THE CLIPPER GROUP
Navigator™
Navigating Information Technology Horizons℠

Published Since 1993 | Report #TCG2013013 | July 23, 2013

IBM's zEnterprise BC12 — More of What You Need

Analyst: Stephen D. Bartlett

Management Summary

When it's time to buy a new automobile, the most important question relates to your driving needs: *how you plan to use it?* Hauling people, adventure traveling, commuting to your office, soccer carpooling, etc.? *What is your style of driving?* Maybe you just want to reflect your self-image, which is purely subjective. When it comes to what is under the hood there are the plug-in electrics, various levels of hybridization of fossil fuel and batteries, gasoline or diesel, turbo or not, and, of course, engine designs including I-4s, I-6s, V-6s, V-8s, even the more exotic V-10s and V-12s still can be found. (Avoid bright orange cars prowling the streets, looking for "prey".) Dismissing the lunatic fringe, the majority of new buyers will be looking at something between an I-4 and a V-8. A V-8 is always great, but sometimes a nice smooth, economical, and less expensive I-4 will do just fine without giving up too many luxury features. **And so it is in the world of IBM's zEnterprise systems. *What are your computing needs and will the more modest, economy model do just fine?***

Exhibit 1 — The New zBC12

Source: IBM

In the world of IBM's mainframe technologies, there is a well-established pattern of announcing its midmarket (or midrange, if you prefer) server products about a year after its lead, or high-end, system of the latest generation. For example, the *zEnterprise 114 (z114)*, followed the *zEnterprise 196*, or *z196*, before that the System *z10 BC* (BC=Business Class) followed the *z10 EC* (EC=Enterprise Class). Now, approximately one year after the announcement of the *zEnterprise EC12 (zEC12)*, IBM follows with the *zEnterprise BC12* (*zBC12*, for short) and, coincidentally, reestablishing the mainframe's, high-end, low-end, naming conventions. (See the zBC12 in Exhibit 1, above.) The question in the minds of many mainframe watchers is: *Other than scale (smaller) and price-performance (lower), are there any real differences between the two?* Often, it is overlooked that IBM sometimes uses the occasion to pop out a few new wrinkles in the family that add to the architecture, performance, availability, scale, or security of *System z*, frequently in concert with enhancements to the supporting operating systems that support the family, particularly *z/OS*.

Today, IBM's System z development, finance, and marketing folks hit the ground running with its

IN THIS ISSUE

first anniversary announcements for the mainframe's twelfth generation. The item at the center of attention is the new midrange/midmarket System z, the *zBC12*, but also there are several new and exclusive features that extend the architecture, a new version of *z/OS Version 2.1*[1] and a new release of *z/VM Version 6.3*, that bring it all together. In this **Navigator**, the details of this announcement are described in detail plus why these are important to your enterprise IT strategy. Please read on for more details.

Technology Transformation is Driving New Business and IT Strategies

There are those that argue which came first, the "Cloud" or the meteoric rise of social media? At this point, it's irrelevant. What is relevant is how enterprises, both public and private, deal with these two phenomena while they are building new services, reaching for new clients or customers, or trying to distinguish themselves from their competition.

- *How do you deal with new delivery models and their demands?*
- *How do you deal with the explosive demand for mobile access in its many forms?*
- *How do you manage and exploit the massive volumes of data that the interconnected world is generating?*
- *How do you mitigate the risk to the security of your systems and your clients' or customers' data as you respond to these demands?*
- *And most importantly, how will you be able to manage and merge these new technological demands cost effectively while maintaining the accelerated pace of our enterprises.*

The challenge to the System z team is to demonstrate how the mainframe not only participates in these transformations but, in fact, leads in being able to deliver them with the most effective products and solutions. Let's start with the new zEnterprise BC12.

Gen 12 Brought to the Midmarket

The announcement a year ago of the zEC12[2] provided a 'heads up" on what we could anticipate for the zBC12, in terms of its technical content and how it would be positioned by the System z folks, and assuming it would follow previous patterns. In that sense, it did not disappoint but there are always a few surprises here and there that help to make the announcement even more exciting. It must also be remembered that since this is also the first anniversary of the zEC12 announcement, what IBM likes to call "GA2", meaning the "general availability date #2", this is also an opportunity to bring new features and function to the product line to address performance, availability, and security as IBM continues to evolve, improve, and extend the System z architecture.

Bringing the z Architecture Forward

Inherited from its big brother is the six-core chip albeit running about 24% slower than on zEC12. This results in a single-engine (CP) MIPS rate of 1064. With a z/OS maximum configuration of six CPs, the total capacity expressed in z/OS MIPS is approximately 4900.[3] The actual maximum capacity is a good deal larger than that taking into account the computing capacities of the special purpose processors that could be included in a maximum configuration for zBC12 of 13 cores, which I estimate to be in excess of 8700 MIPS.[4] Carried forward are the architectural extensions, including:

- The Transaction Execution Facility which improves parallelism and scalability by exploiting transactional memory.
- The Runtime Instrumentation Facility that uses heuristics to reduce JAVA workload overhead.
- 2 GB Fixed Page Frames improves performance for DB2 and JAVA workloads.
- 1 MB page frames for better performance and availability.
- New instructions incorporated to enable broader exploitation of Decimal Floating Point (DFP) arithmetic that improves performance.

---

[1] z/OS Version 2.1 also announced today. For more detail see **The Clipper Group Navigator** entitled *IBM z/OS Version 2.1 — Ready to Go*, dated July 23, 2013, and available at http://www.clipper.com/research/TCG2013014.pdf.

[2] For more detail on the zEC12, see **The Clipper Group Navigator** entitled *The IBM zEnterprise EC12 - Bigger, Better, Faster*, dated August 28, 2012, and available at http://www.clipper.com/research/TCG2012019.pdf.

[3] When compared to a fully configured zEC12, the zBC12 offers about 6% of the capacity of the maximum larger system, which is rated at approximately 78,000 MIPS with z/OS workloads.

[4] Even more if the non-user-accessible processor cores which handle all the I/O processing, encryption, data compression and other firmware that are built into the chip are included in a capacity estimate.

Exhibit 2 — Comparing zBC12 to Its z114 and z10 BC Predecessors

Name	zBC12	z114	z10 BC	zBC12 Compared to z114	zBC12 Compared to z10 BC
Month/Year Shipped	09/2013	09/2011	10/2008		
Type	2828	2818	2098		
Models	H06 / H13	M05 / M10	E10		
Chip Process	32nm	45nm	65nm	n/a	n/a
Cycle Rate (GHz)	4.2	3.8	3.5	+10.5%	+20%
Entry MIPS	50	26	26	+92%	+92%
Single CP MIPS (a)	1064	782	673	+36%	+58%
Total CP MIPS	4900 approx.	3100 approx.	2760 approx.	+58%	+78%
Subcapacity Steps	156 (26 x 6 CPs) <= 4400 MIPS	130 (26 x 5 CPs) <= 2800 MIPS	130 (26 x 5 CPs) <= 2500 MIPS		
CPC Drawers	1 / 2	1 / 2	1	same	+100%
PU SCM per Drawer	2 / 4	2 / 4	4	same	same
Active Cores / PU SCM	4 & 5	3 & 4	3 (b)	+25%	+66%
Total Active Cores	18	14	12	+29%	+50%
Max. User Configurable Cores	13	10	10	+30%	+30%
z/OS CP Max. (c)	6	5	5	+20%	+20%
L1 Cache (Core) (d)	64KI +96KD	64KI +128KD	64KI +128KD	-17%	-17%
L1.5 / L2 Cache (Core)	2MB (e)	1.5MB (e)	3MB (e)	+33%	-34%
L3 Cache (Core / SCM as noted)	24MB (Core)	12MB (Core)	48MB (SCM)	+100%	-50%
L4 Cache (SCM)	192MB	96MB	n/a	+100%	n/a
Entry Main Memory (in GBs)	8 / 16 (RAIM)	8 / 16 (RAIM)	4 (DIMM)	same	+100%
Maximum Main Memory (in GBs)	512 (RAIM)	256 (RAIM)	256 (DIMM)	+100%	+100%
Upgrades To	zEC12	z196, zBC12	z114, zBC12		
zBX Model	2458-003	2458-002	2458-001		
Footprint /Power Envelope	Same as z114 approx. (f)	Same as z10 BC approximately	Baseline		

Core Configuration: All models may include CP, ICF, IFL, zAAP, zIIP, or optional SAP processors.

Notes
(a) CP = Core configured for System z architecture
(b) Only 3 of the 4 Processor units are active
(c) Maximum z/OS CP per CPC, remainder may be IFL, zIIP, zAAP, ICF, or SAP
(d) L1 cache is divided for use by Instructions (I) and Data (D)
(e) L1.5 cache available on z10 server only, equivalent to L2 cache on other models
(f) Some cooling and cabling options extend dimensions to a small degree

Source: IBM, with computations by Clipper

Also, *IBM zAware* is being enhanced to support longer parallel Sysplex distances (up to 3500 km), customized analysis and scoring, and quicker drill-downs to isolate problems. zAware is an innovative feature of the zBC12 and zEC12 that incorporates sophisticated predictive analytics in a self-learning, expert solution that analyzes messages in near real-time to provide insights into the behavior of the z/OS system(s) that it monitors. Through a browser-based GUI interface, operations management can monitor and proactively address issues that could potentially impact continued operations. zAware is implemented on as a unique virtual appliance that operates in a dedicated zAware Mode logical partition, defined on a CP or an IFL.

Flash Express offers reduced access latency, resulting in higher IOPS (I/O operations per second). This feature incorporates storage class memory using SSDs on a PCIe I/O card and, by doing so, delivers accelerated I/O (because it is faster than disk). Flash Express cards occupy a standard PCIe I/O drawer slot and deliver high availability using RAID 10 mirrored card pairs. It also provides 128-bit AES data encryption for added security. Flash Express now is enhanced to improve availability and provides dynamic reconfiguration. z/OS exploits these capabilities to improve system availability by lowering the time to provision, to respond to changing workloads, and to complete dumps. Thus, delays can be drastically reduced when data for the workload can be transferred quickly from Flash Express into main memory. In addition, larger (1 MB) *Pageable Large Pages* are supported as well, often resulting in improved performance without consuming main memory.

Exhibit 3 — Business Class Storage Options

DS8870 Feature/Function	Single Frame Business Class Configuration	Low-Cost Entry Option Configuration (New in June 2013)
Supports Mainframe and Open Data	Yes	Yes
Minimum Cores (controller processors)	2	2
Maximum Cores (Controller Processors)	2	16
Minimum Cache	16GBs	16GBs
Maximum Cache	32GBs	1024GBs
Minimum Cache for Additional, Extra Cost Services (like copy services, SSDs, or I/O Priority Manager)	32GBs	32GBs
Minimum Drives (SSD, Enterprise, Nearline)	16	16
Maximum Drives (SSD, Enterprise, Nearline)	144	1056
Maximum Host Adapters and Ports	4/32	16/128
Protocols Support	Fibre Channel and FICON	Fibre Channel and FICON
Maximum Channel Speed	8Gbps	8Gbps
Warranty	1-4 Years	1-4 Years
Upgradable to Enterprise Class Configurations	Yes	No

Note: In addition to the above CKD Storage Options, other IBM Storage options that can connect to the zBC12 are *XIV Storage System* (for z/VSE, z/VM and Linux on System z), *Storwize V7000* via *SAN Volume Controller (SVC)* (for z/VSE, z/VM and Linux on System z) and the new *IBM FlashSystem 820* via SAN Volume Controller (coupled with Enterprise Linux Server (ELS) offering (for a high-performance cloud solution.

Source: IBM

zEnterprise Hybrid Solutions are Further Enhanced

As expected, the *zEnterprise BladeCenter Extension (zBX) Model 003* introduced with the zEC12 is carried over for use with zBC12. The capacity remains the same, but there are a number of enhancements in the *zEnterprise Unified Resource Manager (zManager)* – improving the monitoring and management of the zBX's resources. Included are support for *Microsoft Windows 2012*, enhanced workload policy-based performance management for *AIX*, *Linux* on *System x*, and Microsoft Windows, and improved availability across the ensemble with alerts and notifications for blade virtual servers. Also included are extended capabilities for continuous availability and disaster recovery of zBX blades in GDPS configurations, and support for enhanced communications between zBX and intranets. In addition, there is a Statement of Direction for the implementation of *DataPower* virtual appliances on x86 blades. IBM also is stating that the zBX will be supported in the next generation of System z, which should cover the next three to four years or so; however, we should expect more parts commonality with IBM's PureFlex systems for the servers. With over 275 hybrid units (zBX and *IBM DB2 Analytics Accelerators*) shipped to date to IBM's global customers, the viability of the zEnterprise hybrid architecture offerings should be unquestioned.[5]

Comparisons to Prior Generations

There are significant differences compared to the last two generations of Business Class mainframes and these are summarized in Exhibit 2 at the top of the previous page. The zBC12 continues the form factor that was established with the z114[6] two years ago, that is, one or two Processor Drawers with two Single Chip Modules (SCM) chip carriers in a single air-cooled frame. However, the zBC12 has four or five active cores per SCM compared to three or four that are active on the z114. This combined with the cycle time, increases in the quantity and improvements in the use of cache memory, and other design improvements[7], yields a single CP (z/OS engine) rating of 1064 MIPS with up to 36% performance improvement per core compared to the z114 and up to 58% performance improvement per core compared to z10 BC. Total capacity is increased to approximately

[5] For more on zBX and hybridization, 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.

[6] For more details, see **The Clipper Group Navigator** entitled *IBM zEnterprise in the Midmarket – Revolution or Evolution*, dated July 12, 2011, and available at http://www.clipper.com/research/TCG2011024.pdf.

[7] The zBC12 chip also includes an *Integrated Firmware Processor (IFP)* which is used for infrastructure management of the zEDC Express and the 10GbE RoCE Express features.

4900 MIPS, 58% better than the z114 and a 78% improvement over the z10 BC.[8]

User-configurable cores increases to 13, from 10 on z114, providing greater flexibility in tuning the system to its workload mix and an increase to six CPs from five that can be assigned to z/OS workloads, adding to the zBC12's increased total capacity. Subcapacity granularity has been increased for improved "rightsizing" by increasing the subcapacity setting steps to 156, compared to 130 on the z114. In addition, main memory capacity has been doubled to 512 GBs (but 16 GB is reserved for HSA, the Hardware Storage Area).

In another bold and unexpected move, the System z designers have changed the configuration requirements for zIIPs and zAAPs. With prior mainframes, the configurations that included zIIPs or zAAPs required the presence of at least an equal number of CPs (standard z architecture, processors/cores) for the number of zIIPs and zAAPs (a 1:1 ratio). With this announcement, limited to the zBC12 and zEC12, the ratio has been increased to 2:1, i.e., only one CP will be required to match up to two zIIPs or zAAPs. This change has significant potential to impact the platform's workload economics[9], i.e., possibly by reducing z/OS workload and, thus, reducing or delaying the requirement for additional CPs to accommodate growth. It also suggests to me that IBM might move more specialized workloads to these processors, to help tilt the balance of the price-performance and TCO equations to favor zEnterprise.

In the same 2008 footprint, using essentially the same power and cooling energy, and software, the zBC12 is capable of supporting 77% greater workload with higher availability, improved security, no higher personnel costs, and even greater opportunities for consolidation and simplification of operations.

Business Class Storage Options

In line with Systems z's Business Class offerings is the need for appropriate and affordable storage to go with it. IBM does not disappoint here either, with its DS8870 Entry Class and Business Class options. Both offer lower price points with a reduction in the number of device (RAID) adapters and I/O enclosures to support a given number of drives, thus reducing configuration cost and increasing adapter usage. Check out the configurations in Exhibit 3, at the top of the previous page. Bottom line – System z Business Class users now can get started with an entry-level IBM enterprise-class storage solution at a reasonable price.

Exclusive Innovations to the Current Family

Mentioned earlier was the fact that the occasion of bringing the midrange system to the market also provided an opportunity to introduce completely new features and architectural enhancements across the current family, in this case the zBC12 and the zEC12. The first of these is a new type of data compression, *zEnterprise Data Compression (zEDC)*, which is being introduced exclusively on the zBC12 and the zEC12, with software support by z/OS V2.1. zEDC is a very efficient compression algorithm, up to four times compression ratio, designed to deliver high-performance, low-latency compression. As a result of the compression workload being offloaded to a PCIe I/O feature, the *zEDC Express for zEC12 and zBC12*, there is very low CPU overhead. Use of this new PCIe I/O-based compression facility reduces latency and thus drives better efficiency for cross-platform transfers, storage devices, and network bandwidth. Potentially, it could be used for files not benefiting from on-chip compression. The first use case will be compression before writing SMF log data, which can be quite voluminous. Statements of Direction[10] for this feature include future provision for *BSAM/QSAM* extended format sequential data sets, JAVA data streams, as well as data managed by *DFSMS*.

Also being introduced for the zBC12 and zEC12 exclusively with z/OS V2.1, is a new communications feature, *SMC-R (Shared Memory Communications - RDMA*). SMC-R is similar to *HiperSockets* in that it provides fast memory-to-memory transfers between z/OS processors or LPARS; it can provide significantly improved network performance – often

[8] Based on IBM internal measurements and projections and compared to the z114 and/or z10 BC. This is roughly 50% faster than a z900 (circa 2002). Results may vary by customer based on individual workload, configuration and software levels. For more details, see: https://www-304.ibm.com/servers/resourcelink/lib03060.nsf/pages/lsprindex?OpenDocument.

[9] Because processor cores configured as either a zIIP or zAAP are priced as a fraction of a standard processor. See more detail in the pricing discussion that follows.

[10] IBM Statements of Direction are subject to change or withdrawal without notice, and represent goals and objectives only.

characterized as latency, throughput, CPU consumption, and scalability. It can reduce TCP/IP overhead and network latency up to 80% and, therefore, drives more efficient system utilization. It also has the added benefit of being application transparent which means it can be exploited immediately to achieve these gains. The SMC-R is a new sockets-over-*RDMA (Remote Direct Memory Access*) communications protocol that allows existing sockets applications that exploit TCP to transparently benefit from RDMA for exchanging data over a *RoCE* (pronounced "rocky", which stands for *RDMA over Converged Ethernet),* which is being adopted widely within the IT community. In addition to its application transparency, it leverages existing Ethernet infrastructure and preserves the security and operational model of z/OS TCP/IP sockets. The feature is available immediately for the *IBM WebSphere Application Server (WAS),* IBM *CICS TS,* and others. Exploitation of SMC-R networking requires at least two *10GbE RoCE Express for zEC12 and zBC12* features.

Enterprise Linux Server Reintroduced

The concept of an *Enterprise Linux Server (ELS)* on System z certainly is not new. The ecosystem has now grown to over 3000 ISVs that provide applications for that environment. The first models of a System z that were offered with *IFL*-only configurations date back to January 2002.[11] For the purposes of this discussion I am not including configurations that include standard System z architecture processors, commonly referred to as *CPs*, running with the z/OS or z/VSE operating systems.[12] The announcement of the zBC12-based ELS brings significantly greater scale and improved price-performance to this solution. Up to 13 processors (cores) may be configured as IFLs supporting *Linux for System z.* Up to 40 virtual servers per processor can be deployed yielding up to 520 Linux server images in a single footprint, which is workload and configuration dependent.

The computing capacity is balanced with up to 496 GB of main memory. The integrity of the Linux servers are protected by the exceptional reliability of System z processors and the isolation provided by Logical Partitioning to provide the highest levels of availability and security. *Red Hat and SUSE* Linux distributions are supported. IBM is offering a pre-packaged Enterprise Linux Server solution that includes all necessary hardware, hypervisor, and maintenance[13], with a total cost of ownership that is promised to be attractive when compared to *x86*-based distributed servers with a typical software stack.

To position the significance of this announcement further, a zBC12-based ELS represents a 62% increase in Linux workload capacity when compared to the z114. This is a result of both the zBC12's faster cores and its three additional user-configurable cores.

Another exciting possibility is presented here by considering a hybrid solution with the ELS supplemented to include a zBX that can bring not only *zLinux*-enabled workloads into one frame but also x86-based Linux workloads[14] into a tightly coupled and common management structure and discipline with the zManager.[15] This hybrid Linux consolidation solution will offer significant savings in personnel, software, environmental costs, and standardization that will offer mainframe level qualities of service.

System Software and Middleware Leverages the New Technologies

In support of the zBC12 announcement, there is a new version of z/OS V2.1[16] that strengthens high performance data serving, batch processing, and business analytics, a new release of *z/VM V6.3*, as well as the aforementioned updates for Linux on System z from Red Hat, SUSE, and OpenStack. The other supporting operating systems include *z/VSE V5.1* and *z/TPF V1.1*, which now will have added compatibility code (but there are no new features at this time).

z/VM V6.3

The new release of z/VM V6.3 includes focus on the Enterprise Class with enhancements that provide improved performance for large scale resources. The number of virtual servers that can be supported is increased with the four

[11] The IFL, *Integrated Facility for Linux*, is the special-purpose processor that hosts *Linux on System z.* It does not execute standard System z programs. Its origins date back to 1999. The first IFL-only mainframe was the z800 Model 0LF.

[12] Clearly, a System z server running a standard operating system along with Linux for System z programs can be an "enterprise Linux server", that is, a consolidation point for all or most Linux workloads.

[13] The Linux stack must be obtained from a qualified distributor, who in turn provides program support.

[14] Those workloads that have not been enabled and certified for Linux on System z.

[15] Linux for System z virtual images are managed by z/VM, not zManager.

[16] See footnote #1.

times expansion of real memory support from 256 MB to 1 TB. Additionally, *HiperDispatch*[17] now is supported, which will result in much greater efficiency and improved performance for z/VM in large-scale deployments. In addition, support for *OpenStack* enables z/VM to be managed with open cloud architecture-based interfaces.

DB2

DB2 for z/OS performance improvements will result from deeper integration with zEnterprise architecture – for example, by using 1 MB Pageable Large Pages on Flash Express and also by leveraging use of 2 GB fixed pages.

IMS

Users of *IMS 13* will benefit from the extended addressability provided by the Long Displacement Facility in the System z architecture and stand to gain from the newly introduced SMC-R, Shared Memory Communications, which dramatically improves performance of cross-system TCP/IP communications. IMS JAVA application developers may take advantage of Flash Express for 1 MB Pageable Large Pages, and the Transactional-Execution facility, which can be used by C and C++ applications as well, helps to eliminate soft-ware locking overhead and drives higher transaction throughput.

Pricing the New Technologies

From generation to generation, IBM's mainframes consistently have delivered improved price/performance for the "stack", usually in the range of 15-20% when using a "stack" model. The "stack" price/performance (usually expressed in $/MIPS) is the sum of the improvements in the pricing of the hardware, a "typical" configuration of Monthly License Charge (MLC) software (e.g., z/OS, DB2 for z/OS, etc.) and monthly maintenance charges on the hardware following the first-year warranty period.

In the case of the zBC12, the improvements come in two flavors. First, notice (see Exhibit 2 on Page 3) that the Entry MIPS for the zBC12 is almost twice that of the z114, 50 MIPS and 26 MIPS, respectively. However, this entry configuration[18] will be priced the same, $75K, a 92% improvement in terms of $/MIPS. Second, at higher capacities customers can expect to be offered a 5-to-10% improvement in $/MIPS.[19]

One aspect of IBM's continued dedication toward protecting their customer's prior investments is exemplified in their pricing practices with regard to processor upgrades. When upgrading, say from a z114 to a zBC12, the price would be based primarily on the incremental capacity being purchased, effectively giving credit for the all the capacity of the prior processor. "Street" prices are available for special purpose engines and they remain the same as those offered on the z114, $35K for an IFL and $40K for *zIIPs* and *zAAPs*. However, since the capacity of the IFL has been increased by about 36% and the Price-Value-Unit (PVU) remains at 100, the result is in effect and price-performance improvement of 27% for a zBC12 IFL (over a z114).

In the case of MLC software, IBM follows through with its practice of providing *Technology Improvement Pricing*[20], in this case 5%. This would result in *reducing* the monthly license cost for those products by 5% for flat-capacity upgrades from the z114; about 11% if upgrading from a z10 BC. Maintenance price-performance for the standard processor is improved by 2%, but 27% for an IFL, in keeping with its flat price for its increased capacity.

Special features get special attention and a new mode of pricing is being introduced for the zEDC. The price has been split between the hardware component, $12K per card, and a charge for the z/OS-enabling software, which will be variable, based essentially on usage. There is some advantage gained here when installed on the smaller scale zBC12, which by its nature, have lower software charges. It remains to be seen whether IBM will extend this practice to similar features in the future. The 10GbE RoCE card will be priced at $11.5K per card for both the zBC12 and the zE12. Flash Express is quoted as $125K per 1.6 TB mirrored card pair; a maximum of four may be installed. The IBM zAware feature will be made available in smaller increments and lower entry price for the zBC12 than on the zEC12, which is quoted

[17] HiperDispatch, exclusive to zEC12, z196, and z10 EC, keeps processes running near their cached instructions and data thus minimizes the need for transfers of data ownership among processors and processor books (groups).

[18] The reference configuration includes 1 standard core, 8 GB main memory, 4 Ficon Ports, and 4 OSA ports. No software, storage or networking is included.

[19] IBM does not publish list prices for zEnterprise hardware (except for certain features) so no specific percentage improvement has been provided.

[20] In the past this was known internally as the "Tech Dividend" as a means of providing price-performance improvement on MLC software as an incentive for customers to adopt the latest technology systems.

as $40K for up to 10 CPs[21]. On the zBC12, a 2-CP pack will cost $8K, a 4-CP pack, $16K, and a 6-CP pack, $24K.

In addition, IBM promises to be aggressive in pricing cross-industry Solution Editions particularly when competing for new applications and workloads; for example, the Enterprise Linux Server Solution for Analytics.

Conclusion

IBM mainframe watchers in the IT universe tend to make a much bigger fuss over the largest member of the family; last year, it was the zEC12. *Why is that a surprise?* The Enterprise Class always leads the announcement of the new family, has by far the highest capacity, and new under-the-cover technologies are introduced there. The largest enterprises and those most well-known always are the early adopters.

However, it would be a mistake to dismiss the zBC12 midrange/midmarket system as much less relevant to those who currently have System z technology or are considering it for their enterprises. The zBC12 is no less capable, in terms of its capabilities, to address – cost effectively – the new technological demands on IT infrastructures while maintaining the accelerated pace of enterprise activities. IBM does not disclose its sales volumes relative to the mix of Business Class versus Enterprise Class, but the market opportunity for the smaller system is much larger than the high-end (in terms of units), and I would estimate that the actual full life cycle unit sales to be roughly equal between the two.

The Business Class mainframe provides the lowest entry point to the many advantages of System z and there are many use cases for including them in an IT infrastructure that already includes an Enterprise Class system. For example, this might be as a backup or disaster recovery system for a limited number of critical business processes, to support a remote location, to provide a host for specialized services, or to supplement a group of systems in a Parallel Sysplex configuration, and possibly many more. For those data centers who were early adopters of the z10 BC, who may now be very close to the end of their operating leases, in need of the capacity bump, and/or the more advanced technology, the zBC12 should be an extremely compelling offering.

Limited only in its total capacity, the current level of the z architecture also is realized there and therefore is interchangeable with the zEC12. **The zEnterprise BC12 announced today offers a well-balanced collection of improvements for existing mainframe customers and even more reasons for a new customer to consider System z for critical workloads.** It deserves your consideration, so take a long and hard look at it.

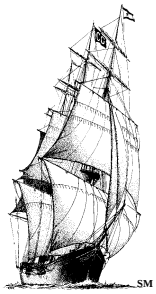

[21] Additional features must be purchased for the zEC12 to cover additional CPs, in multiples of 10.

About The Clipper Group, Inc.

The Clipper Group, Inc., now in its twenty-first year, is an independent publishing and 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, The Clipper Group Calculator, 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.