

IBM's zEnterprise Really Stretches Its Boundaries — New Windows Are Opened

Analyst: Stephen D. Bartlett

Management Summary

You can't read the news on technology without engaging in speculation on what's about to be announced, the significance of what it will mean, who will be affected, and how the marketplace will change. Since the technology is not just institutional but very personal for those of us in IT, we must try to sort out the meaningful from the overlapping waves of bravado and noise.

IBM mainframe sales reps and distributors don't enjoy the pleasures and satisfaction of the whole world's anticipation, speculation, or even the organized illicit intelligence gathering, which seems to be common in Apple's *i-world* these days. There was a time, beyond the memories of most of the IT community, when they did, but the world has changed and new technologies have driven the focus to a broader audience that is less sophisticated in many respects. As a consequence, the "fun stuff" has shifted a good deal from what's the biggest and fastest, toward what is the latest, and usually smaller, innovation – the new gotta' have. Computing, if you will, now is more personal and portable. Does anyone care what all this stuff is connected to and what holds it together anymore? Not so much, apparently.

If a tree falls in the forest and there is no one there to hear it, does it make a sound? Philosophers may argue this question. One might pose an analogous question to IBM's mainframe marketers and sales people; there is anecdotal evidence that the announcements made in the System z family since July of 2010 (namely the *zEnterprise* family of hardware and software products) did not reach as many discerning ears as they hoped or, possibly, they were dismissed as not being relevant. I categorically dismiss the latter argument as false, which I will elaborate on later, but I do believe that IBM now needs to find opportunities to increase their bandwidth a bit and maintain the cadence that they have established. I can suggest at least two reasons why the messaging seems to have fallen short of its goals. The first is that the expansion of the hybridization of the mainframe is dramatic and innovative – a significant branch in the evolution marked by new DNA, if you will, which was completely unexpected and overwhelming. This brings us to the second reason: there seemed to be too much to behold and comprehend. For sure, there were many new elements, perhaps outside the domain of the mainframe's typical audience, which I will review in more detail below. No doubt, some new thinking is required. That is a good thing. Progress requires that we consider new ways to address our problems. In this context, *how does IT continue to support the ever-changing needs of the businesses that they support in a manner that is economically sound, non-disruptive, and open to innovation?*

In particular, *Windows* is coming to the mainframe, not as a native operating environment, but more closely attached and managed than most can imagine. For many and possibly most, this should prompt a reevaluation of your mainframe strategy, even if you don't have one! The future has arrived. Did you notice? Read on, to find out what this really means.

IN THIS ISSUE

➤ Mainframes — Now Slightly Past the Crossroads.....	2
➤ zBX Embraces x86 Windows Applications	3
➤ What Workloads are Appropriate for Deployment on zBX Blades?	5
➤ Conclusion	7

Mainframes — Now Slightly Past the Crossroads

If you are an observer of or active in the IBM mainframe space, it would have been difficult to miss the events of the last 15 months since July 22, 2010, when their newest family, *zEnterprise System*¹, was announced and became the benchmark date for the internal hybridization of enterprise computing. IBM made a substantiated claim to be the industry's only *heterogeneous cloud platform*. If you missed this major architectural extension and have not investigated IBM's new offerings and their implications, then this paper is as much for you as it is for those who consider themselves IT technology leaders.

In a nutshell, IBM's mainframe system architects have successfully implemented an integrated solution that addresses the management challenges resulting from the *sprawl of virtualized heterogeneous servers* (including selection of deployment platform and workload integration), and the many implications thereof. This is not just an issue for CIOs to ponder, but it also has an impact on Line-of-Business (LOB) owners as well, who, frankly, care little about how and where their work gets done – only that it gets done quickly, efficiently, cost-effectively, and in a way that fully supports their mission's objectives.

What are the indicators of the aforementioned sprawl? For starters, application code typically is spread across multiple server platforms, often with different server architectures, and different operating environments. Moreover, often this is accompanied by application data being spread among several storage devices, also frequently of differing architectures. Current (recent) transaction data tends to be separated from historical transaction data (not speaking about archival or backup data here – isolation is a good attribute in that case). As the necessity rises for businesses to have a strong data analytics capability from which they can measure and monitor the attainment of their business goals, the physical and/or chronological isolation of current transaction data from the historical data repository becomes inimical to that process.

As the virtualized images on multi-vendor,

¹ Highlights of that announcement are recapitulated here to establish continuity, but detail can be found in **The Clipper Group Navigator** entitled *The IBM zEnterprise System Reaches Out — Higher, Wider and Deeper*, dated July 22, 2010, and available at <http://www.clipper.com/research/TCG2010034.pdf>.

multi-architecture are perpetuated, *how many management control points are needed to monitor and manage the resources?* **The ideal number is one.** Count the number of links that often are necessary in a multi-tier application, and then include some redundancy to improve the level of availability, now multiply that by the number of applications/workloads that are supported. That number can become very large.

However, a constraint usually imposed on enterprise-class solutions is the recognition that multiple architectures are a practical and useful necessity – the paradigm of *Fit-for-Purpose*². **Diversity usually is a good thing, but presents challenges that must be addressed.** If none of these questions tends to keep you up at night, consider yourself lucky. However, you should make sure that someone is paying attention to these important happenings.

IBM has recognized all of these issues, and more importantly has set the stage for a more complete private cloud solution based on its newest mainframe family. Before I attempt to prove my position, a short recapitulation of the zEnterprise family announcements is in order.

Two Mainframes, a BladeCenter, and a Manager

This story began on July 22, 2010, with the announcement of the zEnterprise System. The mainframe world expected the usual higher performance, expanded capacity, and perhaps some extensions of the zArchitecture. IBM far exceeded these expectations with the *zEnterprise 196 (z196)*, in several respects. It incorporated a new, more integrated chip technology running at clock speeds exceeding all other technologies at that time, 5.2 GHz. Capacity was expanded up to 80 user-accessible engines (out of a total 96) and main memory up to 3TB to yield up to 50,000 MIPS (Millions of Instructions Per Second, which is a measure of system capacity, stated in *zArchitecture* terms)³. Hybrid computing capabilities (the ability to run multiple operating systems and mixed applications on a single platform) are not unfamiliar to the mainframe world,

² IBM often speaks about “Fit-for-Purpose”. It means running an application set on the right (best) server platform(s) and operating environment(s) and storing data on the right storage devices. It is a philosophy of best practices and right choices, both technically (in terms of effectiveness in delivering the solution to users) and efficiency (in terms of cost minimization).

³ IBM never officially quotes MIPS for its processors. The MIPS stated here are those commonly agreed on by industry watchers and analysts, including Clipper.

given that *Linux* on *System z* (known as *zLinux*) has been available for five generations of IBM mainframes.

Completely unexpected was the expansion of mainframe hardware and systems management technologies, which enabled the z196 to directly connect to non-mainframe servers, in particular *Power Systems* and *System x* blades⁴ via the new *zEnterprise BladeCenter Extension*, or *zBX*, connected by two private high-bandwidth networks, and managed by the *zEnterprise Unified Resource Manager*. This will enable many non-mainframe workloads (applications)⁵ commonly running in modern data centers today to be managed from a single focal point with the same control paradigm as the mainframe, the z196, and for the entry and midmarket, the *zEnterprise z114* (*z114*). The first customer shipment of the z196 was September 2010. The first customer ship of a zBX with Power System blades was in November 2010.

Phase 2 – z114 Lowers the Entry Point for Mainframe Hybridization

The z114⁶, which became generally available on September 9, 2011, uses a Single-chip Module (SCM) – whereas the z196 uses a Multi-chip Module (MCM), and the z114 runs at about two-thirds of the z196 clocking rate. It is packaged as a single frame air-cooled system, like the *z10 BC*, instead of the two-frame design required for the much more powerful z196.⁷

IBM attacked the problem of lowering the z114's cost by introducing a more modular two-drawer design resulting in two models, the *M05* and the *M10*, compared to only one model on the z10 BC. The M05 includes a maximum of five

customer-configurable central processor (CP) engines (standard processors running traditional mainframe operating systems, like *z/OS*), plus additional engines configured as *IFLs*, *zIIPs*, *zAAPs*, *ICFs*, or *SAPs*.⁸ The smaller model has the same granularity of CP capacities as does the larger model, but is constrained by fewer engines, less memory, etc., and, importantly, comes at a lower price to the customer with the side benefit of reduced energy requirements. The second drawer contains two additional SCMs, added memory, and additional components for the M10. The larger model includes up to 10 customer-configurable processor engines (cores), with the limitation that – for both models – the maximum number of CPs is five. Therefore, both the M05 and the M10 are rated identically in terms of maximum traditional mainframe performance, estimated to be in excess of 3,100 MIPS⁹.

zBX Embraces x86 Windows Applications

When the *zEnterprise System* was announced in July 2010, we were told to expect the *System x* blades to be supported by *Linux* in 1H2011. Microsoft's *Windows* was not mentioned; some speculated that it was because developing and maintaining the level of control specified for inclusion in a zBX was outside of IBM's domain. It turns out it wasn't; IBM just believed it was less important to its enterprise customers. However, the issue, which some considered a major deficiency (or a lost opportunity), would not go away and lingered long in the blogosphere. The customers and pundits talked and IBM listened, so that this past April, IBM revised its initial announcement by stating its

⁴ Power Systems, formerly *pSeries*, is IBM's industry-leading UNIX offering with its AIX operating system. *System x* (formerly *xSeries*) is based on the x64 and earlier x86 architectures, most commonly running *Linux* or one of the Microsoft *Windows* server operating environments.

⁵ What workloads are best suited for this new environment and how many will there be? This is both an economic question (TCO) and an application one (addressing the affinity of data and processes and the need to manage them all better). A more complete discussion begins on Page 5.

⁶ For more detailed information about the *zEnterprise 114* announcement and comparisons to the *z10 BC*, which it supersedes, see **The Clipper Group Navigator** dated July 12, 2011, entitled *IBM zEnterprise in the Midmarket - Revolution or Evolution?*, which is available at <http://www.clipper.com/research/TCG2011024.pdf>.

⁷ Carried into the z114 from z196 are superscaler design, improved cache structure, new "out of order" instruction execution, and over 100 new hardware instructions, as well as other design and security enhancements.

⁸ In the Mainframe world, an *engine* is a *processor core*, which commonly are called *cores* in the x86 and RISC world. A standard engine, or CP, is one that runs IBM's mainframe operating systems, *z/OS*, *z/VM.*, *z/VSE*, etc., without restrictions. There are also specialty engines: the aforementioned IFL, the zAAP for JAVA offload (z Application Assist Processor), and the zIIP for database acceleration (z Integration Information Processor). An ICF (Integrated Communications Facility) is also available to enable intersystem communications among separated mainframe systems. SAP (System Assist Processors) offload and manage I/O operations. The two models of the z114 and the five models of the z196 provide a great deal of flexibility of engine mix. Some restrictions on combinations may apply.

⁹ The z114 is the first midrange *System z* ever offered over 3000 MIPS. With the additional capacity of optional IFLs, zAAPs and zIIPs, the z114 can be even more powerful. Customers also have the option of up to 130 sub-capacity settings ranging from 28 to over 3,100 MIPS.

intention, as a statement of direction (SOD)¹⁰, to defer the System x blades with Linux by one quarter to 3Q and to support Windows server operating environments on the System x blades in 4Q2011.

The addition of the z114 to the z196 within the zEnterprise System family opens up a whole range of new and exciting potential for the mainframe hybrid solution. This complements the existing portfolio of POWER7 blades (running AIX applications), System x Blades (running Linux applications), the *DataPower XI50z* for zEnterprise appliance¹¹, and the *IBM Smart Analytics Optimizer solutions*¹² – which may be mixed and matched in the zEnterprise BladeCenter Extension, the zBX.

Now, IBM has announced the support of Microsoft Windows on the System x blades on the zBX, specifically the *BladeCenter HX5 7873* dual-socket 16-core blade.¹³ Virtualization for zBX's System x blades will be provided by an integrated hypervisor (based on Red Hat's *KVM*) that is delivered via the *Unified Resource Manager*, as detailed in Exhibit 1, in the next column.

The supported configurations will include SSDs, Ethernet connection, and Fibre Channel adapters with either 64 or 128 GB of memory. There is no price premium attached to the System x blade configurations as a result of being certified for installation in the zBX and the new blade is to be ordered and fulfilled through IBM's System x distributor channel. As the HX5 blade family goes, the supported configurations are fairly modest and at the lower end of middle capacity for the System x product family. When System x blades for zBX were announced last month (simultaneously with the z114), only

¹⁰ SODs are usually promulgated by IBM to inform its customers of a product or technology enhancement that is beyond an announcement horizon, typically about 1 year, and thus is not committed.

¹¹ Specialty hardware that provides secure connectivity, gateway functions, data transformation, protocol bridging, and intelligent load distribution at lower cost.

¹² An appliance-like add-on that accelerates select queries to a degree that enables the integration of business intelligence into operational (real-time) processes.

¹³ The HX5 is IBM's fifth generation of the System x, providing optimal performance for database and virtualized applications. The HX5 7873 includes the *Intel Xeon E7-2830* 2-socket, 8-core processors running at 2.13 GHz.

More detail on this new family of processors can be found in **The Clipper Group Navigator** entitled *Redefining the High-End Server — IBM Upgrades System x for Large Workloads*, dated May 20, 2011, and available at <http://www.clipper.com/research/TCG2011019.pdf>.

Exhibit 1 — Red Hat's KVM to Provide Virtualization for zBX System x Blades

On May 3, 2011, IBM and Red Hat, Inc. announced an agreement to develop products and solutions jointly, based on Red Hat's *KVM (Kernel-base Virtual Machine)*. This was followed on May 17 by the announcement of the formation of the Open Virtualization Alliance, a consortium of BMC Software, Eucalyptus, Hewlett-Packard, IBM, Intel, Red Hat, and SUSE, to foster the adoption of open virtualization technologies including KVM.

KVM was selected by IBM as the hypervisor that will provide the levels of virtualization for zBX System x blades running Linux or Windows images (guests). The choice was based on the capability of KVM to provide the highest levels of performance, scalability, security, and low unit costs. Moreover, it provides the degree of openness and flexibility that most users desire within their infrastructures. Much of IBM's and Red Hat's cooperative developments are to be focused on providing the unique requirements for interconnectivity and interfaces to the Unified Resource Manager. Installation of KVM, as well as *PowerVM*, Power architecture's virtualization software, for POWER7 blades is provided by the Unified Resource Manager.

Sources: IBM and Red Hat

Linux would run on these blades, in particular, Red Hat *RHEL 5.5* (and later releases) and SUSE *SLES 11 SP1* (and, later, *SLES 10 SP4*).

As shown in Exhibit 2 at the top of the next page, the zBX can have up to four racks and up to 112 single-width blades. However, when a mix of blades is desired, which may often be the case, maximums for each type are limited. With 112 blades configured across four 42U racks, this is a very conservative design driven for the most part to assure the highest quality of serviceability and maintenance to which System z customers have become accustomed. It also accepts the fact that, with its current configuration limits, the hybrid mainframe cannot take over the whole data center, unless, of course, it is modest in scope.

Exhibit 2 — The zEnterprise BladeCenter Extension (zBX)

- One to four – 42U racks – for a total capacity for up to 112 blades
- Up to 112 PS701 Power blades
- Up to 28 HX5 System x blades
- Up to 28 DataPower XI50z blades (double-wide)
- Up to 56 IBM Smart Analytics Optimizer blades

Source: IBM

What Workloads are Appropriate for Deployment on zBX Blades?

Applications with Data Affinity

First, you need to know that *any* application certified to run on POWER7 with AIX or Linux or Windows on System x blades will run on these blades when installed in the zBX. Primary consideration for deciding what applications to put on zBX should give priority to any part of your application portfolio that requires reference to a centralized data server hosting *DB2*, *IMS*, or *VSAM* databases or databases running on an IFL, most frequently *Oracle* databases. We call this rationale *Data Affinity*, as it is the closeness of these zBX applications to mainframe data that gives justification for deployment on zBX. Examples would be an *SAP* application server, which could be running on zBX POWER7 blades or a *WebSphere* application server front-end to a core *CICS* or *IMS* transaction processor.

Let's take a look at an online retail shopping transaction as an example. A simple mouse click from a shopper places significant demands on the IT infrastructure. When the shopper adds an item to their cart, a number of activities are initiated. Recognizing the shopper's history and buying preferences is the province of a data mining application, which then makes use of this data available *during* the shopping session. Then the shopper's billing and shipping preferences need to be retrieved and warehouse picking tickets must be dispatched to the correct location. In this example, the Web portion of the transaction may take place on a System x server, running Linux or Windows, the data mining application may run on a POWER7 blade or the Smart Analytics Optimizer, and the inventory and billing portions may take place on a System z running *DB2* for *z/OS*.

Applications with Management Affinity

By integrating the three server hardware platforms into one closely-linked system that is centrally managed by the Unified Resource Manager¹⁴, a great deal of complexity and risk has been removed from the environment. Security is increased, and the number of switches, cables, and adapters has been reduced. Over a large portfolio, reductions of this nature could be an order of magnitude lower or more. We call this *Management Affinity*, as one rationale for using zBX is a more tightly managed application across the different tiers (applications running in different operating environments), in this case facilitated by the singular Unified Resource Manager, as described below.

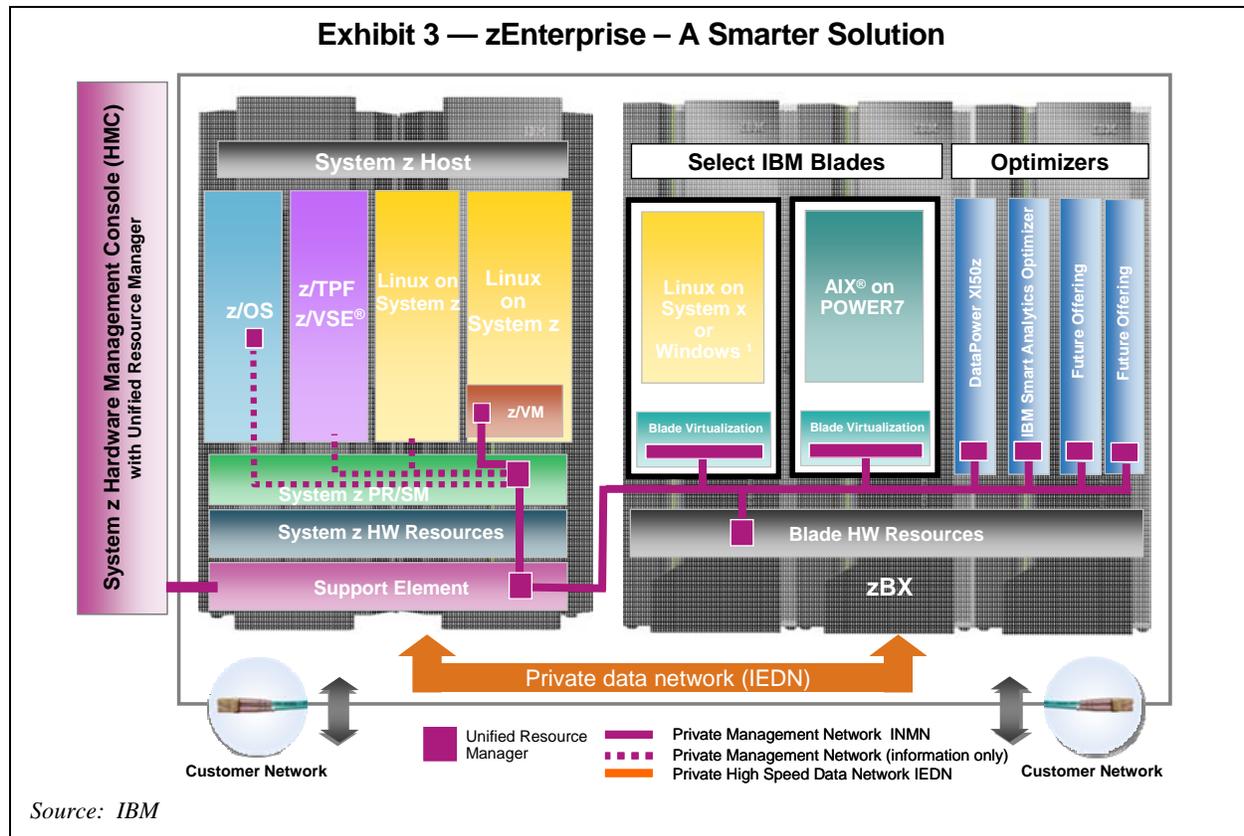
Unified Resource Manager

The third component of the IBM zEnterprise System (beyond z196/z114 and zBX) is the zEnterprise Unified Resource Manager; this is the glue, the software management piece, in the form of imbedded firmware, that unifies the processing resources of the z196 or z114 and the zBX through its hardware, platform, and service management functions. (See Exhibit 3, on the next page.)

The Unified Resource Manager is firmware distributed as a no-charge feature of the z196 or z114. It provides workload awareness to optimize system resources in accordance with the policies assigned to a particular workload. It is a rational and complete central management structure that is unique in several ways.

- It provides total systems management across heterogeneous resources, System z, POWER7, and System x.
- It integrates through a single point of control, common skills for resources, and reduced complexity of day-to-day operations.
- It monitors via a new dashboard for CPU resources and energy management.
- It simplifies installation by auto-discovery and configuration of resources and workloads through a single interface.
- It improves network security with lower latency, fewer hops, and less complexity.

¹⁴ In our earlier bulletins on zBX (and in IBM disclosures), *zManager* was the name used as a synonym for Unified Resource Manager. More recently, that term (*zManager*) is no longer being used, because more is being managed than work on System z.



- It improves control of access through the management of hypervisors, KVM and *PowerVM*, as firmware.
- It extends service and support management by enabling virtual machines and blades to perform hardware problem detection, reporting, and a call home feature to alert IT administrators and dispatch IBM service.

The Unified Resource Manager Now Extended with New APIs

Now, IBM has announced that they are extending the functions of the Unified Resource Manager in several ways. Dynamic discovery and configuration of resources will be extended beyond virtual servers and networks to include, in the future, storage resources as well. No changes to customer network or storage architecture are required when distributed server workloads are migrated to zBX blades, but these resource management capabilities could very well lead to improved data center efficiency, reduced redundancy, etc., thus leading to an improved ROI. Programmatic access to the Unified Resource Manager has been enhanced with new APIs. The announcement also includes an open-documented interface that can enable discovery, monitoring, and management of zEnterprise

resources using external tools written in common scripting languages, e.g., *Perl* and *Python*. IBM Tivoli has stated its intention to enhance the *Tivoli Integrated Service Management System for System z* using these APIs. Other Independent Software Vendors (ISVs) have expressed interest in taking advantage of these APIs, as well, according to IBM.

Managing Heterogeneous Environments on zEnterprise

A primary objective of zEnterprise is heterogeneous platform integration. *IBM Systems Director* is a valuable tool for managing the hardware resources for IBM Power Systems and System x servers, but is not a tool intended to manage System z hardware. For System z, these hardware and firmware management functions are provided through integrated System z firmware, and hosted in specialized System z management appliances known as the *System z Support Element (SE)* and the *Hardware Management Console (HMC)*.

The Unified Resource Manager consists of a set of platform management functions that extend the traditional System z hardware management appliances, delivering the highest possible levels of management integration across all elements of

the zEnterprise system, including the POWER7 and IBM x86 blades. Integration with existing System z management firmware leverages the maturity and resiliency of these appliances, preserves and extends the customer investment in the System z management infrastructure, and protects the customer investment in the management practices, automation, and operational controls built upon the use of the System z management tools. Where Unified Resource Manager stops is below the level of operating system management, so it doesn't deploy, configure, or manage operating systems in the virtual containers it creates.

IBM System Director and Unified Resource Manager both provide function at the hardware and platform management level. IBM System Director remains the management tool-of-choice for stand-alone IBM Power Systems and System x platforms. The zEnterprise Unified Resource Manager provides an integrated solution for the zEnterprise system. And, as advanced platform management functions are introduced through IBM System Director and *VMControl*, they will logically extend the management capabilities of the zEnterprise system, and provide consistent higher-level platform management capabilities across all members of the IBM Systems family.

The Power of the Private Network

Connectivity among the virtual server elements of a zEnterprise ensemble is provided through two high-speed networks, the *Intranode Management Network (INMN)* and the *Intra-ensemble Data Network (IEDN)*. The INMN uses 1000BASE-T (1 Gb/second) technology to interconnect the System z host *Hardware Management Console (HMC)* and *Support Element* with all the virtual servers (in the zBX), in order to provide the necessary management services. It is through this private network that the Unified Resource Manager operates. The IEDN uses 10GbE (10 Gb/second) technology to provide the data paths from each of the virtual servers within the zBX to the System z host. (See again Exhibit 3 on the previous page.)

The underlying technologies of these networks are by no means unique in the industry and that is their underlying advantage, i.e., they are ubiquitous standards. However, beyond providing high-speed connectivity, these networks are contained solely within the ensemble and completely isolated from the enterprise's networks, thus providing the highest levels of secure communications. The close coupling of the zBX

blades to the computing elements of the System z eliminates switches, cables, and inter-server hops, each of which may present opportunities for lost or compromised data, lower availability, and, inevitably, higher costs.

Close to a Cloud in a Box

A enterprise cloud strategy must include a comprehensive view of all the workloads that are likely to be involved be they ERP, e-mail and collaboration, development and test, business analytics, or high-volume transaction processing – you name it. It must provide a high degree of abstraction through the use of virtualized resources, and must be available and secure while maximizing the enterprise's ROI. Most other cloud offerings we have examined try to do this with a single (uniform) server and storage architecture. Reasonable people may reasonably disagree, but accepting one-size-fits-all as suitable for a flexible enterprise cloud strategy seems dangerously optimistic, if not completely flawed. A zEnterprise system-based strategy based on its architecturally flexible and centralized management structure would appear to be a more rational alternative.

Conclusion

At the time of this writing, over 80 zBX units and more than 400 zBX blades have been shipped to more than 60 customers, all of which have z196s installed as their host.¹⁵ Extensive testing and experience has been gained with the installation of a number of POWER7 blades running AIX, Smart Analytics Optimizers, and DataPower XI50z blades. Customers are now just beginning to ramp up the newly-available System x blades with Linux but we won't begin to see these blades running Windows until later this year. The evolution to the hybrid mainframe, more like an extended revolution, will take some time before it becomes an established paradigm, which I have no doubt that it will. I base this on my belief that IBM has much at stake here relative to the continued survival and growth of the mainframe but also the fact that this is the right solution for many enterprises at this time. Moreover, IBM is the only company that is capable of pulling this off because of their vast resources and the range and depth of their technology portfolio. Expect to see a significant acceleration of the uptake of zEnterprise hybrid technology as experience is gained with System x

¹⁵ According to IBM.

blades, particularly those running Windows, where the application portfolios run wide and deep.

If you see the potential benefits of using System z within your enterprise to deliver a hybrid solution, if you want to improve the agility of your business to respond to the rapid changes in today's environment, if you want to do this with consistent controls across your applications and platforms, or if a private cloud infrastructure is on your horizon, the zEnterprise hybrid architecture should be high on your list for serious consideration.



About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

Stephen D. (Steve) Bartlett is a Senior Contributing Analyst for The Clipper Group. Mr. Bartlett's interests include enterprise solutions including servers, system software, middle-ware, their underlying technologies, and their optimal deployment for responsiveness to emerging business requirements. In 2010, he joined the Clipper team after over 42 years with the IBM Corporation as an account and program manager in large system sales, product development, strategy, marketing, market research, and finance. During that time, he received several awards for his contributions in innovative market research and contributions to the business. He has a B.S. from Rensselaer Polytechnic Institute, and an M.S. from Union College.

- ***Reach Steve Bartlett via e-mail at steve.bartlett@clipper.com or at 845-452-4111.***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, The Clipper Group Voyager, Clipper Notes, and "*clipper.com*" are trademarks of The Clipper Group, Inc., and the clipper ship drawings, "*Navigating Information Technology Horizons*", and "*teraproductivity*" are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosures

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

After publication of a bulletin on *clipper.com*, The Clipper Group offers all vendors and users the opportunity to license its publications for a fee, since linking to Clipper's web pages, posting of Clipper documents on other's websites, and printing of hard-copy reprints is not allowed without payment of related fee(s). Less than half of our publications are licensed in this way. In addition, analysts regularly receive briefings from many vendors. Occasionally, Clipper analysts' travel and/or lodging expenses and/or conference fees have been subsidized by a vendor, in order to participate in briefings. The Clipper Group does not charge any professional fees to participate in these information-gathering events. In addition, some vendors sometime provide binders, USB drives containing presentations, and other conference-related paraphernalia to Clipper's analysts.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.