in V8 and V9 and are planned for future products. These include DB2 enhancements specifically focused on warehousing.

As business intelligence workloads integrate into business systems, a key consideration is the ability to effectively control the variety of workloads that a business has to run. The Workload Manager (part of the z/OS operating system) is functionally rich, enabling significant control when consolidating workloads. Its capabilities provide for dynamic BI mixed workloads, BI and transactional workloads, and a combination of data warehouse and data mart workloads on the same system. The facility deals with issues such as varying levels of query concurrency, resource adjustments to exploit the available capacity and handle different work priorities, protection against "killer" queries, maintenance of consistent response times regardless of the workload, and dynamic resource allocation to react to loads. This is particularly important in a business intelligence environment due to a wide range of query run-times and the varying levels of sophistication in the user community.

Having a single, centralized copy of data accessed by all users makes sense, but can be a nightmare if you cannot control and differentiate between the critical work and less important queries in the system. Simple queries for operational support must execute quickly; queries for strategic analysis can tolerate a slower response, and complex analytical queries and data mining work can afford an even longer elapsed time before delivering the results. It's these fine points in managing workloads that differentiate the mainframe from other server choices. Today, these differences are no longer just critical in OLTP environments; they are mandatory in business intelligence environments as well.

Capacity on Demand (see Myth 7) follows a tradition of mainframe innovation technologies that subsequently get adopted in other environments. For example, partitioning was first developed on Amdahl and IBM mainframes in the 1980s, and predecessors of workload management were first offered in the 1970s. These were then subsequently introduced into other IBM and competitor platforms.

In addition, the mainframe has an operations model that addresses security, resilience, workload and capacity management, storage management, and business process integration. This model has proven to be nearly infallible. From a security perspective, the z/OS operating system has the highest security rating from the government. With today's business environment, this must be a key factor when considering servers that host your business's most critical asset, its data.

On the other hand, the IBM mainframe had to catch up to the other server platforms in terms of Java and open programming. Here, investments over the last 15 years have

allowed the mainframe to now provide an open programming environment that fully supports Java and C++.

The fact is that the mainframe is ahead of many of the distributed platforms in being functionally rich, with sophisticated features that ensure the highest application availability. Very often, when servers portray their reliability, they proclaim “mainframe-like” performance – a recognition that the standard in many areas is still being set by the mainframe.

Myth 4: There is minimal data integration support for the mainframe.

Both IBM and SAS have made significant investments in world-class data integration and data quality technologies. Advances in mainframe hardware and software technologies have optimized the z/OS platform for enterprise data integration and warehousing.

The Information On Demand strategy formulated by IBM years ago is now beginning to affect the marketplace and has proven to be one of the most sophisticated in its completeness of vision and technological support. SAS is compatible with IBM’s IOD data integration and quality technologies on System z.

The mainframe environment for BI leverages popular capabilities, such as parallel processing, without requiring any design changes and simultaneously supports batch and real-time operations. Both are important to evolving BI requirements such as operational BI.

The latencies associated with data integration and transfer can also greatly affect an enterprise’s ability to use its BI to make better decisions. A lot of enterprise data is still hosted by the mainframes. When the raw data is resident there, the extract, transform and load (ETL) process is streamlined since the data does not need to be migrated to another platform. Having a single system involved in the majority of data integration, clean up and transformation means the overall process is easier and less costly to manage.

With the advent of operational BI and the prominent role of the ODS, having all or most of the data resident on the mainframe means that a real-time ODS, the goal of many implementers of this technology, becomes feasible. Near-real-time data can easily be kept in sync as well with the same or different DB2 subsystems.

Myth 5: There are few mainframe professionals who know BI

With the open programming model, many of the business intelligence skills available in server environments are completely transferable to the mainframe environment. Very little additional training is needed. And, for a company that has a mainframe in place, the IT department already has a staff with appropriate programming and

MYTHS ABOUT MAINFRAME BUSINESS INTELLIGENCE

maintenance skills. Both types of human resources can easily be used to create and support BI applications for the mainframe platform.

To ensure expertise into the future, both IBM and SAS are partnering with clients and schools alike to maintain the pool of mainframe-savvy IT professionals. Both companies are working with various universities to include System z and SAS training within their curricula. Educators in over 300 schools are investigating or actively teaching mainframe technologies today, and the number of participants is continuously growing. Two examples are the IBM Academic Initiative for System z and the z/OS Basic Skills Information Center. Fully qualified BI specialists are produced from these initiatives.

When we examine mainframe skills, it is important to analyze the skills needed for different types of work specifically performed on the mainframe.

COBOL programmers – There may very well be a shortage of COBOL programmers within the US, but there is an ample supply in the off-shore programming market. Regardless, the shortage of COBOL programming skills is of concern for companies that need to maintain COBOL-based systems. However, this is not a major issue in the business intelligence industry since COBOL is rarely used as the programming language for new business intelligence applications. Mainframes can utilize Linux and Java to further reduce any dependence on COBOL programmers, particularly in the business intelligence environment.

Systems administrators, managers and operators – There is a shortage of people for these positions, but it's not restricted to the mainframe environment. Job posting sites show a far larger number of unfilled positions of this nature for the UNIX and Windows world than for the mainframe world.

IBM has taken steps to reduce the complexity of managing the mainframe environment. It is simplifying mainframe management and providing a Windows-oriented, easy-to-use graphical user interface to both enable lesser-skilled people to manage mainframes and to appeal to people who are accustomed to Windows.

Myth 6: The mainframe can't perform as well as newer solutions

Query performance is influenced by a variety of factors, including work executing on the system, the disk environment, the database design, indexes, summary tables and so on. Mainframes have demonstrated the ability to support BI workloads over the years. Often, platform decisions are made on the execution time of a single query on a server. That is an unrealistic measurement, as a BI workload does not process queries serially. Processing a mix of queries (simple and complex) is the typical performance environment. The ability of a server to handle a workload of mixed queries is a more appropriate and realistic way to evaluate a BI system. It's also an extremely challenging evaluation to perform. The complexity of the BI workload makes it difficult to run realistic tests. To manage this complexity, the platform requires a

MYTHS ABOUT MAINFRAME BUSINESS INTELLIGENCE

workload management function, and the mainframe has one of the most robust managers available.

Innovations useful to BI environments include the new Optimization Service Center that helps database administrators improve the performance of database queries. System z manages these workflows quite efficiently by prioritizing queries according to their importance and performance requirements. For example, long-running scans may run at a lower priority while shorter queries can get their results fast without waiting for the longer running ones to finish. No query can monopolize the processing or memory resources. Additionally, the mainframe has the ability to differentiate between users – some having higher priority than others.

New online utilities to simplify and speed up common database management tasks are quite useful to BI administrators. These include the ability to set workload goals so that the mainframe will manage processing, memory and I/O resources available to the z/OS image to achieve those goals. For example, an administrator might set a goal that 80 percent of all queries should complete in five minutes or less. If this goal is not achieved, the system generates a report to the administrator that indicates recommended actions to meet the goal.

Version 9 of the Query Management Facility puts a modern face on query management as well, enabling data visualization and a new generation of executive dashboards. These innovations, in particular, permit the mainframe environment to support all the evolving trends in BI, such as operational BI, business performance management (BPM) and collaboration.

Finally, consider unplanned outages. These happen frequently in the PC environment – the PC locks up, gets slower and slower, needs to be rebooted, and so on. When do you have to reboot a mainframe? Because of their legacy in the operational environment, these machines are much more fault-tolerant. A reboot is rare, with many applications operating years without outages. That level of availability is unique to the mainframe. Why else would so many vendors declare their machines have “mainframe-like” availability?

Myth 7: Mainframe administration is more costly and complex

Mainframe administration activities already exist for the mission-critical operational applications running in that environment. Proven administration support is in place to handle disaster recovery, security, regulatory compliance, performance management and application management. The incremental addition of the business intelligence workload does not substantially add to the existing mainframe administration workload. This fact alone substantially reduces the overall burden and cost for BI environments when compared to separately managed, disparate BI environments.

IBM adheres to a System z management simplification strategy that makes it easier to install, configure, administer and service. The automation of many of these tasks

MYTHS ABOUT MAINFRAME BUSINESS INTELLIGENCE

makes them more straightforward and accelerates System z administration, again reducing its costs. The goal in the mainframe environment is to minimize administration costs with extensive self-optimizing capabilities. IBM's updated user interface is common across all IBM environments, ensuring seamless (and less costly) administration. This interface is more like the UNIX and Windows environments. In addition, system performance and security monitoring is simpler with the advent of HealthChecker, which automates tasks such as software distribution and installation with the aid of a wizard.

The integration of business intelligence solutions into the operational work flow increases the visibility of both planned and unplanned outages. There are two dimensions to this issue. First is the impact of the failure on work disruption. The mainframe's reliability in this regard is widely recognized. It continues to set the bar against which other environments are compared. Next is recovery, which is not needed as often as in other environments due to the mainframe's reliability and growth path. Recovery entails not just getting the platform and application systems up and running. The database must also be restored, tested to ensure that it's not corrupted and synchronized to reflect business activities that transpired while it was out of service.

The mainframe's innovative hardware features and z/OS operating recovery capabilities have evolved substantially over its 40-year history. Downtime (in the rare instances that it occurs) is often reduced to seconds. Component failures are often invisible, with features such as dynamic disk-swapping that keep the applications and data available to users.

Anticipation of the future workload is very difficult. In the business intelligence environment, this difficulty is compounded by the unpredictability of growth, which brings about business capabilities that were not anticipated during development, and by an increasing tendency to evolve from a strategic BI environment to an operational BI environment.

In the server environment, companies can find themselves needing to install entirely new systems. This could happen either when the system is configured to its maximum or when the server technology has evolved to the point that the vendor can no longer upgrade the original system. At that point, the server replacement can be very complex, and the applications, whether they are a mission-critical OLTP system or a business intelligence environment, suffer some downtime during conversion. Thorough testing is also critical to minimize the risk of failure.

Capacity increases are made simpler with features such as Capacity on Demand, which enables hardware upgrades with minimal, if any, disruption. Frequently, the upgrades can be completed in hours while the system is running without having to install another system, install the software, redefine the database, and unload and reload the data. With Capacity on Demand, the additional capacity is already resident on the initially installed system. It simply needs to be activated after it is "purchased." The operating system immediately recognizes the new processor and begins to assign

work to it. This nondisruptive upgrade is of great benefit to administrators who have to manage the ever-growing size of their data warehouses. This is simplicity at its best.

IBM has also undertaken a five-year program focused on simplifying the management of the mainframe. This initiative is aimed at leveraging GUI front ends for management tasks while leveraging the more common skills in the market. The goal is to make it very easy to learn in a matter of a few months how to manage the mainframe.

Myth 8: Mainframe BI solutions are inherently more complex than server-based solutions

The complexity of business intelligence solutions is mostly driven by their scope and magnitude. The workload is often very intensive and, with knowledge about the workload characteristics, any skilled mainframe administrator can plan for the needed capacity to handle any current and future needs. (See Myth 7 for more on the ease of mainframe administration.)

That said, SAS' BI offerings make it simple for users across an organization to access and leverage these analytic capabilities just as they would other data sources. For example, SAS^{®9} provides targeted interfaces that align with individual skill levels to deliver intelligence in just the right business context. SAS^{®9} not only expands the capabilities of SAS, but also reinvents the way people interact with BI software. SAS developers created personas tied to these unique skill levels then conducted extensive field testing, one-on-one with users, to design and deliver user interfaces that fit the precise needs of those personas.

Meanwhile, IBM has made it easy to handle one of the more difficult and complex aspects of a BI environment – increasing the capacity of that environment as the data volumes grow. In the mainframe, capacity can be increased without necessitating an expansion into additional servers due to processing power limitations. As mentioned, the Capacity on Demand feature has removed a major hurdle for many data warehouses – the ability to increase capacity without disrupting the BI environment. The simplicity of this one critical function pushes the mainframe environment to the top of anyone's list in terms of an easy-to-administer BI environment.

Another reason for BI complexity comes from the fact that these environments are constantly evolving – hardware and software components are constantly being upgraded to newer versions. These updates can cause significant technological headaches as administrators try to ensure the compatibility of the various components. Within the mainframe environment, system and software upgrades are traditionally backward compatible. Not only is this a significant factor in reducing the costs associated with moving from one release to another, it is also significant in reducing the complexity of such a move.

Another complexity-killer is having the operational and BI environments residing on the same platform (See Myth 1 for more on this). This results in several benefits, including reduced complexity and cost; shared processes, tools and procedures; and

streamlined compliance and security. Needless to say, the reduced administrative efforts in this environment reduce not just the complexity but also the overall costs of the environments.

Myth 9: The mainframe is too inflexible to effectively handle BI applications or to quickly respond to evolving requirements

The inflexibility, when it exists, is usually not due to the computing facility. The mainframe environment provides extremely high availability and an environment that has proven to be stable for the most mission-critical applications. Its data integration capabilities (see Myth 4 for more) are based on the same well-thought-out industry data models created by IBM for such verticals as banking and retail. These models not only give implementers a productivity boost but also ensure that most, if not all, of the necessary attributes to support BI have been included and documented. Having a well-documented and thorough set of models to use for implementation certainly reduces the need to constantly make changes to the ultimate database schema.

Popular hub-and-spoke architectures, such as the Corporate Information Factory, can be fully deployed on a mainframe environment. The data warehouse hub itself can take full advantage of the mainframe database management system and its optimization prowess. The Capacity on Demand feature ensures that no disruptions will occur as the data warehouse expands and grows.

An agile staff with responsive processes can deploy data marts quickly due to their co-resident status with the warehouse. BI implementers can structure them so that the organization can quickly adapt them to meet changing business needs. The mainframe also offers sophisticated virtualization capabilities that:

- Share processor memory, I/O and network components among multiple operating environments. These capabilities allow administrators to quickly isolate workloads, share resources among workloads and enable communication for workloads internally with an in-memory TCP/IP network.

- Simplify your IT infrastructure by supporting hundreds of concurrent applications with z/VM and supplying management tools for operation, maintenance and accounting.

- Create, provision, deploy and manage virtual servers for optimum performance and flexibility.

In addition to these capabilities, the mainframe has industry-leading development and acquisition technologies to meet the requirements for data security and availability, encryption, auditability of all data acquisition processes, and timeliness of data loads. These ensure that changes to existing or implementation of new BI capabilities will be handled expeditiously.

Myth 10: Releasing a BI solution into production in a mainframe environment is excessively laborious

By the time you read this myth, it must be obvious that the mainframe environment has all the necessary tools, techniques and technologies in place to create a sophisticated and easily sustainable BI environment. Most of the difficulties in releasing BI applications into production come from a lack of a rollout plan with processes that ensure all the moving parts fit together properly and seamlessly.

Production release processes need to reflect the importance and complexity of the application. Many of the steps in the release process, such as the testing and certification of the application, are totally independent of the platform. The robust mainframe environment supports migration from development, to test, to production environments quickly, and once these processes are established, their complexity is due to the rigor required to ensure a problem-free release of the business intelligence application into production – not because of the technological platform being used.

Fortunately System z has many features used in support of its operational applications that also simplify and enhance the release of BI solutions. For example, the mainframe environment has the DB2 Audit Management Expert to help auditors and DBAs alike. Some of its features include the ability to:

- Have centralized auditing tools that bring together information from many sources into a correlated, coherent view of the BI environment.

- Provide a separate administration user interface to allow product administrators to easily define users and groups, assign privileges, define data collection policies and perform other administrative tasks.

- Enable auditors to collect, view, analyze and report on data and save it into an audit repository.

- Automatically generate reports against this audit repository and to export this data into other applications such as Excel for further analysis.

Other production support mechanisms been discussed in the previous sections and include substantial support in the areas of performance, continuous availability, architecture and ease of use.

Summary

Should you consider the mainframe for your BI platform? After reading this paper, dispelling many of the myths surrounding this venerable workhorse, your answer should be a resounding “Yes”; the mainframe is indeed a viable BI environment. IBM and SAS continue to expend significant resources and effort to ensure that the mainframe and its BI capabilities are modern, sophisticated and have all the same capabilities found in alternative platforms.

MYTHS ABOUT MAINFRAME BUSINESS INTELLIGENCE

Given that data volumes and workload processes are growing at amazing rates, it is remarkable that mainframe data center staffing levels have not changed significantly despite these increases. This is a testament to IBM's and SAS' mainframe strategy and philosophy to ensure easy installation, configuration, administration, service and ultimately use for BI analytics.

The System z platform has demonstrable reliability, scalability, security and the excellent mixed workload performance capabilities needed to support traditional strategic and tactical BI as well as near-real-time operational BI. SAS cut its teeth in the mainframe environment and continues to be the leader in innovative analytics and BI applications for that platform. This alliance makes the combined SAS and System z platform worth investigating as a solution for your BI needs.
