THE CLIPPER GROUP
Navigator™

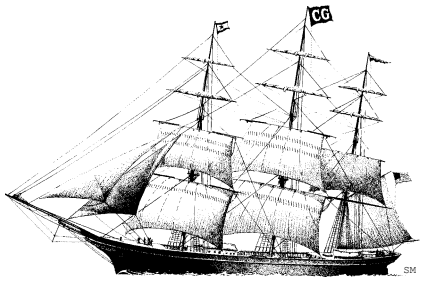

Navigating Information Technology Horizons℠

Published Since 1993 Report #TCG2011024 July 12, 2011

IBM zEnterprise in the Midmarket — Revolution or Evolution?

Analyst: Stephen D. Bartlett

Management Summary

Sometimes revolutions are good: the American War of Independence, IBM's *System 360*, the Intel microprocessor, and the Internet. Sometimes they are bad: almost every nationalist movement in Europe in the first half of the 20th Century, mobile communications everywhere 24x7 (particularly in the hands of sub-adults), and phenomena that claim to be revolutionary and aren't, e.g., the *Shamwow* and even cloud computing. Sometimes evolutions are to be preferred; revolutions can be so messy and disruptive to the status quo. One could also argue that evolutionary progress always follows revolutionary progress, be it philosophical, social, political, scientific, or as plain as dishwasher soap. Philosophical disputes most often involve whether the subject in question is one or the other. **Sometimes it really matters – revolutions often change our lives and our institutions; evolutions just require a little adaptation.**

In the world of IBM's mainframe technologies, the masters of that universe have a well-established pattern of announcing their *midmarket* (or *midrange*, if you prefer) server products about a year after their lead, or high-end, system of their latest generation. For example, the *z10 BC* (BC=Business Class) followed the *z10 EC* (EC=Enterprise Class), before that the *z9 BC* followed the *z9 EC*, and before that the *z890* followed the *z990*, etc. Now, approximately one year after the announcement of the *zEnterprise System* (with the *zEnterprise 196, or z196*, at its core), IBM follows with the *zEnterprise 114*, or *z114*. Hold on a minute, that doesn't look right! There is no "BC" suffix! When IBM seems to break a pattern, especially with the names and numbers of their products, they are usually attempting to convey a brand or marketing message. Consistent with the z196 naming, z114 appears to suggest *zEnterprise, 1st generation, 14 cores*.

Today's announcement by IBM of the zEnterprise 114 as its entry-level offering that reinforces, broadens, and expands upon the announcement of the zEnterprise System announced last July (with all the razzle-dazzle and fanfare it deserved for its revolutionary aspects of hybrid computing). **The z114 continues the evolution of the revolution initiated at that time and adds several enhancements that should prove to make it just as interesting.** Moreover, those aspects that make the zEnterprise System more attractive to existing mainframe customers and potential new ones are enhanced in the tradition established through many past generations. Thus, today's announcement is both revolutionary and evolutionary but, ultimately, that is for the reader to decide, so please read on.

IN THIS ISSUE

- The Midrange You Were Expecting 1
- Continued Innovation of the Hybrid Generation 6
- Fitting In 7
- Conclusion 9

The Midrange You Were Expecting

This story began on July 22, 2010, with the announcement of the zEnterprise System.[1] The mainframe world expected and got higher performance,

[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/TCG2010033.pdf.

Exhibit 1 — zEnterprise 114 Model M10

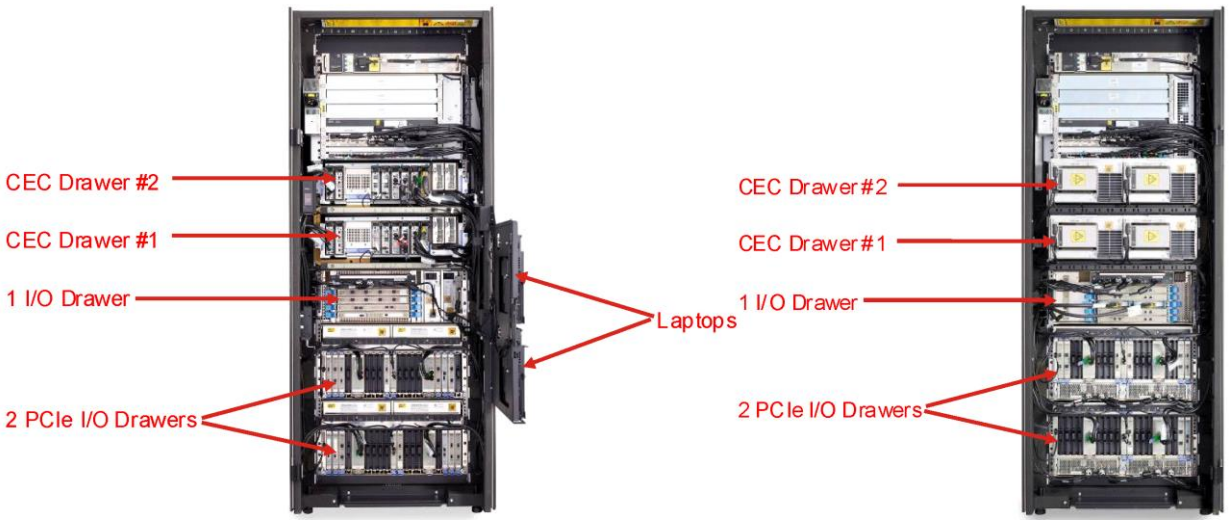

Front Layout Rear Layout

Source: IBM

expanded capacity, and perhaps some extensions of the architecture. IBM far exceeded these, 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.[2] Hybrid technologies were not unfamiliar to the mainframe world, given Linux on *System z* (also called *zLinux*). This is open-systems Linux running on a specialized zArchitecture processor called an *IFL (Integrated Facility for Linux)*, which has been available over the last five generations.

What was completely unexpected was the expansion of mainframe hardware and systems management technologies, which enabled the zEnterprise 196 to directly connect to non-mainframe processors, in particular *Power* and *System x* blades[3] via the new *zEnterprise BladeCenter Extension*, or *zBX*, connected by a private high-bandwidth network, and managed by the zEnterprise *Unified Resource Manager*. This will enable many[4] non-mainframe workloads (applications) commonly running in modern datacenters today to be managed from a single focal point with the same control paradigm as the mainframe, the z196, and now, for the entry and mid-market, the z114. All the elements of the zEnterprise System announced last year are continued and extended in the z114 announcement, now with the option of using a smaller, less costly, mainframe server. The first customer shipment of the z196 was September 2010. The first customer ship of a zBX with Power blades was last November. Expect the z114 to be generally available on September 9, 2011.

What's New in the z114?

The z114 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's 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. (See Exhibit 1, above, for pictorial views and also see Exhibit 2, at the top of the next page, for a side-by-side tabular comparison of the z10 BC, z114, and z196.)

IBM attacked the problem of lowering the z114's cost by introducing a more modular two-drawer design resulting in two models, the *M05*

[2] IBM never officially quotes MIPS for its processors, preferring to state new announcements in terms of relative capacity, *ITRs – Internal Throughput Rates*. The MIPS stated here are those commonly agreed on by industry watchers and analysts, including Clipper.

[3] 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* operating environments.

[4] The big question is what workloads are best suited for this new environment and how many there will 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).

Exhibit 2 — Comparing the z10 BC, z114 and z196

	TYPE		
DESCRIPTOR/FEATURE	**z10 BC (2098)**	**z114 (2818)**	**z196 (2817)**
Models	E10	M05, M10	M15, M32, M49, M66, M80
Clock Rate (GHz)	3.5	3.8	5.2
Uni-MIPS (1 zProcessor)	673 (est.)	782 (est.)	1200 (est.)
MIPS Range (estimated).	26 - 2760	26 - 3100	26 - 50000
Maximum # CPs (z)	5	5	80
Total Engines (Cores)	10	10	80
Max. Main Memory (GB)	248 DIMM	256 RAIM	4,000 RAIM
Sub-capacity Settings	130	130	15
zBX Connect	No	Yes, full range	Yes, full range
Total Engines including zBX	N/A	906	976
Water-cooled Option	No	No	Yes
Raised Floor Required	No	No	Yes

Sources: IBM and Clipper Analysis

and the *M10*, compared to only one model on the z10 BC. Each drawer contains two 4-core processor chips (although only 7 cores actually are used) and two System Control Processors (SCP). Each core uses 96 MB of L1 cache, half that of the z196, and 12 MB of L3 shared cache, again half of that in the z196. This is consistent with the lower number of processors and their reduced cycle-time. 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*.[5] 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. The extra costs of internal components needed to accommodate the maximum configuration are avoided, 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 MIPS[6]. By virtue of the many improved design elements of the z196[7], but at a more modest clock rate of 3.8 GHz, the z114 has a 18% faster uni-processor rating and a 12% increase in total capacity over the z10 BC, and more than a 60% advantage over the z9 BC.

Because of these and other enhancements, customers are likely to benefit with up to 18% improvement in traditional z/OS or z/VSE workloads and up to a 25% improvement[8] in CPU intensive workloads using the enhanced C/C++ compiler. Similar to the z10 BC, customers have the option of up to 130 sub-capacity settings[9] for their CPs, allowing them, combined with improvements in software pricing (more about that later), to optimize the cost of the hardware and

[5] 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 M10 also includes two spare processor cores, a first for an IBM midrange offering. Some restrictions may apply, depending on the mix.

[6] Even though a marginal increase over z10 BC, IBM never before has offered over 3000 MIPS on a midrange System z. With the additional capacity of optional IFLs, zAAPs and zIIPs, the z114 can be even more powerful.

[7] Carried into the z114 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] The z114 will exhibit up to 25% performance improvement, based on measurements and projections for CPU intensive workloads, which is provided by multiple C/C++ compiler-level improvements when going from XL C/C++ V1R9 to V1R12.

[9] View these as sub-models of the M05 and the M10 using from 1 to 5 CPs with up to 26 capacity steps. As a result, if a user needs about 200 or so MIPS, this can be accomplished through several different sub-models: 2 CPs at the 11th step, 3 CPs at the 8th step, or 5 CPs at the 5th step, for example.

software for their particular needs. At the same time, users can respond quickly and flexibly to the changing demands of their businesses or public enterprises.

Astute observers of the IBM mainframe scene will recognize that the performance bump for traditional mainframe workloads moving from the z10 BC to the z114 is substantially more conservative than the difference between the z9 BC and the z10 BC. This seems to be the result of two primary factors. First, IBM's choice of a more conservative design and packaging for the z114 results in a lower cost that could be passed on to the customer. Second, IBM's understanding of the largest target market, those with earlier generations of Business Class mainframes, since the improved z114 still offers substantial capacity and all the options, and more, of its predecessor.[10] Although no performance claims are being made by IBM relative to the z9 BC for non-traditional workloads, there should be substantially higher relative performance gains for CPU-intensive workloads, and also those workloads where IFLs, zIIPs, and zAAPs are involved (i.e., Linux on z, database-intensive, and JAVA (CPU-intensive applications), respectively).[11]

As we have come to expect from IBM, there are improved physical factors in the latest generation. The non-raised floor option is still available, but the footprint is marginally smaller by 9 inches in depth but, alas, still larger than a breadbox, but still much smaller than the space for the distributed servers that it can replace.

More significantly, the z114 M05 is expected to help lower energy demands by as much as 15% compared to the z10 BC; the M10 may come in about 7% better with the equivalent number of I/O adapters. As introduced with the z196, an overhead cabling option is available and additional and substantial data center energy savings may be gained with the new optional High-Voltage DC input. However, there is no water-cooled option available (as there is with the z196), as this would likely yield little or no benefit to the customer, especially because the processors are running at a lower clocking rate than on the z196, and thus runs cooler).

The z9 BC and z10 BC can be upgraded to either model of the z114, when available. The special-purpose engines and memory may be carried over at a nominal incremental cost (more about that below).[12] And, as expected, the z114

Exhibit 3 — zAAP and zIIP Specialty Engines

Unlike IFLs that were designed for IBM's System z family to allow low-cost entry into the open systems Linux marketplace, zAAPs and zIIPs have a more limited focus. These specialty engines were enabled to improve significantly the price-performance of narrowly focused functions with System z-enabled middleware. These engines are much less expensive than CP engines and lightened z/OS workload demands and therefore, reduce z/OS and other middleware costs, and the code runs much faster than on a shared and possibly capacity-constrained CP engine.

zAAPs are designed to enhance JAVA workloads, for example, IBM *WebSphere Application Server*, and XML workloads, specifically *z/OS XML System Services*, for validating and non-validating parsing, in fact all JAVA-based applications. zIIPs enhance a range of functions including: Communications Server IPSec processing, *Global Mirror* (a.k.a. *Extended Remote Copy*) for systems storage, *HiperSockets* (acceleration of intra-system TCP/IP messages), DB2 *DRDA (Distributed Relational Database Architecture)* requests, and Intra-Server communications using the industry-standard *CIM (Common Information Model).*

As of September 2009, IBM effectively eliminated the distinction by enabling "zAAP-on-zIIP" capability, though some minor configuration restrictions apply. It is worth mentioning that several ISV's have announced and are delivering functions within their product lines that exploit the zIIP's capabilities, with IBM's blessing and cooperation.

[10] For much larger capacity requirements, a z114 can be upgraded to a z196. A maximum configured z114 with 5 CPs would fall into the 3-4 CP engine range of the z196 with a commensurate increase in MIPS capacity.

[11] Although special-purpose (specialty) engine capacity cannot be measured in traditional MIPS, they do derive substantial benefits from all the much-improved internal performance factors that apply to a conventional engine plus they always run at full speed, even when the CPs are at sub-capacity settings. One should expect that IBM mainframe marketeers and sales reps will be using these arguments in their upgrade proposals to z9 BC incumbents.

[12] Upgrades from older IBM mainframe generations are also possible but would be a complete replacement of the CPU frame. Enterprises should consider this alternative, as attractive financial terms and allowances are likely to be available,

M10 is upgradable to the z196 M15, as well.

If I/O excites you, then you happily will welcome a huge step towards more open mainframes with the introduction of PCIe I/O Infrastructure for the z114, as well as the z196. The Peripheral Component Interconnect Express protocols were created in 2004 by a number of leading systems vendors and have become quite common in the open systems and UNIX worlds. IBM Power Systems and System x servers have incorporated this standard for some time and now it is time, some would say overdue, for the mainframe to follow suit. PCIe switches are used to create multiple endpoints out of one endpoint by a sharing architecture and protocol. In the z114, I/O drawer implementation facilitates better density, improved granularity, and is more energy efficient, as well as closer in conformance to open systems standards.

The z114 incorporates the internal data and management network of the zEnterprise System, *OSA-Express4S* for 10/1 Gigabit Ethernet, and *FICON Express8S* for storage devices.[13] Other enhancements to I/O include improved *Sysplex* clustering performance through the new HCA-3 *InfiniBand* Coupling Links.

On the Soft Side

In support of the z114 announcement, there are new releases of the supporting operating systems: z/OS, z/VM, z/VSE, and updates for Linux on z from Red Hat and Novell SUSE. (More information follows.) z/OS Version 1 Release 13 of has several new features consistent with the themes of simplification, modernizing batch, predictive failure analysis, performance, and security standards that IBM is highlighting. The *z/OS Management Facility* has a new face adding new software deployment and disk management tasks and many enhancements that help create a more productive and integrated z/OS experience. Enhancements include the ability to clone z/OS images and deploy software more easily, define new storage volumes to *SMS (System Managed Storage)*, and the enablement of a web-based *ISPF (Interactive System Productivity Facility)* interface for system programmers.

Autonomics are improved to facilitate speedier and, thus, earlier error detection. Prior releases introduced *Predictive Failure Analysis (PFA)* and *Runtime Diagnostics*. The new z/OS release provides additional monitoring of disk and file latch contention. There are updates to shorten the batch window, simplifying batch programming, and provide more flexibility in deploying batch applications. *z/OS Batch Runtime*, a new base component, with associated new functions is intended to form the foundation for a powerful, integrated, and modern batch application development, deployment, and runtime environment. The z/OS Batch Runtime provides the framework for JAVA-to-COBOL interoperability, for transactional updates to *DB2*, and for sharing database connections between JAVA and COBOL, enabling the extension of COBOL batch applications with JAVA. In addition, JES2 JCL enhancements are designed to make programming JCL easier and to give more control of batch applications. A new *z/OS Jobs REST (Representative State Transfer)* API (Application Program Interface) has been added that allows the submission, querying, and retrieval of output for z/OS batch workloads from any non-z/OS system and is intended to make z/OS batch processing much more accessible to distributed systems and web-based processes.

z/OS Release 13 also supports new encryption and compliance standards and keys. Function has been added or enhanced to address the need for stronger hashing and cryptographic algorithms and enhanced interoperability with distributed servers. In addition to the current support for *SNA Advanced Peer-to-Peer Communications (APPC)*, *RACF Remote Sharing Facility (RRSF)* is designed to support the use of TCP/IP connections. This will help to improve usability, simplify network configuration, and the security of RACF data shared between RACF nodes in the RRSF network. There are significant improvements to I/O performance for any z/OS UNIX workload using shared *zFS (zFile System)* in a Parallel Sysplex. Applications that use zFS, such as *z/OS UNIX System Services* and *WebSphere Application Server* for z/OS are expected to benefit. Also, look for updates to the C/C++ compiler, an optional-priced feature, which will provide several new functions and performance enhancements.

For z/VSE customers, who are natural prospects for the z114, 64-bit virtual addressing now is announced and will result in reduced memory constraints and the ability to exploit more "data in memory" potential. This new function will be

---

with savings in software, maintenance, and operational costs over an older mainframe system.

[13] OSA is for Open Systems Adapter. OSA-Express4S has demonstrated increased application performance by improving network access. FICON (FIbre CONnection), IBM's mainframe Fibre Channel Protocol. FICON Express8S improves performance for large sequential I/O operations and can help reduce the duration of backup/copy processes.

shipped with z/VSE V5.1 during 3Q2011.[14]

Continued Innovation of the Hybrid Generation

The addition of the z114 to 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, supported by *AIX*, the *DataPower XI50z* for XML data, and the *IBM Smart Analytics Optimizer* – which may be mixed and matched in the zEnterprise BladeCenter Extension, the zBX. (See Exhibit 4, to the right.) Now, IBM is formally announcing support for the addition of System x blades, specifically the *BladeCenter HX5 7873* dual-socket 16-core blade.

The HX5 is IBM's fifth generation of the System x providing optimal performance for database and virtualized applications. The *HX5 7873AC1* includes the *Intel Xeon* E7-2830 2-socket, 8-core processors running at 2.13 GHz and drawing 105W.[15] Initially the two supported configurations will have two 50 GB SSDs, a 10 Gb Ethernet, and 8 Gb Fibre Channel adapters and will vary only by the amount of memory included: 64 or 128 GB. List prices for the blade will start at less than $10K each, street price. There is no premium attached to the System x blade configurations because it has been certified for installation in the zBX.[16] As the HX5 blade family goes, the supported configurations are fairly modest and at the lower end of middle capacity for this product. Initially they will be supported by Linux that will be KVM-based and provided by Red Hat *RHEL 5.5* and later releases and by SUSE *SLES 11 SP1* and, later, *SLES 10 SP4*. The new blade is to be ordered and fulfilled through IBM's System x distributor channel.

When the zEnterprise System was announced last July, we were told to expect the System x blades to be supported by Linux in 1H2011. On announcement day when generally many details go unspecified, one of the first questions to be addressed to senior executives was *What about Windows?* (on the System x blades, of course). To say that the initial response could be inferred to be somewhere between "why on earth would you want that?" and "we don't think our customers will need that capability…" would be the more polite way to put it. However, the question would not go away and lingered long in the blogosphere. It struck many observers that the issue absolutely had to have been an option that the IBM architects and developers considered seriously. After all, *what is the dominant operating system on x86 servers? Sorry! Time's up.* Ultimately, it came down to setting priorities, balancing resources, and do-ability in the near term.[17] This past April, IBM revised its initial announcement by stating its intention, as a statement of direction (SOD)[18], to defer the System x blades with Linux by one quarter to 3Q and to support MS Windows server operating environments on the System x blades in 4Q2011.

Exhibit 4 — z114 Connecting to zBX

Source: IBM

In the current announcement, there are no changes to the configurability of the zBX but now with the specification of the HX5 blades, we can be more definitive and provide an illustration. As shown in Exhibit 5 (on the next page), the zBX can have up to four racks and up to 112 blades. However, when a mix of blades is

[14] For a review of recent z/VSE capabilities and its importance to IBM's mainframe strategy detail can be found in **The Clipper Group Navigator** entitled *IBM Continues Extension of z/VSE — More Function for Midrange Main frame Users*, dated April 8, 2011, and available at http://www.clipper.com/research/TCG2011013.pdf.

[15] 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.

[16] This compares to Power7 PS 701 Express blade prices of $13K - $22K that also may be used in the zBX.

[17] Often IBM can be a bit paranoid about discussing potential products that are outside a 1-year horizon; they are not considered "announce-able".

[18] 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.

Exhibit 5 — 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

desired, which may often be the case, maximums for each type must be observed, although the maximum total is 112 single-width blades. For example, a valid maximum configuration using all the currently-available blade types could be:

- 28 PS701 Power[19] *plus*
- 28 HX5 System x *plus*
- 14 DataPower XI50z *plus*
- 28 IBM Smart Analytics Optimizer

In this example, only the HX5 System x Blades are at maximum configuration.[20]

Even though the z114 announcement is focused primarily on IBM's midmarket System z customers and growing businesses looking to leverage mainframe technology for the first time, the potential of the zEnterprise System, anchored by a z114 with a fully-configured zBX, has a staggering potential with over 3,000 MIPS. It would allow as many as 10,000, or more, virtual AIX, Linux, MS Windows, DataPower, and Smart Analytics Optimizer images centrally provisioned, controlled and managed by the Unified Resource Manager. *Does this sound like cloud computing*? Of course, it does!

What has been the track record of the zEnterprise System announcement so far? IBM has not shared exact numbers but we can infer that the z196 has been phenomenally successful delivering 58% MIPS growth in 4Q2010, the highest rate in a decade, and a 34% growth rate in the first quarter this year. Clearly, a big bump in the growth rate was expected, though perhaps not at a rate that even may have surprised IBM. Nevertheless, industry watchers, and certainly mainframe skeptics, are more curious about the acceptance and experiences with the zBX.

Earlier[21], I said that the hybridization of the mainframe represented the fifth major milestone of IBM's enterprise-class computing and, therefore, represented a revolutionary step. To date, well over 50 zBX units have been installed with the z196. The companies represent several different industrial sectors that are in pilot and early stages of enablement of Power, IBM Smart Analytics Optimizer, and DataPower blades. Most are very lightly configured (since they are being tested for new uses). There should be a substantial ramp-up of zBX market activity when the HX5 blades become available.

Fitting In

Since before most people currently in the mainframe business can remember, IBM customers and prospects have steadily and sometimes vociferously expressed concern about System z's systems software and middleware pricing, those prices being primarily driven by the capacity of the system on which they were installed. This phenomenon appeared particularly egregious in the midmarket and entry-level systems, and rightfully so. Mainframe-focused ISVs also got their share of the wrath as they commonly used a similar pricing model. IBM mainframe architects and pricers have taken enumerable creative steps, too many to list here, to alter this perception and reality, and in many cases, have gone to non-mainframe market pricing models for certain workload segments particularly in those dominated by distributed systems.

In the last few generations prior to the introduction of the zEnterprise 196, it was customary for IBM to provide improved price-performance for its MLC software by defining a lower MSU[22] rating for the equivalent capacity of the prior generation. The effect was a lower monthly charge for the licensed software product(s), at a given level of capacity. Typically this was, from generation to generation, in the

[19] Additional information about Power7 blades for zBX 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.

[20] This seems to be artificially low, based on this blade's physical requirement. IBM wanted to be conservative initially because of the potential systems management challenges inherent with a non-IBM operating system, i.e., Microsoft Windows operating environments.

[21] See more information on this in **The Clipper Group Navigator** entitled *zEnterprise Breaks Through Pricing Barriers — Big Rewards Will Promote Growth*, dated September 24, 2010, and available at http://www.clipper.com/research/TCG2010041.pdf.

[22] MLC–Monthly License Charge, which typically applies to z/OS, z/VSE, and key middleware such as DB2 and CICS. MSU–Million Service Units, IBM's preferred metric (instead of MIPS).

Exhibit 6 — z114 Pricing Compared to the z10 BC

Component	Approx. % Increase z114 over z10 BC	z114 Pricing (Street)	z10 BC Pricing (Street)	% Price Reduction (z114 Over z10 BC)	z10 BC Upgrade Costs ($K)	% Price Performance Improvement
IFL	16% (in MIPS)	$35K/Eng.	$47.5K/Eng.	26%	$5/Eng.	58%
zAAP/zIIP	16% (in MIPS)	$40K/Eng.	$47.5K/Eng.	16%	$6/Eng.	40%
Memory (Traditional Workloads)	3% (in TBs)	$1.5K/GB	$6K/GB	75%	$.75/GB*	N/A
Memory (New Workloads)	3% (in TBs)	$1.5K/GB	$2.25K/GB	33%	$.75/GB*	N/A

Note (*) – 8GB or 16GB to carry forward free on upgrade, depending on z10 BC memory configuration

Source: IBM, with Clipper computations

range of 5% to 10%. The pattern was broken for the relevant software installed on the z196 with the introduction of a new pricing metric, *Advanced Workload License Charge*, or *AWLC*, a redefinition of the prior of the software-pricing curve.[23]

Likewise, with the z114 announcement, the MLC price curve[24] is being redefined again, relative to that which applied to z10 BC and prior generations. This new pricing metric is called *Advanced Entry Workload License Charge*, or *AEWLC*. The new metric provides an *average* price-performance improvement of about 5% compared to the z10 BC at the same capacity. All z114 customers will benefit from this lower price curve, but those at the lower capacity ranges will see a much larger improvement. This is a highly-targeted effort on IBM's part to remove or at least alleviate one of the major cost hurdles with which midmarket customers have been challenged. As with *AWLC*, *AEWLC* applies only to MLC software installed on zEnterprise 196 and z114, respectively.

One-Time Charge (OTC) software running under Linux (such a WebSphere Application Server, Cognos, or Content Manager) also benefit by lowering the Processor Value Unit (PVU)[25] of the IFL from 120 to 100, a 17% reduction. There are no software charges for zAAPs and zIIPs, a most attractive feature and, in some circumstances, their implementation may actually result in lowering the cost of z/OS and other middleware, as the functions they provide result in offloading operating systems workloads from CPs to zAAPs and zIIPs (i.e., an offloading of MIPS).

Where IBM is willing to quote "street prices", new lower prices have been announced for the specialty engines, memory, and hardware maintenance. (Refer to Exhibit 6, above.) z114 IFLs are priced at $35K per engine, a 26% reduction; zAAPs and zIIPs are priced at $40K per engine, a 16% reduction compared to the z10 BC. Additional increments of memory will be priced at $1,500/GB, in increments of 8 or 16 GB, depending on the model. Upgrade charges for specialty engines and memory will apply when carrying them forward from prior generations: $5K per IFL and $6K per zAAP and zIIP.[26] At this upgrade price, the customer would be buying the equivalent of about 110 MIPS of capacity for about $50 per MIP (*$50* – not *$50K*).[27]

Monthly maintenance charges for hardware, which applies upon the expiration of the one-year warranty period, have been reduced as well. For

[23] To read how IBM broke new ground in pricing its mainframe systems, see **The Clipper Group Navigator** entitled *zEnterprise Breaks Through Pricing Barriers — Big Rewards Will Promote Growth* , dated September 24, 2010, and available at http://www.clipper.com/research/TCG2010041.pdf.

[24] Really not a smooth curve but a step function with several tiers.

[25] PVUs are a non-z capacity metric assigned to normalize OTC prices.

[26] Specialty engine upgrade charges are $10.5K and $12K, respectively, when upgrading from the z9 BC to z114.

[27] As a reminder, this is still a very good deal. Where else can you get a credit for what you bought previously (in this case, for previously purchased MIPS)? Certainly not for open systems, where a three-year-replacement-and-buy-anew cycle has become the norm for many data centers.

the same capacity, the maintenance of z114 M05 compared to the z10 BC is 5% lower and is even a better deal with when the capacity grows. Maintenance charges for IFLs have been lowered by 48%, now down to $9K/engine/year. However, maintenance charges for zAAPs and zIIPs have not changed, remaining at $17K per engine per year.

The biggest price kicker for this announcement is that the customers for the z114 M05 can expect quotations of about 75% of the entry price of a z10 BC E10, lowering the entry price substantially below $100K. Contrary to the mythology perpetuated by some noisy critics, mainframes are not million dollar machines affordable by only very large corporations. This is a very aggressive price decrease and reflects not only IBM's extensive efforts to reengineer their offering to lower its manufacturing costs, but also to reflect the realities of the market where this system is targeted. As these price reductions apply to any particular customer situation, they obviously will be dependent on a number of factors, including engine quantities and mix, memory, and I/O configurations, and the software to be installed. The additional cost of the zBX and its blades, if this is also included, will be substantially, or just as likely, fully offset by the costs associated with the displacement of distributed servers, their software, networking, and operations personnel and resources. You have to investigate this carefully to see your potential economic benefits.

To Cloud, or Not to Cloud?

This question is on every CIO's agenda, or it should be. Within the next few years, every enterprise, regardless of size and technical sophistication, will be participating in cloud computing in some way, shape, or form. For the moment, putting aside how cloud computing should be defined, the current discourse strongly suggests that it will be everywhere, will affect everyone on the grid, and will evolve in ways that no one can imagine fully. Frankly, to this observer, it is a fascinating phenomenon to watch although I must admit that it has a bit of "been there, done that" taste to it having been bred and matured in the mainframe world. To those who are unable or unwilling to use the words "cloud computing" and "mainframe" in the same breath, you should really get over this biased thinking, as it does not serve you well. The announcement last July 2010 and this further evolution of the zEnterprise hybrid-computing model have slowed that down considerably.

The announcement of the z114, along with the other enhancements to the zEnterprise System portfolio, now provides:

- An entry-level, enterprise-ready, highly scalable, heterogeneous pool of virtualized resources.
- Management protocols as a single unified system.
- Multiple heterogeneous resources that can be activated, allocated, prioritized, and retired on demand.
- Utilizes automated service delivery tools.
- Provides maximum utilization of resources, improved ROI, and low cost service delivery.
- Continues to deliver the highest levels of security, resiliency, and manageability.

The focus of much of the discussion surrounding the hybridization of mainframe computing has been the mixing and matching of the zArchitecture with Power and x86. Not to be left out of consideration are the impact and breadth of computing models also embodied in the IBM Smart Analytics Optimizer addressing business intelligence/data warehousing requirements and the DataPower XI50z Appliance to provide XML hardware acceleration, streamlining, and improved security for critical service-oriented architecture (SOA) applications. Look for updated *Solutions Editions* (bundled solutions) that incorporate the z114 and zBX, with a choice of blade technology, although there are none at the time of the z114 announcement. These bundled multi-year offerings usually can be expected to include a mix of hardware, system software, middleware, maintenance, and technical services, all of which might facilitate rapid proof-of-concept and an accelerated ROI. These offerings are likely to include cloud computing, enterprise Linux, application development, data warehousing, Tivoli management, and possibly others.

Conclusion

It would be a mistake to view this announcement as just a plug-in at the low end of the zEnterprise System product family. This view suggests that the z196 will do all the heavy lifting as the vehicle for the growth of the System z-based hybrid-computing model. The full acceptance, confidence, and business returns on the innovations of System z mainframe hybridization are still in the future, the sooner the better for both clients and IBM. One can only imagine the discussion when IBM's System z sales and

financial executives were kicking this year's forecast around the table. After agreeing on the fairly straightforward questions of how many z196's and z114's were expected to be sold, the much harder question would have been – how many zBXs will be sold with each? The more enlightened position would not be to view this case from the central mainframe outward but quite the opposite.

The more modest specifications of the z114 hosting a fully-enabled zBX addresses a very broad market that includes enterprises that have System z installed but are likely to be surrounded by a larger array of multiple architectures, fractious computing models, and proliferating data. These enterprises have the most to gain by exploiting application and data affinities; these enterprises have much to gain from rationalizing the management of these heterogeneous resources within a centralized model. And these enterprises have even more to gain from retaining the traditional values and qualities of service of IBM's mainframe architecture. It would be a mistake to consider the new midrange processor as only a minor player in the hybridization of the System z architecture. In fact, the zEnterprise hybrid computing model may be even more beneficial and impactful to small and mid-sized organizations as "the dream machine" – providing the ability to integrate and manage classic mainframe with AIX, Linux on System x and, soon, Windows workloads.

Of course, this new opportunity does not diminish the importance of IBM's traditional market dynamics of replacing the old with the new, accommodating organic growth, and expanding the role of the mainframe in all the enterprise's application portfolio. With the renewed emphasis reiterated by System z marketing and sales executives regarding the importance of gaining new placements and competitive wins in growth regions, such as the BRIC countries, as well as traditional areas, the z114 and the zBX, tied together by the Unified Resource Manager, will be a compelling offering.

Is z114 a revolution or an evolution? As an analyst, I see both, but with one eye looking backward (to the BC base) and the other eye looking forward (to the many new possibilities). Truly, the answer is in the eyes of the beholder. What do you see? What can you imagine?

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, middleware, 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.