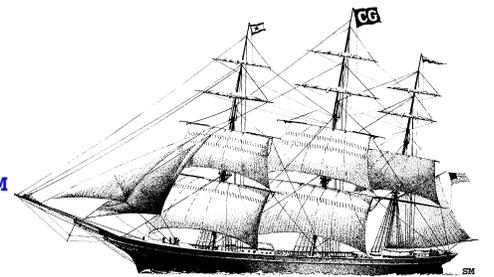


THE CLIPPER GROUP Navigator™



Navigating Information Technology Horizons™

Published Since 1993

Report #TCG2010041

September 24, 2010

zEnterprise Breaks Through Pricing Barriers — Big Rewards Will Promote Growth

Analyst: Stephen D. Bartlett

Management Summary

It is impossible to make rational choices or judgments – those grounded in facts or hard data and without bias – on a subject of which one knows little or has no experience. It becomes *judgment by hearsay*. This is what IBM Mainframe marketers and sales persons have been battling ever since the age of the dinosaurs. Look back to the period in the early 1990s when some industry pundits pronounced that Mainframes were dead. UNIX servers were thriving and hordes of youthful, aspiring computer scientists were becoming enamored with the very accessible and low-cost technology servers being driven principally by Intel and Microsoft. The possibilities and the potential seemed endless. The cost of acquisition (usually the *total cost of acquisition*, or *TCA*, but sometimes, the *total cost of ownership*, or *TCO*) was at the core of this argument.

Well, we all know how that tale played out (ten years or so later): most enterprises were drowning in hordes of distributed servers, storage elements, network interconnections, accompanied by commensurate growth of IT administrators to keep it all running, energy bills that were approaching critical stage, and service and governance challenges were beginning to appear unmanageable. Meanwhile, the dinosaurs lived on.¹ Mainframes, referring here to *IBM System z* and its predecessors, continued to be the core information systems of many of the medium and large enterprises that recognized the value of security, openness and adaptability, resiliency, virtualization capabilities, the most mature management tools, unexcelled ability to manage very large data stores, and a legacy that now spans nearly 50 years. Clearly, there are believers and non-believers, what The Clipper Group has dubbed as *diversity believers* and the *universal believers*.²

Although you might think so, I am not setting out to continue the argument of “We versus Them”³ because the reality is that virtually all datacenters that rely on Mainframe technology also rely on UNIX and/or x86 servers as well. Nonetheless, during the postlude between the announcement of the zEnterprise System – which is summarized below – and the delivery of the same in the third and fourth quarters of this year and the first half of 2011, IBM has been quietly revealing some of its *zEnterprise* pricing details, many of them innovative.

Have you ever asked the question – *Hey, that’s great technology, but how much does it cost?*” This seemingly straightforward question often has a complicated answer, whether you are talking about a car or computer. Here is a simple

IN THIS ISSUE

➤ Pricing for The zEnterprise System	2
➤ Pricing for zEnterprise BladeCenter Extension (zBX)	6
➤ Pricing for the Unified Resource Manager (zManager)	6
➤ Conclusion	7

¹ Laughing at the doomsayers, IBM enjoyed codenaming the z990 development program as *T-Rex*.

² See **The Clipper Group Navigator** entitled *Capturing The Rewards of Server Diversity - IBM's New Approach*, dated July 22, 2010, and available at <http://www.clipper.com/research/TCG2010032.pdf>.

³ Although, clearly, at that time I was one of the “We”.

The Clipper Group, Inc. - Technology Acquisition Consultants ♦ Internet Publisher

One Forest Green Road ♦ Rye, New Hampshire 03870 ♦ U.S.A. ♦ 781-235-0085 ♦ 781-235-5454 FAX

Visit Clipper at www.clipper.com ♦ Send comments to editor@clipper.com

answer. With the *zEnterprise* announcement, IBM has continued to improve the price and performance of the *System z* stack substantially: the hardware, the software, and the maintenance costs and we expect that to continue. Two cost-related challenges need to be discussed for *zEnterprise*:

1. Improving the perception and reality of Mainframe TCA/TCO, and
2. Pricing for a totally unique platform to support enterprise computation and information management.

If you are interested in more detail, then please read on to learn more.

Pricing for the *zEnterprise* System

The first eight months of 2010 has been a very busy period for IBM as they continue to demonstrate its server and storage technology leadership. These new *POWER7*, *System x* (x86), many storage systems (both disk and tape based), and now the *zEnterprise* System announcements have provided IBM's clients and potential customers much to consider. First, let's focus on the latter and summarize what was announced on that 4th Thursday of July.

Allow me to predict, somewhat boldly, that this date will be marked as the fifth great milestone of enterprise-class Mainframe computing –

1. The announcement of *System/360*, which was the introduction of *general purpose* systems⁴ [1964]
2. Large-scale virtualization [beginning in 1965 (with the *System 360 Model 67* but achieving widespread use in 1970 (with the *System/370*)]
3. The dyadic processor⁵ [1980]
4. The integration of Open Systems technologies, particularly *Linux*, within established Mainframe operating environments [1999]
5. And now [in 2010], the integration of the all the most important datacenter server technologies under common management – nothing less than the redefinition of traditional Main-

frame computing (and, as you will see, its costs)

The three principle elements of the *zEnterprise* System are the *zEnterprise 196*, or *z196* (an enhanced yet traditional Mainframe), the *zEnterprise Blade Center Extension*, or *zBX* (for running *AIX* and *Linux* workloads under the control of *zEnterprise*), and the *zEnterprise Unified Resource Manager*, or *zManager* (to control the workloads on the *z196* and *zBX*).⁶ The *z196* is the *traditional* (Mainframe) part of this system and is the natural evolution of the *System z10 EC*⁷, but incorporating many new advancements, not the least of which is a new benchmark for Mainframe price/performance. The *z196* incorporates the most advanced and fastest 4-core microprocessor technology running at 5.2 GHz, with a super-scalar pipeline, 100 new instructions, and 30 MBs of local and shared cache memory on each chip. It scales up to 96 *z/Architecture standard* or *specialty engines*⁸ – only 80 of which are available to the client⁹ – to a maximum capacity of over 50,000 MIPS, a 60% improvement in total capacity over its predecessor.

The second part of the system is the *zBX*¹⁰, which enables the private high-speed data and management networking for up to 112 *POWER7* blades – each with 8 cores at 3.0 GHz, local cache, main memory, disk storage, and an interconnection fabric, running under *AIX* and *PowerVM*. *zBX* is scheduled for availability in 4Q2010. This will be followed in 1H2011 by *System x* blades running under *Linux*. Also available for installation within the *zBX* is the *IBM Smart Analytics Optimizer (ISAOPT)* providing offloading of very-high performance queries using parallel vector processes, as well as the *WebSphere DataPower Appliance*.¹¹

The third element of the IBM *zEnterprise*

⁴ For a retrospective, see *The Clipper Group Captain's Log* entitled *The Beginning of I.T. Civilization - IBM's System/360 Mainframe*, dated March 30, 2004, and available at <http://www.clipper.com/research/TCG2004028.pdf>.

⁵ A type of multiprocessor that includes two processors which operate under control of the same copy of the operating system. [From McGraw-Hill Dictionary of Scientific and Technical Terms. McGraw-Hill Companies, Inc., 2003. Answers.com 06 Sep. 2010. <http://www.answers.com/topic/dyadic-processor>]

⁶ See *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/TCG2010033.pdf>.

⁷ EC = Enterprise Class, i.e., the class of larger Mainframes. The smaller ones were called *Business Class (BC)*.

⁸ Each a “core”, in more common terminology.

⁹ The additional engines (cores) are used for spares, internal management, and I/O acceleration. These “hidden cores” are responsible for the Mainframes high-performance and near-linear scalability.

¹⁰ Essentially, a special-purpose *BladeCenter* rack, which will use standard *POWER* and *x86* blades.

¹¹ The *System x Blade* and *WebSphere DataPower Appliance* support were announced as Statements of Direction (SODs) and, as such, are a declaration of IBM's intent, which is subject to change or withdrawal.

System is the *zEnterprise Unified Resource Manager*, or *zManager*, for short. *zManager* is the glue, the centralized management delivered as firmware that unifies the processing resources of the z196 and the zBX through its hardware, platform, and service management functions. The *zManager* is required for full exploitation of the *zEnterprise System*. Specifically, the *zManager* will provide:

- Hypervisor management and creation of virtual networks
- Operational controls, service and support for hardware/firmware
- Network management of private and secure data and support networks
- Energy monitoring and management
- Workload awareness and platform performance management
- Virtualization management – a single view of virtualized resources across the platform

With this announcement, IBM brings forth and embraces its vision of *integrated hybrid systems*. Its attractiveness lies in the system being designed and built with well-known and well-understood but heterogeneous computing elements, which bring their unique strengths to bear upon the common goal, that being to provide a cost-optimized mix of power, efficiency, and resiliency to the delivery of enterprise information technology services.¹²

A Frontal Attack on Cost

The System z family frequently has been accused by skeptics and *universalists*¹³ as being *too expensive*. That begs the question: *Too expensive for what?* It's not too expensive if your enterprise must absolutely have its core data and revenue producing transaction operations that are customer facing be secure and continuous, with virtually guaranteed qualities of service.

As was pledged by the *Mainframe Charter* in 2003, System z has delivered increased value, innovation, and an ever-expanding community of support both inside and outside IBM. The price/performance of the System z hardware, software stack, and annual maintenance has improved at a compound rate of approximately 20% over these

¹² Also see [The Clipper Group Navigator](#) entitled *IBM zEnterprise Builds Business Value in an Age of Colossal Computing*, dated July 22, 2010, and available at <http://www.clipper.com/research/TCG2010034.pdf>.

¹³ Universalists believe that the benefits of a single architecture (almost always x86) always outweigh the benefits of diverse architectures. See the paper referenced in Footnote #2.

7 years. The platform has expanded its flexibility to allow customers to consider the Mainframe for virtually *any* workload type or requirement and deliver optimum performance. As a consequence, thousands of applications from Independent Software Vendors (*ISVs*) are now available for the platform. To meet all competition for non-traditional Mainframe workloads, System z continues to offer *Solution Editions* that package hardware (either a standalone system or as a workload partition), prepaid maintenance, a comprehensive and targeted middleware stack, and storage and services, as required. These solutions offer competitive acquisition prices, leadership TCO, and Mainframe values to a variety of critical business activities.¹⁴

Rewarding Growth

To provide you with the specifics, let's start with the *zEnterprise196*. IBM rewards growth – a principle that has been integral to their Mainframe strategy for some time. z196 clients can expect the price/performance of the z196 to be improved in the range of 10% compared to the z10 EC. However, those who add substantially greater capacity may do better.¹⁵

There is bigger news for specialty engines and main memory, elements that more directly affect the comparable costs of IBM System z for non-traditional workloads, those currently running on or considering *UNIX* (more specifically, *AIX*, *HP-UX*, and *Solaris*) but also *Linux* applications on x86 servers.

Pricing for Specialty Engines

For z196, specialty engines¹⁶ have been *reduced in price* by 20% to 27% from the z10, which is significant. Announced U.S. street prices are:

- *IFL*¹⁷ – from \$75K to \$55K (a drop of 27%).
- *zAAP*¹⁸ and *zIIP*¹⁹ – from \$125K to \$100K (a decrease of 20%).

¹⁴ IBM System z will continue to have Solutions Editions offerings for SAP, ACI, WebSphere, Data Warehousing, Security, GDPS, and Application Development. As requirements dictate, other solution sets may be added, according to IBM.

¹⁵ IBM does not publish list prices, or street prices, for System z's "traditional" features and capacity, with some notable exceptions – specialty engines (IFLs, zIIPs, zAAPs, main memory, and I/O connectivity components).

¹⁶ *Specialty engines* have a "special purpose" and only can be used for that purpose, i.e., they are not standard Mainframe engines capable of running traditional (a.k.a. "legacy") applications on traditional Mainframe operating systems (like *z/OS*, *z/VM*, *z/VSE*, and *z/TPS*).

¹⁷ Integrated Facility for Linux.

Exhibit 1 – Specialty Engine and Memory Upgrade Pricing

Component	Approx. % Increase z196 over z10	z10 Pricing	z196 Pricing	% Price Reduction	z10 Upgrade Costs (%)
IFL	30% (in MIPS)	\$75K/Engine	\$55K/Engine	27%	30%
zAAP/zIIP	30% (in MIPS)	\$125K/Engine	\$100K/Engine	20%	30%
Memory (Traditional Workloads)	100% (in TBs)	\$6K/GB	\$1.5K/GB	75%	50%
Memory (New Workloads)	100% (in TBs)	\$2.25K/GB	\$1.5K/GB	33%	50%

Source: IBM, with Clipper calculations

- Main memory, the cuts are even deeper, one might say aggressive: the street price is reduced 75% from \$6K/GB to \$1.5K/GB for *traditional* (a.k.a. *legacy*) workloads.
- Main memory for *non-traditional* (i.e., *new*) workloads has been reduced from \$2,25K/GB to \$1.5K/GB as above.

(See Exhibit 1, above.) An enterprise might realize even deeper discounts if they make significant upgrades with either of these items.²⁰

But wait – there is a hitch! One of my pet peeves is paying for stuff that used to be free, but now gets added to the tab: the air pump at the gas station, checked airline luggage, maintenance of a checking account, and now IBM has introduced upgrade charges to these items. It isn't as bad as it first sounds. In prior generations, let's say when upgrading a z9 EC to a z10 EC, if you carried over your IFL engine feature (or zAAP or zIIP), it costs you nothing additional, and you got a more-performant engine, to boot! No other system vendor had this practice. The customer would have received a 30-40% performance increase built into the brand new technology and with new warranted parts without paying IBM a dime, not even a handling fee! With this announcement, IBM is asking enterprises to pay for that portion of the *added value received* when they upgrade and "carry over" the engine (to the more performant one) – in the case of a z10 EC IFL to a z196 IFL, it's 30% of the purchase price, or \$16.5K.²¹ A z196 engine is about 30% faster than a z10 EC engine (and, thus, it is no coin-

cidence that the price also equals 30%). If we use an estimate the performance of the z10 engine at 920 MIPS and the z196 engine at 1200 MIPS, IBM is charging for the difference, 280 MIPS, at a rate of about \$60/MIP! In addition, the first 920 MIPS are free (if upgrading from an IFL on a z10)! IBM effectively is offering a "trade-in" allowance of \$38.5K. Upgrade charges for zAAP and zIIP engines use similar arithmetic as the IFL upgrade, but starting with a higher street price for those engines, \$100K, which is a 20% reduction from an IFL on a z10.

Pricing for Main Memory

Main Memory has an upgrade charge as well, but remember that the 75% reduction in street prices is for *traditional workloads* and 33% for *new workloads*.²² Memory carried over from the previous generation will be charged 50% of the purchase price.²³ Remember however, that the client receives all new memory cards incorporating the latest technology and with a new parts warranty. The economics of this exchange benefits any client that upgrades their memory by 15% or greater from an older Mainframe. This is, essentially, a "gimme".

Maintenance Pricing Also Gets Attention

The zEnterprise System announcement addresses maintenance pricing on two fronts, but particularly for IFL specialty engines – where IBM most aggressively competes and is compared to Linux workloads on other (mostly x86)

¹⁸ Specialty engine for Java offload.

¹⁹ Specialty engine for database offload.

²⁰ In all cases where prices are quoted herein, all apply to U.S. only; other geographies will vary.

²¹ For an upgrade from a z9 IFL to a z196 IFL, the upgrade fee is 60% of the purchase price, or \$33K.

²² It was common practice to charge lower \$/GB for memory upgrades associated with non-traditional workloads. With this announcement IBM has normalized the memory for all types of workloads to a common rate, \$1500/TB. This does not preclude the potential for discounts to this rate at the discretion of IBM. Remember, IBM rewards growth.

²³ 70% and larger memory upgrade charges are typical in the industry.

servers.²⁴ On a per engine basis, the maintenance price of an IFL engine has been reduced 48%, compared to the z10 EC to \$763/month or \$9156 per year (per IFL engine). On a MIPS basis, the maintenance has been reduced 61% to \$0.87 per MIPS/month (about \$10/MIPS/year), which is about 12% of the maintenance charged for a standard z196 engine. This IFL maintenance reduction will significantly affect the TCA/TCO balance equation, as IBM System z embraces non-traditional workloads.

Standard engine maintenance prices are not ignored in this announcement, either. Maintenance prices have been reduced for the z196 by about 5%, for a no-growth upgrade²⁵, and average about 15% for MIPS growth exceeding 30%. Everyone who upgrades will enjoy a price reduction for maintenance, but *growers* will do much better.

And the Software Part of the Equation – the Largest and Perhaps Most Debated

How does one value very tight integration and exploitation of the underlying hardware by the operating system and the middleware? How does one value the capability to incorporate new instructions that allow compilers to optimize the execution of applications with unique operational characteristics? How does one value the simplification of many simultaneous operations and performance benefits plus the easy-to-learn and use graphic interface? How does one value the automatic discovery and configuration for fabric-attached FICON disk and tape devices? How does one value the capability to monitor the rate of error records in order to predict the possibility of a failure and then to use run-time diagnostics to immediately identify the source? The new *z/OS Release 12* contributes substantially in these areas (and elsewhere), while delivering better price/performance.

From generation to generation, *z/OS* users have been used to receiving better price/performance by an improved MIPS-to-MSU ratio. MSUs²⁶ are surrogates for MIPS that were created to flexibly offer improved price/performance for MLC software, such as *z/OS* and *DB2*, when installed on newer technology systems – what was commonly referred to as the *technology dividend*. Conveniently, it also enabled IBM to

avoid the slippery slope of quoting MIPS for each of its System z processors. (*Your performance may vary – and it usually does.*) The MIPS/MSU ratio for the z10 EC averaged approximately 7.7:1. That has not changed for the z196.

In lieu of increasing this ratio, and thus lowering the price per unit of capacity, IBM has created a new pricing metric, *Advanced Workload Charge*, or *AWLC*, which at this time only applies to the z196. Two additional tiers have been added at the top of the capacity-pricing definitions, now with a total of eight, with the highest tier covering 1,976 MSU and above. Clients of the z196 electing *AWLC* will enjoy lower cost per MSU ranging from approximately 2-3% at the lower tiers and ranging up to 10% at the highest, all compared to the z10 EC. This new metric provides a lower software price to their most advanced customers and a more flexible tool for IBM to tune the price/performance of selected software products with more granularity and discrimination, ultimately benefiting all of their customers. All the other terms that applied to *Variable Workload Charge*, or *VWLC*, will also apply to *AWLC*, including sub-capacity charges based on actual utilization. Everyone who upgrades will enjoy a price reduction for MLC software, but higher capacity growers will do much better.

There are also cost savings and an improved carbon footprint that are accrued from the indirect elements, the physical characteristics, of the new *zEnterprise 196*. Performance per unit of energy is greatly improved and new and innovative energy management functions are included. A High-Voltage DC power option is available, and system cooling is enhanced and now includes a water cooling option that reduces the overall datacenter cooling loads. These elements are addressed in more detail in an earlier Clipper report.²⁷

The TCA/TCO implications are very clear. Price/performance is improved on the total hardware, software, and maintenance stack by as much as 20%, with prices clearly linked to increased performance and capability. Acquisition costs and short term ROI are improved considerably in those arenas where *z* technology goes head-to-head against non-Mainframe alternatives for new workloads. At the same time, there are no compromises to IBM's history of delivering consistent leadership in the highest qualities of service most valued by enterprise customers.

²⁴ Only after the 12-month warranty period has expired.

²⁵ No growth implies no additional capacity (MIPS) is included in the upgrade from a z10 EC to a z196, an occurrence that is rare, as well it should be.

²⁶ MSU=Million Service Units

²⁷ See the report referenced in Footnote #5.

Pricing for zEnterprise BladeCenter Extension (zBX)

The zBX pricing presented only a minor challenge, even though it is a completely new component of the System z family. The zBX “box” has its own machine type and model, 2458 Model 002, and may include from 1-to-4 *BladeCenter* racks with the capacity to hold up to a total of 112 POWER7 blades, and eventually, IBM System x blades²⁸. The 2458 Mod 002 includes a TOR²⁹ switch, BladeCenter chassis, power distribution units, a private 10Gb Ethernet data network (to the z196), and the connectivity for the private system control and management network. IBM has advised that the zBX will be priced at approximately a 15% premium above a similar infrastructure configuration for a POWER7 or System x BladeCenter. For this small premium, the value received by the customer includes:

- A fully-integrated system with tested and supported hardware and code that requires no customer setup or assembly nor code downloads from various vendor sites to ensure a functioning machine.
- An installed system ready for customer use.
- An established service method to provide tested code updates which requires no customer action to update the various components or concern for their matching levels.
- A direct interface to IBM to address problems encountered with vendor hardware and code.
- Support for “Call Home” and diagnostics for the System Services Representative (SSR) requiring no customer intervention.
- 24x7x2-hour Warranty Support with no customer action required to replace parts or update code.

As for the IBM POWER7 blades, at this time only the *BladeCenter PS 701 Express* have been certified for use in zBX and it will be priced without a premium relative to the same blade commonly available from IBM and its dealers today.³⁰ Prices will range from about \$13K to \$22K per blade depending on configuration.

Pricing for the Unified Resource Manager (zManager)

The pricing for the zEnterprise Unified Re-

source Manager presented greater pricing challenges for IBM, because it is a unique “*service*” entity, not a separate utility program, and visible only because it is ordered as a hardware feature code. It will provide more extensive services delivered as firmware, or microcode, that has no precedence in the market. Strictly speaking, this is not the first instance of IBM pricing microcode; system capacity upgrades are essentially microcode loads that “turn on” processor engines that are already installed on the system. But the value of the resource to be priced in this instance involves three important elements: (1) the level of governance, (2) the type of resource being governed, and (3) the scope of the governance – these being unique to System z.³¹

The level of governance is to be priced in two bundles named *Manage* and *Automate*. The *Manage* bundle incorporates core operational controls, virtual server and network provisioning, hypervisor management, and energy monitoring. The *Automate* bundle *includes* all of the *Manage* capability *plus* performance monitoring and reporting, performance management, workload context recognition, energy controls, advanced energy management, availability monitoring and reporting, and availability management. IBM claims that the *Automate* bundle functionality is unique and superior to any other resource management capability available.

This year, the resources to be connected are the POWER7 blades, IFLs, and the IBM Smart Analytics Optimizer (ISAOPT), and they are priced *per connection* based on the level of governance desired. At the *Manage* level, POWER7 blades are priced at a one-time-charge (OTC) of \$1,600 per connection. At the *Automate* level, POWER7 blades are priced at \$5,000 per connection (OTC). IFLs have no charges (\$0)

Exhibit 2 — OTC\$ per z/BX Connection

Resource	<i>Manage</i>	<i>Automate</i>
POWER Blade	\$1,600	\$5,000
ISAOPT	\$1,600	\$1,600
IFL	\$0	\$3,400

Source: IBM

²⁸ A target date of 1H2011 has been stated.

²⁹ TOR=Top-of-Rack.

³⁰ However, PS701 Express Blades for inclusion in the zBX will have unique part numbers.

³¹ Software and microcode support for new features and functions for the System z family are normally “buried” within

associated with connecting them at the *Manage* level; this is essentially the same function currently provided when running under z/OS and z/VM today. However, at the *Automate* level they are priced at \$3,400 per connection (OTC). The ISAOPT is priced at \$1,600 per connection regardless of the level of governance (*Manage* or *Automate*), because the latter functionality is not extended to the ISAOPT (because of its unique processing and management characteristics). (See Exhibit 2, on the previous page.)

However, one important limitation is that all connections must be at the same level of governance, whichever is chosen by the client. It should be noted that there are no Service and Support (S&S) charges, normally associated with OTC software pricing, for the connection bundles. You may expect that as other resources, such as System x blades, are added to the mix that similar charges will be announced.

Conclusion

There is some consensus that announcement of the zEnterprise System will prove to be a game-changer in all enterprise datacenters.

- *It changes the definition of what constitutes Mainframe computing;*
- It changes how the cost of Mainframe computing should be evaluated.
- It will change the way datacenters are planned and managed as computing *silos* will begin to disintegrate.
- It will expand the universe of those who should and will consider System z.

Coupled with this innovative announcement, IBM had to deal with the long standing issues of the cost competitiveness of Mainframe technology while it continued to expand its technological capacity to be a platform for which *any* workload was feasible. Along with the traditional methods of improving price/performance that users have come to expect, IBM also addressed particular issues that have been barriers to Mainframe adoption in the past. New entities, the zBX and the zManager, which now open the zEnterprise System to a greatly expanded universe, had to be priced simply and clearly in order to relate them to their respective values. This, IBM seems to have done very successfully.

The diverse portfolio for the zEnterprise System is not yet complete and may not be for some

time. The potential applications of the zBX concept are still being defined and the capabilities of the Unified Resource Manager are very likely to be expanded in scope and granularity well beyond what has already been announced. All conscientious CIOs should watch this space closely and remember that they were present at the time of the fifth great evolution of the Mainframe.



their respective products. They are considered integral to the technological innovation and progression of the platform.

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.