



The Mainframe: It's Baaack!

RFG believes the dynamics of infrastructure ownership are rapidly changing; central IT groups have, or will soon have, ownership of and responsibility for the entire underlying infrastructure in data centers. This watershed event is being driven by the desire for consolidation, cost savings, power and cooling constraints, risk management, and virtualization. This will allow IT to be able to maximize utilization of hardware resources, and to devise a cohesive hosting strategy that leverages existing investments. One result of this shift has been, and will continue to be the re-emergence of the mainframe as a platform of choice. Indeed, a recent study regarding database management system (DBMS) trends indicates that 25 percent of respondents were moving data to mainframe-based DB2. IT executives and infrastructure architects need to re-evaluate their current assumptions in light of the new dynamics, and then fold the new conclusions into the architecture and plans for application development, operations, and data center constructs.

Business Imperatives:

- The value proposition for the mainframe in large enterprises has been greatly underestimated. Moreover, it will be improving over the next few years. Some of the value components that provide significant leverage include data sharing, disaster recovery, Linux, partitioning, power and cooling, server pool workload balancing and prioritization, security, specialized processors, total system utilization, and virtualization. Before IT executives and architects decide to offload applications from the mainframe or select platforms for new applications, they should have a complete understanding of the mainframe value proposition and total cost of ownership (TCO) components. Three strategically important roles that IT executives should consider for the mainframe are information hub (data serving), security hub, and server pool master control and management hub.
- The new wave of consolidation, partitioning, and virtualization occurring on Unix and x86-architected platforms is an affirmation to the mainframe model, and a move away from the distributed one-application-per-server mentality that has driven up the cost of computing. The multi-purpose role of mainframes, with its inherent superiority in partitioning and virtualization, provides the same options now being considered on non-mainframe platforms, but with stronger and more proven technology. IT executives should consider the value of the 20-year advantage that mainframes have over other platforms in this area. RFG believes the mainframe technology investments that IBM Corp. has made (and will continue to make) will ensure technology superiority for the foreseeable future. With this in addition to the power and cooling advantages of the mainframe, IT executives and staff should consider the mainframe as a "first tier" option for hosting new applications and acting as a central hub for security, server pool management, and consolidated workloads/data.
- Service-oriented architecture (SOA), standardization, use of wrappers, and other reuse techniques present new opportunities to mine the value of legacy applications on the mainframe. For many enterprises, SOA will help unleash the millions of lines of legacy code that are optimized for the mainframe for reuse by other applications. In addition, the expansion of Java and other languages that are platform-independent will help address the development issues that have hurt the mainframe in the past. IT executives and architects should consider mainframes in the context of "black box" application servers that can satisfy most application requirements, including availability, reliability, and scalability.

RFG views 2006 as a watershed year for mainframes, in that the assumptions of infrastructure ownership and preferred platform characteristics and requirements are permanently shifting. Virtualization has become a mainstream concept, and is being implemented in x86-based servers as well as Unix and System i (previously known as iSeries) servers. **VMware**, and **EMC Corp.** company, has become the *de facto* standard for virtualization on x86-based servers, and is driving consolidation of applications onto blade



server platforms. This shift has resulted in the loosening of control of infrastructure ownership by lines of business (LOBs).

However, the power and cooling issue is now forcing IT executives to rethink the decision to virtualize on blade servers. While technologies are maturing to mitigate this issue (e.g., **Hewlett-Packard Co. (HP)**'s recent c-class server announcements), RFG has worked with enterprises that are consolidating to the mainframe, with its strong virtualization technology. At one such firm, they consolidated 600 **Intel Corp.** servers running Linux onto a single mainframe. This shift has eliminated one floor of servers, cut power and cooling costs by 80 percent, and reduced the administrative staff (and costs) by 10 administrators. No other platform has clearly demonstrated the vertical scaling (scale-up) savings as the mainframe. It can consistently achieve the highest utilization rates, and hardware/software costs go down as the processing capacity increases.

The one problem non-mainframe clients have in performing this TCO evaluation is the inability to compare platforms meaningfully. MIPS (the traditional mainframe processor speed metric of millions of instructions per second) versus megahertz (MHz) is meaningless, and the additional capabilities of the mainframe (e.g., inboard channel directors and sophisticated workload managers) have no counterpart on other platforms. When all the factors are considered in a TCO evaluation, the mainframe should emerge as the least cost solution when being compared against 100 or more Unix or Linux servers. RFG expects to see a number of these types of infrastructure swings over the next several years, as architects and IT executives study the cost/benefit analyses of the server platform alternatives.

The Information Hub

Another element driving this mainframe movement is the need for a "single version of the truth" of corporate data. Today, enterprises have a number of different copies of enterprise information that they spend 25 to 50 percent of the time copying and syncing. (See the RFG Research Note "[Data Architecture: The Key to Information Services](#).") Nonetheless, reports and information presented to business users may be inconsistent and inaccurate.

This master data management (MDM) issue has percolated to the top, and is now on the radar of most IT and business executives. Many executives and architects recognize that this problem and phenomenon is the result of moving from centralized databases, historically on mainframes, to distributed databases that are endemic to distributed computing. They also realize that this can be reversed by consolidating data and constructing an MDM environment. Enterprises will be able to solve this through the creation of an information hub (or information server) where data sharing is optimized or a grid structure that allows access to the data from any server. There is only one platform today that allows for unlimited data sharing with the levels of protection and security required – and that platform is the mainframe. IBM has been retooling the mainframe over the last five years to make this environment one of its primary roles.

RFG expects to see enterprise architects pushing for corporate information to be managed as a common, corporately-owned data pool (similar to the concept of a network cloud), and that the gateway to that pool will be through a processor optimized for that role. Ultimately, SOA implies an architecture that will create both application and data independence, hindered for years by the local ownership of data and traditional coding methods. General acceptance of this view will not come easily, as LOB architects seek to keep control of their data. The latter view sub-optimizes the value of information for the enterprise, which results in multiple copies of data and the loss of the holistic knowledge of business partners, customers, employees, financial impacts, inventories, products, relationships, suppliers, and vendors. While it is likely that companies built around a federated model of distributed businesses are likely to maintain a business-



unit-centric view of data, RFG strongly recommends that business and IT executives support the move to information as a corporate asset.

Specialty Engines

Another new trend that is emerging among large organizations is the convergence of the desire to run Linux applications on a mainframe and the associated need to simplify disaster recovery. This is not just a process issue, but is also a risk management and Sarbanes-Oxley (SOX) controls compliance requirement. Major corporations are finding that the ability to deal with thousands of distributed servers from a controls and disaster recovery standpoint is pushing the limits. These companies are finding that if they take their Linux applications and operate them on a mainframe they reduce their TCO, simplify their disaster recovery effort, and trim down their residual risks to manageable levels.

IBM has helped to make this movement possible by the introduction of the Integrated Facility for Linux (IFL) processor on the mainframe. This was the first of the specialized engines that IBM has introduced. Versions of IFL specialty processors have been around for a number of years, but the latest incarnation of the IFL came on the market in 2001. It allows for the addition of Linux applications to a mainframe environment without purchasing additional general purpose (GP) processor capacity. Hence the user does not get hit with an increase in software costs, and can expect equal or better performance, reliability, and recoverability. IBM has priced the IFL processor at USD 95,000 for the System z9 Business Class servers and at USD 125,000 for the System z9 Enterprise Class servers.

IBM has also introduced additional specialized processors that are targeted at optimizing the mainframe for two other environments. The System z Integrated Information Processor (zIIP), announced this year, is designed to enhance the role of the mainframe as an enterprise information hub. What IBM did was to determine where the performance bottlenecks were, and eliminate or reduce them through the use of this specially designed processor. Here again, users pay a reduced price for the hardware (versus standard mainframe GP processors) and avoid any added software charges. From a processing standpoint, this is another central processing unit (CPU) that assists in the execution of (primarily DB2) code. These specialty engines further separate the mainframe and its true capacity from alternative platforms, providing capacity that is not counted in the base capacity measurement, but relieves the burden of GP processors. Additionally, the mainframe utilizes over 100 input/output (I/O) channel processors, an encryption processor (if desired), and other specialty processors (also as desired). This parallelism allows for better performance and throughput than other servers that tout multi-processor and multi-threading capabilities.

In the same vein, IBM announced the System z Application Assist Processor (zAAP) back in 2004. It is a specialized processor for Java execution environments. This engine, which is priced the same as the zIIP and IFL processors, is currently priced at USD 95,000 for the System z9 Business Class Servers and USD 125,000 for the Enterprise Class. Companies that are running a significant amount of Java code on the mainframe will find that by offloading these Java applications from the GP processors to the zAAP engine, the cost of computing will be lowered. In addition, MIPS will be freed up for other applications on the general purpose System z processors.

RFG does not expect IBM to increase its line of specialty processors in the next few years. Instead, RFG expects to see further optimizing in these engines and the software that supports them. Moreover, RFG expects the IBM Software Group to increase its focus on the mainframe and optimization of applications in both the general purpose and specialty processors. IBM definitely has and should continue to have a leg up on its competitors in the specialty processor category, as it has been developing, delivering, and integrating them on the mainframe far longer than any of its major competitors on any other platform.

**The Master Control Point and SOA**

One such area that IBM will need to address is workload balancing, optimization, and provisioning. RFG expects that the mainframe will become the enterprise server pool master control point. That is, enterprises should view all the servers as a single processor pool, and allocation of work and provisioning of server resources should be managed from a single control point. This will require enhancements to its virtualization engine and Tivoli integration. These products and their respective databases will need to become more intertwined, or combined, so that the mainframe can handle all the autonomic computing, end-to-end application monitoring, orchestration, prioritizing, provisioning, and workload balancing across all servers in the data center. IT executives and architects should expect IBM to deliver on this vision, and should be building their infrastructure architecture and data centers accordingly.

SOA on the mainframe has already been announced, and this will be another area on which RFG expects IBM to focus. IBM recognizes that its clients, especially those in the financial services industry, have billions of lines of code of mainframe legacy applications that remain the core of their business processing. These applications cannot be changed overnight, but need to be upgraded for use and reuse in an SOA world. To protect this investment, IBM will deliver improvements in SOA functionality for mainframe applications to make them available for reuse and competitive in today's agile environment.

Moreover, it should be noted that when performance is a key application criteria, architects and designers should seriously consider having as many SOA components as possible on a mainframe complex. By having the application and data components on the same processor or in the same sysplex, the networking performance issues can be minimized. (See the RFG Research Note "[Network Management 101 for IT Executives](#).")

Looking Back, Looking Forward

Enterprises moved away from the mainframe in the 1990s because of issues with application backlogs, control, and cost. Now the equation has inverted, and companies are struggling with application backlogs and quality, availability, complexity, control, costs, and security of distributed systems. These challenges are being exacerbated by the new regulatory pressures and the implementation of risk management practices and processes among large enterprises. The new paradigm cannot be ignored and must be addressed.



RFG believes business and IT executives and architects must revisit their infrastructure hypotheses, and adjust them for the shift that is now in the process of occurring. The primary issues of people, performance, power, price, and risk need to be factored into new application development platform selection, data center design, infrastructure architecture, and operations management. Enterprises will continue to employ the gamut of servers from mainframes to PCs. However, IT executives and architects should create a vision that defines the role of mainframes in terms of both overarching "command and control" capabilities (e.g., security, data hubs) as well as application hosting.

RFG analyst Cal Braunstein wrote this Research Note. Interested readers should contact Client Services to arrange further discussion or an interview with Mr. Braunstein.