Enterprise Computing Advisors

CLABBY ANALYTICS

Advisory

Choosing IBM zEnterprise for Next Gen Business Analytics Applications

Introduction

A half-decade ago, if you randomly asked 10 CIOs to characterize how they use their mainframe computing environment, 8 of the 10 CIOs would have characterized their mainframe use as being "fenced off". Mainframes ran a traditional mix of application workloads, most notably classic (high volume) workloads (often CICS-based); back office jobs (favoring a batch computing reference architecture); and "heavy-lifting" business applications (e.g., complex ERP, SCM application suites that have been highly optimized and integrated with an enterprise class DBMS environment for optimal performance). Meanwhile, departmental and office workloads operated as separate islands of information — walled off from the mainframe.

Now, let's fast forward to the end of 2011. By our estimate (this advisory has been jointly authored by information technology research analysts Brad Day of *Enterprise Computing Advisors* [ECA] and Joe Clabby of *Clabby Analytics*) *40% of the current System z customer base is either piloting or has actually deployed "new workloads" on their System z's*. The mainframe is becoming a large consolidation environment for all types of workloads — including traditional batch/transactional workloads as well as modern Java/Linux and business analytics workloads.

> ***Mainframe customers tell us that there are two major reasons why they are deploying new workloads on IBM's System z. The first reason is "superior economics" (System z costs less to operate — and can cost less to acquire — than a cacophony of underutilized Unix- and x86-based servers). And the second reason is "system design strengths" (that manifest themselves in high quality-of-service levels and lower risk).***

ECA and *Clabby Analytics* have structured this Advisory in two sections. Section One takes a closer look at why information technology (IT) buyers are placing new workloads on System z. (We focus specifically on the impact of running new business analytics workloads on z). We explain what IBM and its independent software vendor (ISV) partners have done to make its System z more attractive to CIOs. We then explore why System z offers superior economics when running various workloads. We follow this discussion by examining some of the technical aspects of System z that make it ideal for running new workloads. Section Two discusses the impact that running business analytics workloads on z can have on an enterprise (deploying on a System z can lead to greatly improved business process flows resulting in more business and/or improved customer service). And we conclude with a recommendation that more CIOs consider hosting certain new workloads on IBM's hyper-efficient System z.

SECTION ONE

Understanding z's Evolution: A Critical First Step

Five years ago, most mainframe enterprise architects would have received a cool reception if they brought a "let's run new workloads on the mainframe" message to their CIOs. Back then, new workloads were being placed on Unix-based systems architectures, or on Linux/Windows

x86 distributed computing environments. But since then, several factors have changed the historical Big Iron biases of CIOs regarding mainframe architecture:

- System z hardware and software design breakthroughs (first triggered by the introduction of Linux on the mainframe, then followed by the debut of the IBM System z10 in 2008 and by the new hybrid zEnterprise system in July, 2010) have enabled the mainframe to be deployed in new ways;
- The growth of the independent software vendor (ISV) ecosystem. Hundreds of ISVs now consider the System z to be a Tier 1 strategic computing platform;
- The optimization of IBM's software stacks to exploit z native hardware/software elements; and most importantly,
- IBM's push to continually improve the economics argument for "Why z?" with a more improved set of pricing metrics, and mechanics (such as "specialty engines"), as well as specialized hardware/software/services packaging (such as "solutions editions") — to significantly reduce the life cycle costs of running a mainframe computing footprint.

Improvements in hardware and software design, the growth of the ISV community, as well as IBM's efforts to deploy its own integrated software stacks on z and drive down System z pricing, have all contributed a change in perception about the mainframe at the CIO level. It is this change in perception, combined with superior economics and a superior systems design that is leading to the deployment of new workloads on System z.

What Makes System z "Uniquely Qualified" to Process Certain New Workloads

As stated previously, the primary factors driving the deployment of new workloads on IBM's System z are related to superior economics as well as system design strengths. In this section, we examine these factors more closely.

Superior Economics

Research conducted by Dr. Howard Rubin, founder of Rubin Worldwide, shows that huge cost savings are possible when deploying large, high QoS transactional and business analytics applications on a System z. For example, some of Dr. Rubin's findings include:

- IT processing cost per transaction savings in the credit card industry can be up to 44%;
- In the consumer loan business, a mainframe can lower IT processing costs by up to 31%;
- In the hospital industry, IT cost per bed can be lowered by 24% using a mainframe;
- Airline processing cost per airline passenger can be lowered by 20%;
- In the oil/gas business, a mainframe can lower the IT cost per barrel of oil by 23%
- In the auto manufacturing industry, IT cost per new vehicle can be lowered by 20%; and,
- In retail, the IT cost per retail store can be lowered by up to 25%.

Further, Dr. Rubin's research illustrates that a System z can help lower IT cost per megawatt hour by up to 25% as compared with distributed systems solutions. (See Figure 1 {next page} for a graphical representation of these savings).

What our own research shows is that IBM is able to offer superior economics on System z by enabling its customers to achieve higher utilization rates than can be achieved on other systems; by improving hardware price/performance; through lower energy costs; and, surprisingly, through lower software costs (software costs are usually tied to the number processors the system runs — and System z uses very fast processors that can be highly utilized, enabling customers to run more work on fewer processors. This results in having to buy fewer software licenses).

Figure 1 — System z "Best Economics" — Typical Savings By Industry

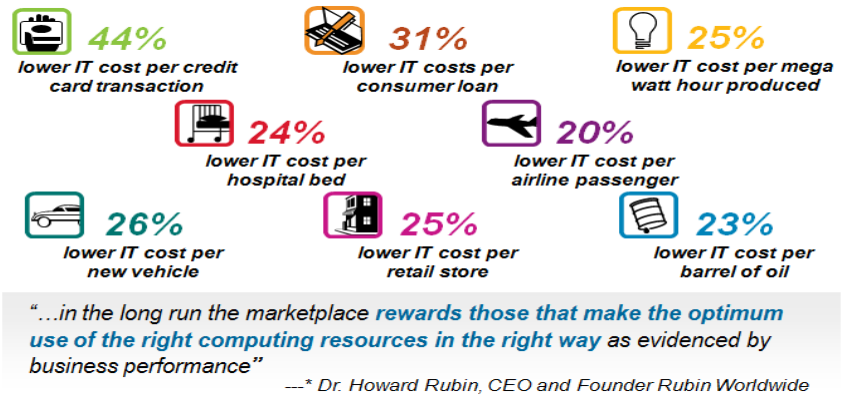

Source: Rubin Worldwide — September, 2011

System Design Advantages

The way we see it, IBM's System z has three distinct design advantages when compared to competing distributed systems architectures. These design advantages can be found in: 1) the System z network/bus/I/O subsystem structure; 2) the System z processor design; and 3) the ability of the System z to recover unused processing cycles and rapidly return them to a virtualized pool (leading to outstanding utilization rates).

A closer look at the System z network/bus/IO subsystem design shows why mainframes communicate more efficiently than distributed systems designs. Distributed systems rely heavily on networks to pass information between servers and to fetch data from databases. And networks can easily become congested — especially during peak workload time. System z, on the other hand, has a huge, high-bandwidth internal bus and a supporting subsystem of specialized input/output processors that enable applications to work efficiently and cooperatively without the network congestion that plagues distributed computing designs. Further, the System z can be tightly-coupled to external storage that supports backend data stores, delivering data at high speeds to the processors performing the work.

Next, consider the System z microprocessor. Microprocessors handle three basic tasks: parallel, serial, and data-intensive computing tasks. What makes the System z stand out is that it is good at all of these tasks. By comparison, Oracle/Sun servers are pretty good at parallel computing tasks, but not as strong in serial and data-intensive situations. x86 processors are also pretty well-balanced, but don't hold a candle to the System z from a scalability and speed perspective (at 5.2 GHz — the z microprocessor is the fastest in the industry). Processing speed is important since a fast processor can do more work than a slow one — and the ability to do more work lowers costs because it forestalls the need to purchase additional servers and software licenses.

In addition to looking at the System z main processors, we also looked at the characteristics of IBM's zIIP specialty processors (IBM builds several "specialty processors" for System z that are optimized to perform certain tasks). It should be noted that the zIIP processor is the same microprocessor used for general purpose computing (described above) — but it is driven by different microcode. This microcode has been designed to expeditiously serve IBM's DB2 database, delivering *exponential increases in database and business analytics performance*, as compared with using general-purpose processors.

System z also has a huge system design advantage over distributed computing architecture when it comes to utilization rates. What our research shows is that mainframes offer far better utilization rates than midrange and low-end servers, primarily because the mainframe system design does a far better job of returning resources to a resource pool where they can be used by other applications that need processing power.

Our point is that System z offers enterprises better value than servers that can only be operated at a 40-60% utilization rate, as mainframes can consistently operate at a 100% utilization rate.

From a systems design perspective, it is important to note that an IBM System z can be deployed as a hybrid computing environment that is capable of integrating workloads that span z/OS, Linux on System z, z/VM, AIX on POWER, Linux on System x (IBM's x86 server line) — as well as Microsoft Windows operating environments in a single unified system. What is particularly impressive about this environment is that a mainframe and attached blades within the zEnterprise BladeCenter Extension can all be managed from a single point of control using IBM Unified Resource Manager. (A centralized approach to management is far more effective and efficient as compared with having to manage activities on a bevy of distributed servers).

SECTION TWO

Market Assessment: The Role of Business Analytics — and About How to Deploy BA Applications

In the previous sections we examined customer environments where System z offers superior economics as compared with distributed systems environments. We also examined some of the unique design characteristics of System z that give it distinct performance and cost over a distributed systems approach. In this section, and through the remainder of this report, we switch gears and take a closer look at the implications of running a new workload on System z. We have chosen to focus specifically on business analytics in order to demonstrate this impact.

First, we must note that as we examine customer scenarios in the business analytics world, what we have found is that most enterprises seem to be making big mistakes in the way that they view the role of business analytics, and in how they deploy business analytics systems. Many enterprises view business analytics as a point solution that can be used by a marketing department to get a better sense of customer requirements, or by financial analysts seeking to analyze business performance. Business executives at these enterprises have been unable to take full advantage of the massive amounts of data being captured by their businesses. Marketing and finance departments should not be the only departments that can benefit by analyzing that data. Further, these executives have not maximized the positive effect business analytics can have on streamlining business process flows.

As for business analytics deployments, the problem we see is that many enterprises are deploying business analytics solutions on a departmental basis. Data is becoming siloed in departmental databases, where it cannot be accessed by the extended enterprise that could benefit from the insights in that data. Further, departments are deploying their own servers to analyze their own copy of data — an approach that we see as costly and inefficient. Finally, it has been estimated that only 8% of employees who work for companies that have a traditional data warehousing environment can access the information within that warehouse (think of all the business insights that are lost by limiting/restricting access to warehoused data).

The Relationship of Business Analytics and Business Process Flow

One of the first lessons taught in business school is that process flow efficiency lowers SG&A (sales/general/administrative) costs — and that the resulting cost savings pass directly to the enterprise bottom line in terms of profitability. One conclusion that can be drawn from this point is that enterprises that drive efficient process flows can be more profitable than enterprises that use manually intensive or inefficient processes.

Business analytics helps executives and analysts make more informed decisions. These more informed decisions can lead to improved customer care; better customer insight; improved visibility into the root causes of problems; improved operational efficiency and performance; the ability to manage risk better — and can provide a means to more effectively streamline business process flows.

As examples of this concept, consider how business analytics is influencing financial analysis, call center, sales, and strategic investment activities:

- ***Deep business analysis*** — using business analytics software, coupled with statistical/performance management products, analysts can spot trends and anomalies in business data. These tools enable analysts to conduct "what if" analysis and perform predictive modeling that enables analysts to predict potential threats and identify new opportunities. This software combination also provides business executives with greater visibility into operations and corporate financial activities — helping C-class executives more effectively lead their companies.
- ***Call Centers*** — to be most effective, call centers need to be able to access real-time information, and they must have an understanding of the customer that they are serving. To better meet customer needs, call center managers must make quick decisions based on the information being presented by a customer, and the information available in that manager's information systems. Business analytics software helps filter information, simplifying decisions for a call center manager. Additionally, information from customer interactions (such as call volume statistics, actual call counts, number of calls completed, transferred, abandoned, or dropped) can be captured and analyzed, providing enterprise managers insights into their call center performance.
- ***Predictive cross-sell and up-sell performance management*** — predictive cross-sell/up-sell analytics tools can be used to help analysts design solutions that will keep customers coming back, and spending more. Insights from historical information can be analyzed to develop predictive models of customer's buying behaviors. With this data, workers in a call center or sales clerks operating point-of-sale devices could be made aware of products that a customer might want to buy — and capitalize on an opportunity to upsell products to those customers.
- ***Strategic business management*** — IBM Cognos software can be used to build linkage between business analytics applications and business outcomes (and can also enable business analytics to be integrated into core operational processes. For example, Cognos offers the ability to enable mobile notifications when business actions need to be taken. This push of information creates a partnership between the information systems and key decision makers. The recent delivery of Cognos on the z/OS platform enables tighter integration of business analytics within the operational processes that drive business. IBM's SPSS enhances the analytic capabilities on the System z platform with predictive analytics, data and text mining capabilities. Together, these business analysis tools help enterprises easily define and model their strategic objectives, and enables business information to be quickly passed to those individuals that impact business results.
- ***Strategic investment management*** — one of the best examples of the linkage between business analytics and process flow, can be found in some advanced offerings available from IBM's Cognos organization. IBM's Cognos Planning and Cognos Business

Intelligence products automate and structure the entire business planning process and builds a linkage to strategic investment analysis. By using the Cognos Strategic Investment Management Blueprint, an enterprise finance department can improve its ability to fund and manage profit and growth programs, while ensuring they are aligned with strategic business objectives. These business analysis tools help enterprises easily define and model their strategic objectives. This enables financial targets to be quickly passed to those assigned to execute the related funding initiatives.

What is particularly interesting about this analytics-based blueprint is that: 1) it features heavy use of business analytics/modeling tools; and, 2) it is highly efficient from a process and management point of view. This kind of turnkey analytics solution can be used by enterprises to build plans that support customer objectives. These tools save time and effort (because business analytics software works in tandem with workflow software). As a result, the strategic planning process can be greatly compressed and better managed —enabling an enterprise to more accurately build and execute a strategic investment plan, and then monitor that plan as it is executed.

NOTE: In each of the above cases, business value is derived by pulling data from a variety of data sources and transforming it into a strategic asset. The strategic assets can be used to streamline performance, improve customer service, create competitive differentiation, and drive new growth opportunities.

The Deployment Situation: Siloed/Distributed vs. Centralized/Scale-up

Business analytics software enables business executives to quickly and thoroughly analyze data — leading to improved customer service, better governance, improved asset management, and increased operational efficiency. But to date, most business analytics activity has taken place at the departmental level where data is funneled into *distributed departmental silos* and used by a *select few analysts* to evaluate supply chain activities, to delve into financial models, or to review and predict customer buying behaviors.

Both ECA and Clabby Analytics believe that this siloed approach to business analytics needs to change. Enterprises would be better served by placing data into a common repository where multiple decision makers, suppliers, and customers can access that data through a secure server to address their specific analytics requirements.

By using this "centralized" data management approach, decision makers throughout an enterprise could use business analytics software to develop new insights, more efficiently predict customer behaviors, better analyze supply chain activities, and more.

But building a centralized database is only part of the solution. Business executives who want to enable dozens, hundreds, or thousands of analysts, managers, and employees across an enterprise to analyze large databases, would be best served by using a* scale-up architecture *as opposed to numerous distributed servers, as is usually the case today.

In our view, a scale-up architecture offers the best economics for large scale business analytics workloads for the following reasons:

1. Scale-up servers offer faster access to data due to large internal busses and adapters that can operate at channel speeds;
2. Better security due to fewer access points;
3. Significantly lower management costs as compared with having to manage and secure dozens of smaller servers; and,
4. Richer qualities of service (QoS) as compared with distributed computing approaches.

A Real World Example: IBM's Blue Insights

One of the industry's best examples of a scale-up business analytics environment can be found at IBM. Like many other enterprises, IBM had islands of data that was accessible by relatively few researchers and analysts. IBM decided to create one large centralized business intelligence solution that was highly scalable for its distributed workforce. Its goal was to create an environment that would provide access to data for thousands of people in order to capitalize on a broader set of insights.

To build this environment, IBM chose to deploy its database on an IBM System z using the IBM Cognos 8 BI software. The actual system design is a private cloud which the company calls "Blue Insight".

A closer look at Blue Insight finds that today, over 200,000 people have access to IBM databases for business analytics purposes. Users can allocate their own virtual servers and run jobs — and when complete, these servers are returned to a virtual resource pool. This is possible because IBM has simplified the access to its Blue Insight cloud, and it gives users simple controls for setting up and running tasks. This means that IBM no longer has to deploy and manage distributed servers at the departmental level, and this reduces the cost and complexity of running a business analytics environment. Further, data is no longer trapped in silos and departmental databases, but rather is centralized in an easy-to-control environment, where it is accessible by its global workforce. An additional benefit of this environment is that the data is secured at the row level, and users see only what they are authorized to see.

> ***IBM says that it has reduced the cost of deploying business intelligence/business analytics solutions by a whopping 70% using this approach. Further, IBM also points out that it captures data from over 100 different data sources — and this data is highly valuable for real-time decision-making.***

Creating Competitive Advantage Through Integration and Packaging

In previous sections, we discussed how a business analytics workload can exploit z microprocessors, z system design, and z QoS features to achieve performance levels that are exponentially better than competing server architectures. But, to further extend its advantage, IBM has created packaged offerings of integrated business analytics and database software, with underlying System z hardware. These include IBM's Smart Analytics System 9700 and IBM DB2 Analytics Accelerator for z/OS V2.1.

A closer look at IBM's Smart Analytics System shows a turnkey, ready-to-use, deeply integrated and optimized business analytics environment. This environment has been designed to support cubing services, data mining, text analytics, intuitive business intelligence reporting, and other analytic functions. Further, it has high-performance data warehouse management functions that have been well integrated with accompanying storage hardware. And because it uses an IBM System z, this environment is highly reliable.

> ***What is especially interesting about this environment is that DB2 and z/OS (the mainframe operating environment) were built for each other. Accordingly, the DB2 database can take full advantage of the System z hardware components and related instruction set — and it can exploit IBM's zIIP processor microcode instructions, and hardware compression facilities to deliver a high-performance business analytics environment.***

IBM's new version of its DB2 Analytics Accelerator (the IBM DB2 Analytics Accelerator for z/OS, V2.1) is a high-performance "analytics accelerator appliance" that delivers dramatically faster complex business analysis than many competitive systems. Using this DB2 for z/OS "accelerator" approach, users can exploit the quality of service features available on System z (security and high availability), while benefiting from the high performance of this massively parallel system environment. The end result is that this system transparently speeds complex analysis within the secure environment of DB2 for z/OS.

> ***What we found to be especially interesting about this packaged solution is that it is a quick- to-deploy server environment that requires few, if any, changes to existing business analytics applications. Simply deploy it, and substantial increases in performance can be instantly recognized by DB2 for z/OS queries.***

Both of these packaged environments leverage the rich RAS (reliability, availability, security) of the IBM System z platform. Both offerings are scalable, and both offer superior performance as compared with competitive offerings.

Summary Observations

In this *Advisory*, we examined customer claims that IBM's System z offers superior economics for some workloads, as well as claims that the mainframe systems design is superior to distributed computing designs. And we found both of these claims to be true.

Still, with all this evidence that mainframes provide better economics and a superior systems design when compared to distributed systems offerings, there are still those who fail to evaluate mainframe solutions. One of the main reasons is a belief that mainframes cost too much. For IT buyers who still believe that mainframes are too costly, we note (with emphasis) that IBM has driven down its pricing by creating specialized business analytics servers (IBM's Smart Analytics System 9700 and IBM's DB2 Analytics Accelerator for z/OS) — and by creating other packaged solutions for other types of applications (see *Clabby Analytics'* Solution Editions report at http://www.clabbyanalytics.com/uploads/SolutionEditionsFinal.pdf for more details).

We spent a lot of time examining the impact of a properly designed business analytics environment on a given enterprise. It remains our belief that enterprises that take a centralized approach to data management — and then make it possible for hundreds or thousands of users to access that data — are better served in terms of innovation and growth, as compared to enterprises that run siloed business analytics environments.

In the end, it is important to run the right workload on the right system in order to reduce costs and maximize return on investment. And for large business analytics workloads that can support large populations of users, the right systems choice is IBM's System z.

Clabby Analytics
http://www.clabbyanalytics.com
Telephone: 001 (207) 846-6662