THE CLIPPER GROUP
Navigator™

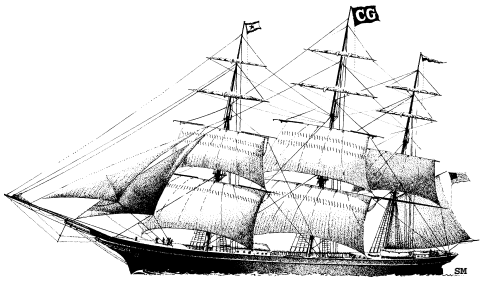

Navigating Information Technology Horizons℠

Published Since 1993 Report #TCG2010033R July 22, 2010

The IBM zEnterprise System Reaches Out — Higher, Wider and Deeper

Analyst: Stephen D. Bartlett

Management Summary

Just when you think you know what's going on and what the future is likely deliver, you unexpectedly get whacked on the head by that proverbial *2x4*[1]. First, you see stars flying around and then you sit back to assess what hit you. IBM's next-generation Mainframe is one of those expected announcements with unexpected significance. Please read on to learn *why*.

Editor's Note: This is second of three Clipper bulletins focusing on IBM's new zEnterprise Mainframe. The other two focus on (a) capturing the rewards of a diverse server environment and (b) the software that enables this new platform, especially the Unified Resource Manager. This paper deals with the "speeds and feeds": how fast; how big; the extension and enhancement of the System z architecture; the underlying technologies; how the pieces fit together; and some rationale for the choices that IBM made.

It's been about two and a half years since the last IBM large enterprise Mainframe announcement in February 2008, for the IBM *System z10 Enterprise Class (z10EC)*, so it must be time for the next. If you were expecting a *z11EC*, well it's here except IBM has named it the *zEnterprise 196*, or *z196*, to keep us on our toes perhaps, but more to identify something very new and different. By itself, the z196 will satisfy the folks waiting for an upgrade to the z10EC; that takes care of the expected vertical growth and improved performance. However, the announcement of the IBM zEnterprise System has a horizontal dimension as well. To accomplish this, IBM adds a new frame (a rack of non-Mainframe processors) – named the *zEnterprise BladeCenter Extension*, or *zBX* – which is connected to the z196 by a dedicated, secure, very-high bandwidth connection. A third element, the *zEnterprise Unified Resource Manager*, or *zManager*, is the glue that makes this amalgamation of diverse servers workable. Essentially, zManager runs independent of *z/OS* and provides unified workload management services for the entire system, referred to as a *zEnterprise Node*. (See Exhibit 1, on the next page.)

To whom is this announcement important? It may seem to be a little presumptuous to say "Everybody", but bear with me as I explain. Clearly, those who already have Mainframes installed should have a strong interest in this announcement. This is the next generation, after all, and that is always widely anticipated, especially after two-and-a-half years. Of course, IBM has a strong interest not only as a clear demonstration of its technological leadership and sophisticated solutions, but also from a business point of view, as a significant revenue and profit generator. Nonetheless, technology for technology's sake ultimately is a *fool's errand* if it does not meet or anticipate the needs of one's customers, those who we strive to serve. Therefore, on this occasion, the Mainframe folks at IBM are reaching out beyond its normal province and well into the other corners of the datacenter with a new proposition. Please read on to learn *how*.

IN THIS ISSUE

[1] For those readers in the rest of the world where the metric system is fully entrenched, a 2x4 (pronounced "two by four") is a structural piece of lumber 2 inches by 4 inches (about 50 mm. by 100 mm) by a longer dimension, usually about two to three meters in length.

Exhibit 1 – The IBM zEnterprise System

A system of systems that unifies IT for predictable service delivery

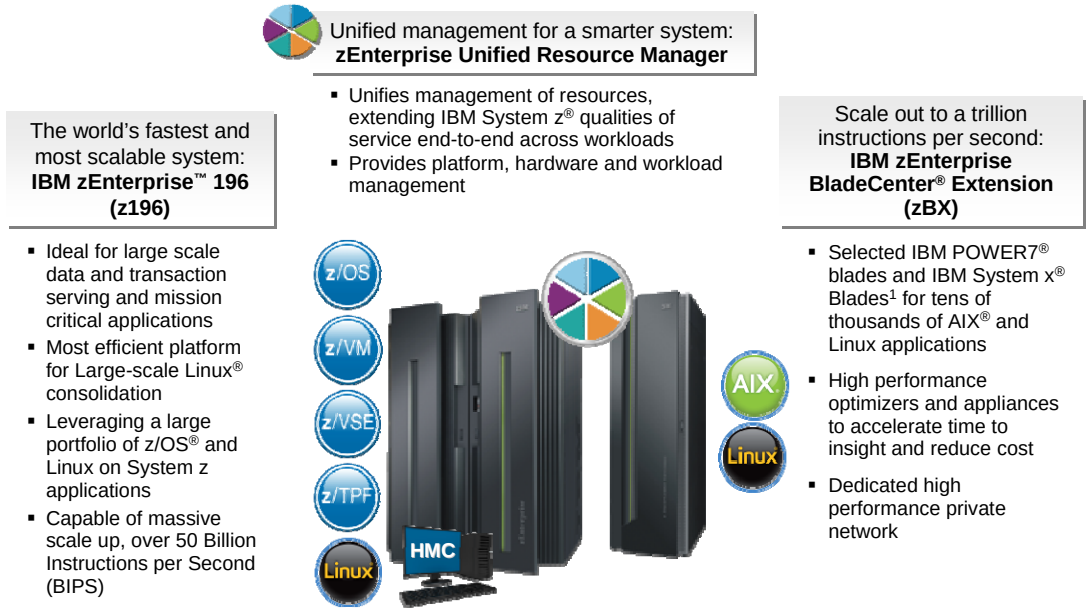

Source: IBM

In the first part of this three bulletin series, two important concepts are discussed: (1) server diversity and (2) workloads.[2]

- *Server diversity* refers to the intentional presence of diverse server architectures[3], because IT decision makers value what each brings to operational efficiency and effective delivery of business solutions and services.
- *Workloads* refer to the applications that run on one or more servers. Running only one application on a server all of the time is easy. Running many applications, especially those that are interrelated, *as needed* (i.e., on demand) or *as scheduled*, across many servers is much more challenging, especially when an important goal is to optimize use of the server assets.

Running many workloads (applications) simultaneously without a hiccup on a single server is what the Mainframe does best; and it has been king of this mountain for many decades, first with applications running on Mainframe operating systems (*z/OS*, *z/VM*, *z/VSE*, and *z/TPF* on standard Mainframe engines[4]). There are specialty engines for Linux hosting (called an *IFL* or *Integrated Facility for Linux*), for Java offload (called *zAAP*), and for database acceleration (called *zIIP*).

Thus, even when only considering the Mainframe, it can be considered *diverse*, as it hosts a variety of operating environments simultaneously. The only restriction is that they must run on the zArchitecture. The increased diversity provided by the specialty engines has been the source of significant growth for System z over the last decade. Even so, if a workload on a Mainframe wanted to share data or integrate processes with applications running on other architectures (like POWER or x86), this had to be done at arm's length and with much forethought and orchestration by the datacenter staff. In effect, each was a segregated (or siloed) environment and the connectivity between them was a network. This was the status quo in a diverse server environment, *until now*.

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

[3] In IBM's case, these are Mainframe (*System z*), Power Systems (formerly *System p*), and x86 (*System x*). Each is based on a different processor (chip) architecture.

[4] In the Mainframe world, an *engine* is a *processor core*, which commonly is called a *core* in the x86 world. There are many cores on a processor. A standard engine, known as a "CP", is one that runs without restrictions, i.e., it will run the Mainframe operating systems listed above. There are also specialty engines, which have a special purpose and are restricted to that special use. Pricing is higher for standard engines and lower for specialty engines, since their use is restricted.

With zEnterprise, IBM takes the next step by offering a more fully-integrated computing environment that incorporates the traditional Mainframe operating environments with *AIX* on POWER architecture and *Linux* on x86. The additional POWER7 and *System x* servers will be delivered in a special *BladeCenter* rack (called the *zEnterprise BladeCenter Extension*, or *zBX*). This will happen in stages, with the traditional Mainframe being delivered first, to be followed by the zBX with AIX on POWER7 blades (later this year) and then, according to a Statement of Direction (SOD), in the first half of next year, Linux on System x blades will be added to the zBX.

The key capability here is that every computing element in the structure can be individually managed and monitored in order to achieve the highest and most effective utilization, ensuring delivery of client services in a consistent and reliable manner. This is done with the qualities of service and reliability normally associated only with System z Mainframes.

Those not possessing a Mainframe might ask, *What's in this for my organization and me?* In the short term or for their next acquisition, it would seem very little. However, IBM is showing its hand here and it should not be ignored; they will continue to drive towards higher and higher levels of integration and System z is carrying the management banner. Beyond its extensions to general-purpose POWER7 and an SOD for System x blades, there is an SOD for a *WebSphere DataPower (SOA) appliance* for zEnterprise, plus there is the already-announced *Smart Analytics Optimizer*.[5]

The Mainframe You Were Expecting

What is the zEnterprise 196 (z196) and how does it compare to prior generations of IBM's Mainframe systems? If you are familiar with the z10EC than the z196 will appear to be its stronger, more powerful younger brother (due to better luck in the gene pool!) and only the discerning eye could tell them apart; they are physically the same size and appearance, plus their environmental factors also are virtually identical. This means that a z196 could easily slip into the same space as a z10EC without any significant data-center modifications.

Again, you will see up to a four collections of processors, packaged into modules called *books*, but this time carrying 24 engines each, making for a maximum of 96 z/Architecture engines.[6] However, only 80 of the 96 may be used by the customer as standard or specialty engines, with the rest reserved for I/O acceleration and internal management (*Systems Assist Processors* or *SAPs*) and spares. As in z10EC, there is co-located cache memory, a large main store, and the same high-bandwidth I/O processors. Also continuing are subcapacity settings, 15 in this case, allowing operation of the standard processors at a "dialed-down" level of performance (to meet lesser needs, with a savings on software costs, as well).

The *Integrated Cryptographic Facility* continues with enhancements providing industry-leading intersystem security. And, as has been IBM's practice for many years, upgrades are available from the *z10EC*, as well as the *z9EC*. IBM protects its customer's assets by not penalizing the customer who chooses to skip a generation. There are five models of the z196 (Machine Type: *2817*) as before but these are designated as *M15*, *M32*, *M49*, *M66*, and *M80*, with the numerics referring to the number of customer-configurable processors available.[7] z196 will be supported and its capabilities exploited by new releases of the same operating systems: z/OS, z/VM, z/VSE, z/TPF, and *Novell SUSE* and *Red Hat Enterprise Linux* distributions for System z.

But It Gets Better – Much Better

Whenever a new IBM Mainframe generation is announced, always the first question asked is *How much faster is it?* IBM is quoting, always with an "up to" (your performance may vary, and it always does, however IBM tends to be conservative):

- 40% improvement for traditional z/OS workloads (batch, *CICS*, *IMS*, *TSO*, and *DB2* transaction workloads mostly)
- 30% additional improvement[8] in CPU-intensive workloads via compiler enhancements; and
- 60% total capacity improvement.

[5] The July announcement includes the z196, zBX, zManager, Power7 blades, and the SOA with SODs for System x blades and WebSphere DataPower appliance. Delivery of the z196 begins in the 3rd quarter with the balance of the zEnterprise system shipping in the 4th quarter. SOD products are expected in first half of 2011.

[6] The z10 had a total of only 80 engines, with a maximum of 64 that are customer usable.

[7] The sum of all your CPs, IFLs, zIIPs, zAAPs, ICFs, and optional SAPs. The first three models are 1-, 2-, and 3-book (MCM) models, respectively, and the last two models include 4 books (MCMs).

[8] This 30% is in addition to the 40% for traditional z/OS workloads. The improvement is cumulative.

Exhibit 2 – the z196 Processor Chip

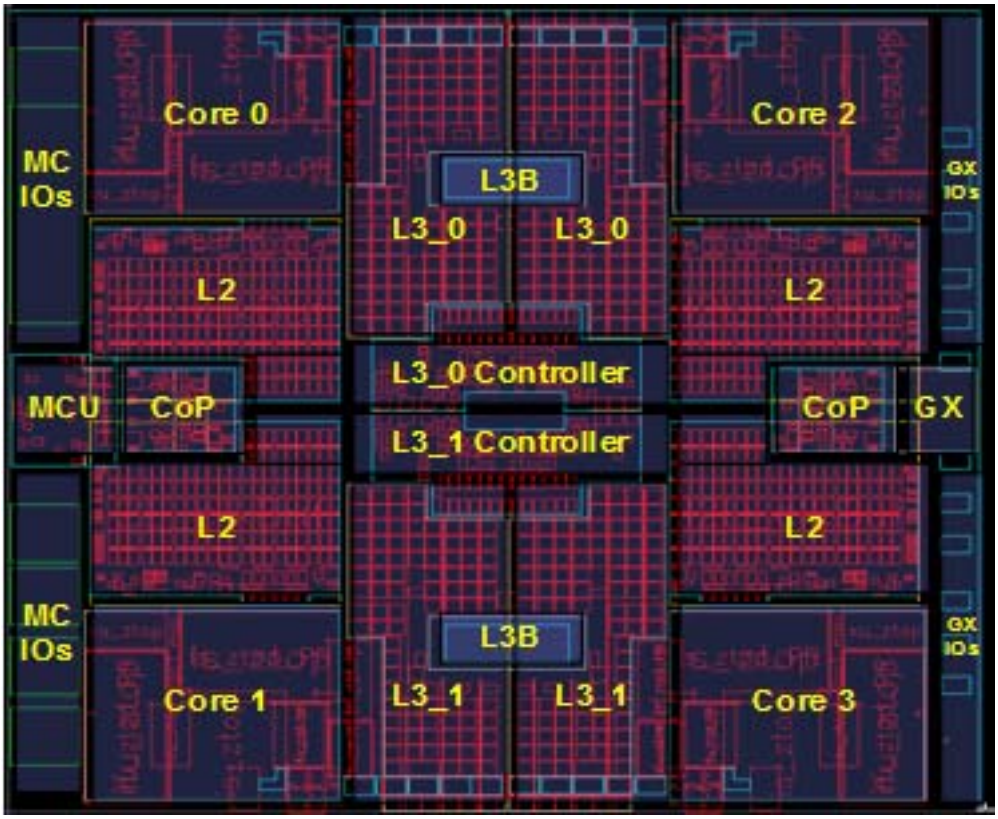

Source: IBM

How this is accomplished is the most fascinating part of the story and usually the most arcane. Nonetheless, let's look at the highlights.

The z196 Processor Chip

Using 45nm SOI[9] technology, the z196 packs in more function into the same space, when compared to the z10. (See Exhibit 2, above.) z196 uses quad-core processors (as in the z10) but the cycle time is boosted to 5.2 GHz, 18% faster.[10] It adds 100 new instructions resulting in substantial improvements to CPU-intensive, Java, and C++ applications, when used with the latest compiler versions that take advantage of these new instructions. Also, expect to see significant improvements for floating-point workloads.

This processor also includes three levels of cache, one more than the prior generation, and there is more than twice as much cache – about 30 MBs, both local and shared – on each chip. More cache improves the "locality" of data, which helps to optimize data-intensive operations by reducing the need to reference data, including program instructions, off the chip, either in L4 cache or in main memory.

Instructions are decoded and executed in what IBM calls a *Super-scalar pipeline*; think of this as a set of carefully designed and orchestrated parallel execution pathways. Up to three z/Architecture instructions can be decoded per cycle. Up to six execution units are available, each specialized for different instruction types and lengths; up to five operations can execute per cycle.

In addition, the architecture introduces Out-of-Order (OOO)[11] execution, the essence of which is that execution of some instructions may continue (in parallel, without waiting) while others are waiting for waiting for data from L2 cache and beyond, and for results from other instructions. This allows better utilization of the

[9] *nm:* nanometers (10^{-9} meters); *SOI:* Silicon-on-Insulator.

[10] As a point of comparison, Intel's *Nehalem EX* has a clock speed of 2.66 GHz for multi-processor implementations. So the Mainframe is inherently faster (more than twice as fast), in spite of its very rich instruction set. The comparison leans even further in favor of the z196, since there are several-to-many SAPs (additional cores) to accelerate I/O and internally manage the server, which must be done by the standard Nehalem processors, effectively reducing the performance capacity available for users' applications. While comparing clock speeds may not be an apples-to-apples comparison, it does indicate why the z196 is so powerful.

[11] Inside joke: spoken as *ooo* (as in loose) – *that's fast!*

execution pipes and hides the latency inherent in most storage and user access. Instructions are begun and decoded in program order, reordered during execution as needed, then returned with their results in the original sequence, all of which is invisible to any software or application.[12]

The z196 Cache

In addition to the four cores included on each z196 processor chip, there are coprocessors for accelerating cryptography and compression, interfaces to the L4 cache, and controllers for the I/O bus and the main memory. This level of integration and design facilitates the first level of the performance improvements for the z196. It is obvious that this is not your father's everyday commodity processor chip.

The next level of integration is found on the *Multi-Chip Module*, the *MCM*, which is similar on the z10. However, there are six processor chip sites, each containing four cores, yielding up to twenty-four active cores.[13] In addition, there are two sites for System Control (SC) chips, each containing 96 MB of L4 cache for a total of 196 MBs of shared eDRAM[14] cache per MCM. Compared to the z10, the z196's 4-level cache is four times larger than the 3-level cache found on the z10[15]. The L4 cache provides the means of very efficient sharing of instructions and data with each of the other four MCMs on a fully populated z196. The new cache design results in significantly improved probability of finding the next required instruction or data element nearby electronically, the speed of light always being a factor in computer architecture.

The z196 Main Memory

The main memory of the z196 has been doubled from the z10 to three terabytes. IBM refers to the z196 memory design as *Redundant Array of Independent Memory (RAIM)*.[16] The custom DIMMs are connected to three of the six processor chips on the MCM via their memory controllers to the L3 cache and, therefore, are spread physically across the MCM. The enhanced design of this memory provides redundancy and multiple levels of checking and retry with the reliability objective of zero memory faults during the life of the system.

z196 Performance

Adding all of these factors together, the nominal performance of each z processor core is approximately 1200 MIPS, compared to about 920 for the z10[17]; with the aggregate capacity of the z196 exceeding 50,000 MIPS. It now becomes more convenient to refer to Mainframe capacities in term of *BIPS*[18]. To put this into further perspective, this single system capacity of a z196 M80 is in the same order of magnitude as the total capacity of all Mainframe systems installed 25-to-30 years ago.

This would be a good point to mention the implications of the above discussions to the System z specialty engines – the *IFL*s, *zIIP*s, and z*AAP*s, in particular. Strictly speaking, these engines share the same hardware as your normal zArchitecture engines, which are measured easily in MIPS, if you will. The specialty engines however, have customized microcode that enables their particular and more narrowly defined functions. Although their capacity cannot be measured in traditional MIPS, they do derive benefits from all the much-improved internal performance factors that apply to a conventional engine. With about a 40% improvement in raw capabilities over z10 engines, the specialty engine customer will get a lot more processing power on zEnterprise.

Water Makes a Comeback – and Other Improvements for the Datacenter

IBM has demonstrated, in a number of ways, its focus on more energy-efficient computing, including grand challenges, such as a more efficiently managed and secure national power grid. This focus now is being demonstrated by new

---

[12] No recompilation is required to get benefits of OOO. However, recompilation with latest compilers (COBOL, PL/1, Java, C/C++, etc.) will improve code optimization for z196.

[13] There may be 24 active cores per MCM but only a maximum of 20 are available for the customer's use; multiplied by four MCMs provides the maximum 80 customer-usable cores configured in model M80. For all models other than the M80, four processor chips on the MCM would have only three cores enabled, while the remaining two would have all four cores enabled. The remaining cores are used for SAPs and spares.

[14] *eDRAM:* Embedded Dynamic Random Access Memory.

[15] In the z10, the three layers were labeled L1, L1.5, and L2.

[16] RAIM might be described as "RAIDed Memory". While this stretches the meaning of RAID (since there are no disks involved), it provides for the z196's main memory a very similar kind of striping, comparing, and protection that RAID provides for disks. Simply put, the z196's memory is now even more reliable (protected against errors and failure).

[17] Most workloads, including LSPR workloads, will see the about the 40% overall processor performance improvement.

[18] MIPS: millions of instructions per second – whimsically also known as "meaningless indicator of processor speed". BIPS: billions of instructions per second. IBM never officially quotes MIPS but prefers ITRs – Internal Throughput Rate, which is based on a common reference point originating a number of generations back.

features exclusive to the z196. They include new energy-management functions and, optionally, water cooling, direct DC power, and provisions for overhead cabling.

As previously stated, the physical, power, and thermal footprint of the z196 are virtually identical to a similarly-configured z10. To enhance the power and thermal management further, new triple-redundant pressure and humidity sensors have been incorporated into the new system. Pressure (altitude) sensors contribute to saving blower power, while the humidity sensors protect the system from abnormal dew point conditions that would produce undesirable condensate within a processor frame.

There is now a *Static Power Savings Mode* which a system manager would want to initiate for a system that will have extended periods of low processor utilization, such as for a backup system with many CBU (Capacity Backup) engines that are standing by (to be put into action) in case of a failure. This enables up to a 20% improvement in power utilization. For pre-installation planning, IBM provides a Power Estimator Tool to give detailed guidance on system energy requirements. In addition, the z196 provides a facility, aptly named *Query Max Potential Power (QMPP)*, which calculates the maximum power requirements based on a number of internal and external variables. Since the z196 does not incorporate a power-capping function, which could slow down processor cycles undesirably at perhaps the most critical periods, the QMPP function provides guidance for planning safety margins for power provisioning. These are all standard features of the new processor.

Like the z10, the z196 continues to be air-cooled at the system level, using front-to-back airflow, but with an efficient, localized liquid cooling for the processors. However, this announcement introduces an external water cooling option – not because of any extraordinary heat densities in the new processor – but to provide an additional opportunity to improve the overall efficiency of the datacenter.[19] IBM says water is capable of removing heat 4,000 times more efficiently than air. In a typical air-cooled datacenter, up to 50 percent of energy consumption and the associated carbon footprint is caused by the cooling systems to keep processors from overheating.

This z196 option incorporates two air-to-water heat exchangers within the back of the *A* and *Z* frames, which adds 4 inches to the back of both frames, and is provisioned by redundant *Water Conditioning Units (WCUs)*. This is a closed-loop system maintained by IBM, but requires connection to the datacenter's chilled-water system. At the system level, this results in approximately 25% less airflow, compared to an air-cooled system. There are further input power savings in the range of 2.5kW, with additional savings of approximately 2.5kW in reduced air-cooling heat load; most of these savings being achieved for 3- and 4-book systems. The datacenter benefits by reduction of overall heat and power load, the elimination of hot spots, and the opportunity to increase compute power density with lower impact to environmental requirements.

Newer datacenter designs are considering High Voltage Direct Current (HVDC) as a power option, because of its potential cost savings and lower overall power consumption in the range of 10%. zEnterprise optionally will operate on HVDC, while continuing to support the wide range of AC power options available now. This is delivered with z196, where the HVDC power option will eliminate the need for an additional DC/AC inversion step in the datacenter, while also providing an energy savings of up to 3% at the system level.

Optional overhead cabling provisions will be available on the z196. Four additional enclosed cable raceways will be offered that are attached at the external corners of the two system frames to facilitate the passage of I/O, fiber optic, and Ethernet cables. This will provide much greater flexibility in datacenter design and layout (an I/O floor above the server floor, for instance). Power and the optional water cooling continue to be bottom exit only.

zBX – The New Kid on the Block

It should be no revelation that BladeCenter servers, containing POWER and/or x86 blades, are likely to be installed in the same datacenter as a System z, since IBM has long advocated the advantages of its servers at every tier. However, the IBM *zEnterprise BladeCenter Extension*, the *zBX*, is specifically designed for direct high-speed connection to the z196 Mainframe, with the zBX systems and workloads managed from the Mainframe.

[19] IBM introduced external water cooling in the era of the System 370 Model 168, circa mid 1970's, but was no longer required when Mainframes abandoned bipolar technologies and adopted CMOS, where heat densities were much lower. The water cooling option described here benefits from IBM's extensive experience in thermal design, for which it has numerous patents, plus recent designs for Power6-based high-performance computing clusters.

The zBX has its own System z-like Machine Type and Model, *2458 Mod 002*, and is ordered in the same fashion. It contains the BladeCenter Chassis, Power Distribution Units, a TOR[20] switch, and IBM blade servers, either the *BladeCenter PS701 Express* (POWER Architecture), or in the future (under an SOD), one or more blades based on the System x (x86 architecture), all integrated and certified for this amalgamated use. A zBX frame may contain one-to-four racks, 28 PS701 Express Blade Servers per rack, with 8 cores per blade, for a total of 896 cores.[21]

A few words about the BladeCenter PS701 Express are in order. This is a single-wide blade server incorporating eight 3.0 GHz POWER7 cores, each core including L2 and L3 cache memory of over 4 MB. Main memory is available up to 128 GB per blade. Also included in the server package is a 2.5" SAS disk drive, at up to 600 GB capacity. Ethernet, Fibre Channel, or InfiniBand connectivity may be specified. This family of blade servers is characterized by it numerous RAS features, virtualization, and systems management tools. This will be supported by the current version and releases of AIX and the *PowerVM* hypervisor (for virtualization management).

The operating systems support is for AIX for the POWER7 blades (expected later this year) and for Linux on System x (expected in the first half of 2011). Interestingly, Linux support for the POWER blades is not on the current horizon nor is support for Windows Server on the System x blades. Although System z management software will be running in the zBX[22], any application that is currently certified for either a POWER (AIX) or System x (Linux) BladeCenter blade will not require any changes to be included as a zBX application. The zBX will connect to the z196 by three networks: two for systems management and the other is 10Gig Ethernet dedicated to moving the enterprise's data. The zBX will have the same level of support as the z196, including problem reporting and resolution, and active hardware and firmware updating, when required.

The third element of the IBM zEnterprise System is the *zEnterprise Unified Resource Manager*, or *zManager*; this is the glue, the software management piece that unifies the processing resources of the z196 and the zBX through its hardware, platform, and service management functions. The zManager is part of the *System Director* family, is delivered as firmware with the zEnterprise, and is required for full exploitation of the zEnterprise System. More specifically, the zManager provides:

- 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 virtualization across the platform

A few examples of the zManager's scope are the integrated deployment and configuration of hypervisors (including z/VM, PowerVM, and an undisclosed virtualization manager on the x86 blade); auto-discovery and configuration support for new resources; monitoring and trend reporting of CPU energy efficiency; wizard-driven management of resources in accordance with specified business service level objectives; and more. The zManager will be used to drive the increased flexibility of the zEnterprise's multi-tier, multi-architecture structure to deliver Mainframe qualities of service throughout the system. Look for IBM to deliver increasingly broader scope of systems integration through this vehicle.[23]

Putting the Pieces Together

With this diverse set of servers now coming together with a single point of control (zManager), IBM has added some new clustering concepts and terminology.

- A ***zEnterprise Node*** is a single z196 and a single zBX instance, with 1-to-4 blade server racks. Thus, later this year, the zEnterprise Node will max out at 80 customer-usable Mainframe cores and 896 POWER7 cores.
- A ***zEnterprise Ensemble*** is a collection of 1-8 zEnterprise Nodes managed as a single virtualized pool of server resources. Thus, later this year, a zEnterprise ensemble will

[20] Top-of-Rack…for easier cabling.

[21] There are two chassis per rack, each housing 14 single-wide blades, for a total of 28 per rack. Thus, in four such racks, there are 112 blades.

[22] The software is delivered as firmware to the zBX.

[23] See **The Clipper Group Navigator** dated July 22, 2010, entitled *IBM zEnterprise Builds Business Value - in an Age of Colossal Computing*, which is available at http://www.clipper.com/research/TCG2010034.pdf.

max out at 640 customer-usable Mainframe cores and 7,168 POWER7 cores.

What IBM has assembled under a single point of control is the largest footprint for commercial processing, capable of an unimaginable amount of application execution and workload management. The maximum capacities of an ensemble are incomprehensible to most of us, but recognize that this kind of amalgamation has been available on the Mainframe for generations via Parallel SysPlex connectivity, especially valuable between separated datacenters. Thus, what is new about the zEnterprise Ensemble is that this stellar management will be extended to incorporate POWER7 and x86 blades.[24]

In an interview with The Clipper Group[25], Martin Kennedy, Managing Director for Citigroup's enterprise systems and infrastructure, said of zEnterprise: *The hybrid architecture provides a great opportunity to integrate our datacenters' heterogeneous platforms. … A very attractive aspect is the expectation of greatly reduced effort and complexity.* He added: *We want to leverage the security of z/OS and the crypto engine … in a really tightly-coupled [private] network. … Today, we have a lot of process and governance surrounding data security. zEnterprise has the potential to make this easier to implement and manage.*

Conclusion

With this announcement, IBM brings forth and embraces its vision of integrated hybrid systems. Its attractiveness lies in this 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 the most optimized mix of power, efficiency, and resiliency to the delivery of information technology services. zEnterprise reaches *higher*, as one would expect from a new Mainframe generation. It reaches wider, as the Mainframe architecture is *broadened* to encompass more diverse and rigorous demands of maturing as well as emerging customer applications. It reaches *deeper* into the datacenter by incorporating and integrating more-highly specialized services that heretofore have been isolated into separate operational silos.

Most exciting is the combined processing power and virtualization capabilities of a thousand or more diverse processor cores under a centralized management scheme, one that is built on the Mainframe's well-established operational philosophies and well-respected qualities of service. The common goal of operational optimization is extremely well fulfilled through IBM's new, unified hybrid-computing solution. Take a close look and reflect on all that zEnterprise has to offer.

[24] As in previous generations, there are restrictions on what older Mainframes can be attached in a Parallel SysPlex. These mostly have to do with lack of Parallel SysPlex Infiniband support on older machines. In general, if you have one or more z10s, everything connects (except the zBX). However, if you have prior generations, the possibilities for connecting to the z196 are restricted (for *z9s*) or unavailable (for *z800s* and *z900s*).

[25] On July 16, 2010.

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.