THE CLIPPER GROUP
Navigator™
Navigating Information Technology Horizons℠
Published Since 1993 Report #TCG2012019 August 28, 2012

The IBM zEnterprise EC12 — Bigger, Better, Faster

Analyst: Stephen D. Bartlett

Management Summary

Think "skyscraper" and the never-ending desire to be the tallest. While there are bragging rights, of course, afforded to the tallest, it also may represent the maximum utilization of the real estate and location, i.e., more square feet of office/residential space on a given footprint. However, once built, it doesn't really have the ability to grow larger or have the ability to roll over the sunk investment in "infrastructure" into something larger within the same physical footprint. Not so with the mainframe – you get more with each generation – always more MIPS; faster, more capable I/O connectivity; and more functionality to enhance designed-from-inception reliability, availability, and serviceability features, and usually in less space and almost always without having to shut down for the transition. Occasionally, IBM adds some things that were not quite expected – for instance the extension of the architecture to provide hybrid solution through the addition in July 2010 of the *zEnterprise BladeCenter Extension (zBX)*, facilitated by the *zEnterprise Unified Resource Manager (*a.k.a. *zManager)*. And now, in August 2012, just 25 months later, IBM brings its next enlargeable "skyscraper" into the world, the *zEnterprise EC12 (zEC12)*. (See picture at right.)

Exhibit 1 — IBM's zEC12

Source: IBM

What's in a Name?

IBM has abandoned the prior naming scheme of the prior generation, the *z196* (first generation hybrid, 96 cores; often called "z11" because it was the eleventh generation of *System z*), and returned to that which was established by the z10 EC, with a swizzle. Now we all can relax, as order has been restored. The focus of this announcement is on central processor scale, granularity, and new and improved optimization functions. We should emphasize here that there is no backing off the long-term strategy of *System z* hybridization. IBM declares that System z is and will continue to be a hybrid system that incorporates and centrally manages System z, *Power Systems*, and x86 (*System x*) architectures as a unified structure. The timing of this announcement probably catches many IBM observers, certainly The Clipper Group, and mainframe

IN THIS ISSUE

customers a bit by surprise. In the past, the refresh cycle for the largest mainframe has been in the 27 to 32 month range. Moreover, many of IBM's most loyal mainframe customers are on 36 to 48 monthly technology refresh cycles. For those who installed z196's with typical operating leases, this means that they are likely to still be on the front half of their deals. (I suspect that IBM will have a financial solution that still can make the zEC12 attractive.) There is, however, a sweet spot that will be tapped as there are many z10 ECs that clearly are ready for a refresh, which we will explore in more depth in this paper. However, at this time, there is less to be said in terms of either enhancements or functional extensions the zBX and *zManager*.

Exhibit 2 — The zEC12 Processor Unit Chip

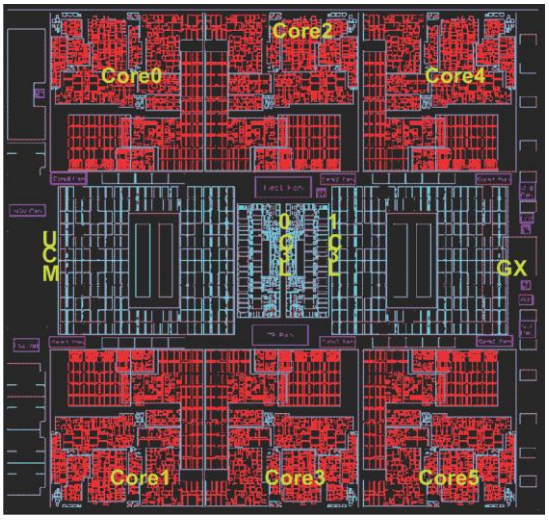

Source: IBM

So why now? Because it was ready; technical and development goals had been reached and there was new function demanded by customers. Mainframe workloads are growing at double-digit rates all over the world and its markets are expanding. *How is IBM achieving the new levels of performance? How have they enhanced the reliability and availability of System z? In the climate of stealth viruses and international cyber-theft, how have they helped make important systems more secure? How have they improved the value equation for my business that continues to reinforce the essential nature of System z?* These questions and more will be addressed. If you are curious, then please read on.

How Do You Build the Tallest Skyscraper?

Of course, the answer is slowly and carefully and with much forethought. In the case of the zEC12 – carefully – but as mentioned above, the recent pace was not casual. Though the story always starts with the processor chip technology, that is just the beginning of the story. (See annotated illustration in Exhibit 2, at the top of this page. The zEC12 processor chip is 598 mm^2.)

For zEC12, the instruction cycle time is ratcheted up a notch, by 6%, to 5.5 GHz, a modest improvement but still the fastest processor in the market by far.[1] (See Exhibit 3, a table on the next page, which includes this comparative data plus much of what is discussed herein.) The highly-integrated processor design includes doubling the cache memory on each chip to 48 MBs and doubling the total cache on each MCM (book) to 384 MBs. Design enhancements include second generation out-of-order execution and multi-level branch predication. As one would expect, the chip continues to include cryptographic and data compression functions, which will benefit from the increase cycle speed and cache enhancements. Hardware functions have been added to specifically address often observed software bottlenecks:

- The *Transaction Execution Facility* improves parallelism and scalability by exploiting transactional memory. This is a new architecture that improves performance by enhancing concurrency of multi-threaded, multi-processor programs; initially it will be exploited by *Java*.
- The *Runtime Instrumentation Facility* uses heuristics to help reduce Java workload overhead.
- 2 GB Page Frames improves performance for *DB2 for z/OS* buffer pools and Java heaps.
- A new set of instructions have been incorporated to enable broader exploitation of Decimal Floating Point (DFP) arithmetic, soon to be exploited by a new PL/I compiler for a performance boost.

These incremental enhancements are on top of the other familiar aspects of System z architecture's design excellence, to name a few:

[1] By comparison, Intel *Xeon* chips are running below 3.5 GHz, with the high end (*E7*) in the range of 1.73 to 2.66 GHz.

Exhibit 3 — Comparing zEC12 to Its z196 and z10 EC Predecessors

Name	zEC12	z196	z10 EC	zEC12 Compared to z196	Compare zEC12 to z10 EC
Type	2827	2817	2097		
Models	H20, H43, H66, H88, HA1	M15, M32, M49, M66, M80	E12, E26, E40, E56, E54		
Chip Process	32nm	45nm	65nm	n/a	n/a
Cycle Rate (GHz)	5.5	5.2	4.4	106%	125%
Entry MIPS	240	240	218	same	110%
Single CP MIPS - See Note (a)	1500	1200	920	125%	163%
Total CP MIPS	78,000 approx.	50,000 approx.	30,000 approx.	156%	260%
Subcapacity Settings	60 (3 x 20 CPs)<= 13.7 kMIPS	45 (3 x 15 CPs)<= 8.9 kMIPS	36 (3 x 12 CPs)<= 5.8 kMIPS	154%	236%
Books (MCMs)	1 to 4	1 to 4	1 to 4	same	same
Chips / MCM	6	6	5	same	120%
Cores / Chip	6	4	4	150%	150%
Total Active Cores	120	96	80	125%	150%
User Configurable Cores	101	80	64	126%	158%
Other (SAP, Spare, Reserved)	19	16	16	118%	118%
Cache / Chip (L1.5 in MBs)			3	n/a	n/a
Cache / Chip (L2 in MBs)	2	1.5		133%	n/a
Cache / Chip (L3 in MBs)	48	24		200%	n/a
Cache / MCM (L4 in MBs)	384	196	48	196%	800%
Entry Main Memory (in GBs)	32 (RAIM)	32 (RAIM)	64	same	50%
Maximum Main Memory (in TBs)	3 (RAIM)	3 (RAIM)	1.5	same	200%
LPARS maximum	60	60	60	same	same
z/OS CP maximum	100	64	64	156%	156%
Upgrades To	zNext (TBD)	zEC12	z196, zEC12		
zBX Model	2458-003	2458-002	2458-001		
Blade Types	PS701, HS5	PS701, HS5	See Note (b)		
Blade Max - See Note (c)	0 - 112, 0 - 56	0 - 112, 0 - 56	N/A		
Footprint /Power Envelope	Same as z196 approx. See Note (d)	Same as z10 EC approx.	Baseline		

Core Configuration: All models can include CP, ICF, IFL, zAAP, zIIP, and additional SAP (optional)

Notes

(a) CP = Core configured for System z architecture

(b) IBM Smart Analytics Analyzer V1.0 only

(c) Max Blades = 112, any mix

(d) Some cooling and cabling options extend dimensions to a small degree

Source: IBM, with computations by Clipper

- Maintains the highest utilization (100% is attainable) with heterogeneous workloads through dynamic resource allocation and workload dispatching and balancing.
- Superior virtualization through a two-tiered approach of *PR/SM* and *z/VM*, with guaranteed isolation and integrity.
- The availability of compilers and middleware optimized to fully exploit the enhanced machine architecture.

Scale is addressed not only by a faster, more efficient chip but now includes six cores, versus four for the z196.[2] The result is a total of 120 cores[3], up from 96, 101 of which are user-configurable as CPs, ICFs, IFLs, zAAPs, zIIPs, or optional additional SAPs, which are discussed further on page 7. In support of this additional capacity, *z/OS* has been extended to a single-system image of up to 100 logical CPs. With the zEC12, IBM now claims a single engine (core) capacity exceeding 1,500 MIPS and a total capacity of approximately 78,000 MIPS.[4]

In real world terms, when compared to a similarly configured z196, the sum of all the parts of the zEC12 is capable of delivering:

[2] For more detail on the z196, see **The Clipper Group Navigator** entitled *The IBM zEnterprise System Reaches Out - Higher, Wider and Deeper*, dated July 22, 2010, and available at http://www.clipper.com/research/TCG2010033.pdf.

[3] All MCMs are capped with a maximum of 30 cores.

[4] IBM internal studies have been used to estimate performance prior to general availability. Relative capacity and performance comparisons are made at equal software levels as measured by IBM Large System Performance Reference (LSPR) workloads using z/OS V1.13. Results may vary by customer based on individual workload, configuration and software levels. See the LSPR website for more details at https://www-304.ibm.com/servers/resourcelink/lib03060.nsf.

Exhibit 4 — IBM zEC12 Delivers Stellar and Well-Balanced Performance

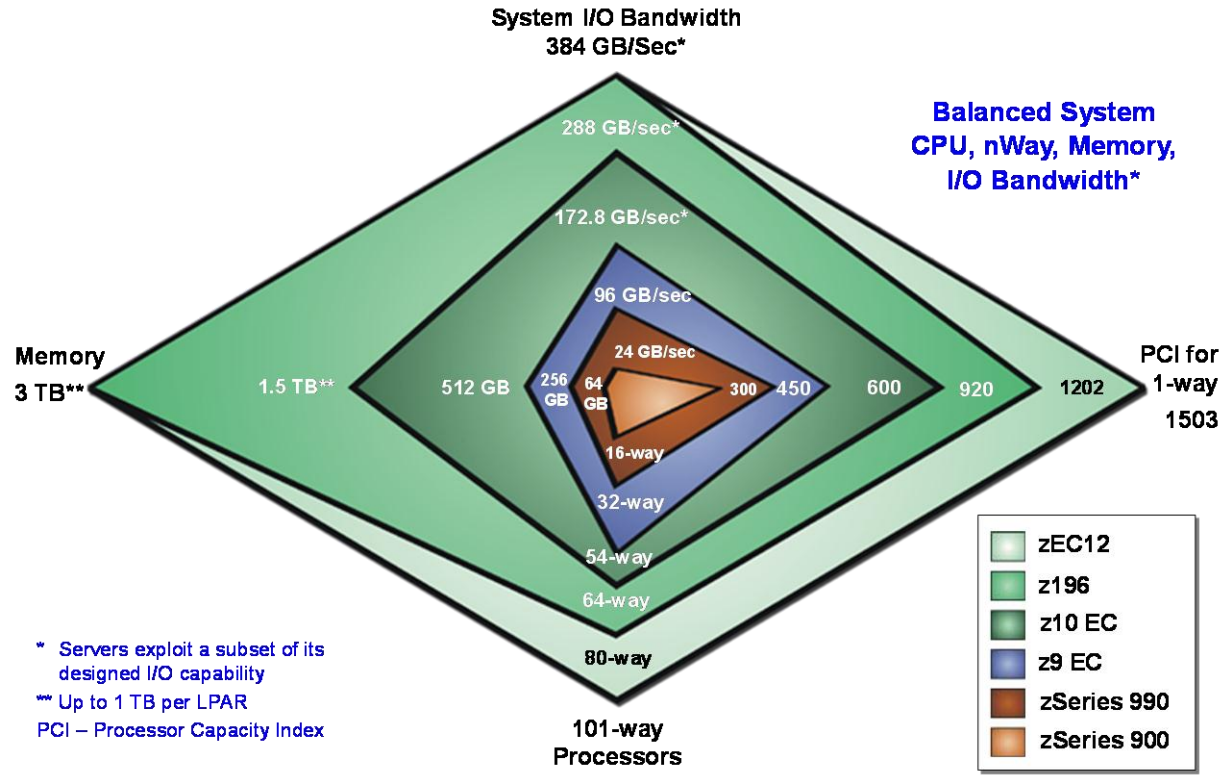

Source: IBM

- Up to 45% improvement for Java workloads.
- Up to 30% improvement in throughput for DB2 for z/OS operational analytics.
- More than 30% improvement in throughput for SAP workloads.
- Up to 30% improvement in IMS throughput owing to faster CPU, larger cache, and improved compilers.
- Up to 27% improvement in compute-intensive C/C++ applications and a 31% average improvement for PL/I V3.4.

All of this results in a very powerful and well-balanced offering, as shown by the chart in Exhibit 4, above. Do note that the four axes represent different aspects of system design and have different scales. This chart shows how, over several generations, IBM has grown performance in these dimensions – often by differing technological means. In summary, the zEC12 has been tuned to a complicated objective function (to be well-balanced for the most demanding workloads) and this has been achieved by a lot more than just growing the processor clock speed to world-class standing (which, by itself, is worthy of accolades).

General availability for the zEnterprise EC12 will be on September 19, 2012. Model upgrades from z10 EC and z196 also will begin on this date.

New Innovations in zEnterprise Availability

IT organizations throughout business and the public sector are being driven to delivery higher and higher service levels. The availability characteristics of System z are well-known as one of its essential foundational elements, but the engineers at IBM have deemed that the prior status quo was not sufficient to continue to deliver unchallengeable world-class leadership. Two new features have been added exclusively to the zEC12 and, interestingly, approach the issue from two very different vantage points.

Flash Express

The first is *Flash Express*. By its name, System z cognoscenti will quickly infer that this feature is an I/O card that incorporates flash memory and doing so it delivers a new, transparent acceleration tier between main memory and storage class memory. (See Exhibit 5, at the top of the next page.) Flash Express cards use a standard PCIe I/O drawer slot and incorporate SSDs as a RAID 10 mirrored card pair. The

Exhibit 5 — Flash Express Card

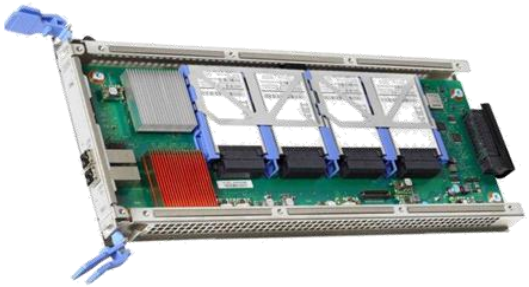

Source: IBM

SSDs are sized to accommodate *paging for all (or multiple) LPARs*, with each pair providing 1.6 TB of usable mirrored storage. A maximum configuration, 4 card pairs, yields 6.4 TB of redundant memory. For security, data is encrypted on the Flash Express adapter using the 128-bit AES algorithm. Its implementation in the system is very simple as it is designed for concurrent replacement or upgrade, requires no capacity planning or time-consuming intelligent data placement, and may be divided into a virtualized resource across several partitions.

The benefits of Flash Express to improved availability and reduced latency are tangible and achievable immediately and invisibly. CPU performance benefits from the use of pageable, large 1 MB pages and DB2 and Java workloads are likely to achieve several percent in improved performance. Applications with random read access or high read/write ratios also will see improvement. Availability will be improved during the workload transition periods, as when the workday begins, and when spikes occur. And, when unfortunately something breaks and a standalone or SVC dump is required, Flash Express accelerates the writing of snapshots, which improves paging performance because of fewer disruptions from diagnostics.

Development efforts continue in this area to enable expanded use and exploitation by middleware and application programs. Besides DB2 for z/OS and the Java SDK, expect other large page exploiters to leverage this feature including *WebSphere Application Server*, *IMS*, *CICS*, and the z/OS V1.3 Language Environment. Flash Express is a separately-priced hardware feature and will be available at general availability on September 19, 2012. z/OS exploitation of the functionality will be delivered via PTFs on December 14, 2012.

IBM zAware

Where Flash Express contributes to enhanced availability and improved performance as a part of the memory hierarchy, *IBM zAware* employs a very innovative and unique approach to availability improvement. IBM zAware is a self-learning, integrated expert solution that analyzes messages in near real-time to provide insight into the behavior of the z/OS systems that it monitors. It incorporates sophisticated analytics and applies expert knowledge and machine learning to alert operations managers of any anomalous conditions that may lead to an interruption of services. In the flood of information that engulfs the management console, IBM zAware identifies issues quicker and more efficiently and it learns and adapts to its unique environment. IBM zAware reduces impacts to business and service levels by minimizing or even eliminating the need for additional diagnostic information, thereby shortening the time to repair.

In simplified terms, IBM zAware frequently samples the z/OS OPERLOG and assigns messages an anomaly score based on such characteristics such as rarity, context, or abnormally high frequency or other unusual patterns (and this list is by no means all inclusive). Performing near real-time analysis, it provides, through a browser interface, high-level and drill-down reporting tools. Quickly, an operations manager can identify conditions that require the highest priority for their attention and remedial actions. The flip side is that, in the absence of detectable anomalous conditions, IT management can gain some comfort that all is well and they now can focus on less urgent issues.

Implementation is straightforward and takes a very modest amount of resources – a base of 4 GB of memory, 500 GB of storage, and about 1% of processor overhead on monitored systems. IBM zAware will run on a general purpose CP, but an IFL-based installation is the obvious choice. The host, a zEC12 server can monitor any other System z server with the appropriate connectivity, and requires z/OS V1.13 + PTFs. It is a separately-priced hardware feature and will be generally available on September 19, 2012.

The Enterprise Standard for Security is System z

Another essential foundational element of the System z architecture and its implementation

is security. Beginning at Day 1, almost five decades ago, storage and fetch protection, storage protection keys, authorized program facility (APF), cross-memory communication management, the system authorization facility (SAF), and the IBM Security Server furnish the bedrock. Cryptographic processing was introduced into the System z architecture via a coprocessor in the early 1990s and continues to establish and maintain System z's leadership in this domain. The financial systems that bind the world's commerce rely on it and, in turn, continue to drive future requirements and development efforts.

Hardware cryptography is built into each general-purpose CP and IFL, and via the new *Crypto Express4S coprocessor PCIe* I/O card that, as in prior generations, incorporates tamper-resistant packaging. Together, they support secure data serving with protection of data whether "at rest" or "in flight". The Crypto Express4S coprocessor is designed to meet FIPS *140-2 Level 4* requirements[5], the most stringent. It has three configuration options:

- As a *Secure Socket Layer (SSL)* accelerator and clear key RSA operations;
- As the host for IBM's *Common Cryptographic Architecture (CCA)* with the *Trusted Key Entry (TKE)* option, which provides flexible key entry and remote key management; or
- As the new *IBM Enterprise PKCS #11 (EP11) coprocessor*.

The IBM Enterprise PKCS#11 firmware has been designed to meet both FIPS 140-2 Level 4 and Common Criteria (EAL 5+) standards. In particular, the zEC12 meets open standards for digital signatures with the new PKCS #11 support. The new Crypto Express4S coprocessor option targets the public sector, where industry standard services are required and certifications are tailored to meet requirements of this market place. It is becoming a mandate by the European Union for high-quality electronic signatures (trusted to the same extent as handwritten signatures), e.g., for smart passports, national identity cards, or other identity-certification applications.

In addition, the CCA is being extended to support the payment card industry with enhancements to support the *Europay*, *MasterCard*, and *Visa* (EMV) global standard for credit and debit cards that incorporate an integrated chip that allows them to perform cryptographic computations, such as sharing keys, signature authentication, and PIN encryption.

Supporting the "Green" Datacenter

With the introduction of the zEC12, there are now *two cooling options* available, providing additional flexibility to meet the particular needs of customers' data centers.

- **Introduced with the z196, there is an external water-cooled option**, which is 4,000 times more efficient than air. It incorporates two air-to-water heat exchanges within the rear of two frames and adds four inches in depth. This is an IBM-maintained closed-loop system, with appropriate redundancies built in, that requires connection to the datacenter's chilled-water system.
- **The newest cooling option, introduced with the zEC12 announcement, is radiator-based air cooling.** The design is essentially a set of water-filled, closed-loop radiators built into the rear panel of the processor frame. This is more efficient then standard air-cooling but less complex and less expensive than the water-cooled option, although also less efficient, but perhaps a good trade-off in many cases.

Also, a non-raised floor option is now available. With this option, a provision has been made to route all power and signal cabling from overhead runners, a feature common in many new datacenters.

zEC12 Continues the Evolution of the zEnterprise Hybrid System

With each conversation with IBM, we at The Clipper Group continue to be convinced that IBM's model for all future mainframe development is focused fully on its evolution as a hybrid system. And by this we mean tight integration of multiple (different) processor architectures (and by implication, multiple operating systems and hypervisors) central managed and controlled as a single unit. *One system, one window* – if you will pardon the expression. And lest we forget (or for the newly-initiated), hybridization of IBM's mainframe architecture first was introduced on a prototype IFL as the host processor for *Linux* in December, 1999, on the *System/390 9672-G6*.

The new extensions to hybridization of zEnterprise (select IBM *System x and Power Systems* blades) are being embraced by enterprise

[5] Federal Information Processing Standards (FIPS) are U.S. government computer security standards that specify requirements for cryptography modules. FIPS 140-2 Level 4 makes the physical security requirements more stringent, and requires robustness against environmental attacks.

customers, at a rate exceeding that for Linux on System z at the same point since its release, but perhaps at a slower pace than IBM mainframe marketers may have hoped. I offer what I believe to be the two most important contributing factors.

1. This was a radical departure from the course which IBM's mainframe customers have become accustomed; it would take some time to understand, evaluate, and fit it into their computing environments; and
2. It lacked a most important element when first announced – the delivery of support for the most ubiquitous of the distributed server platforms, *Windows* on x86 servers.

Now, in a little more than two years after zEnterprise BladeCenter Extension (zBX) was announced, these barriers have been removed and the hybrid model in now in its early growth stage. IBM claims over 150 zBX units sold worldwide, populated by more than 1100 blades integrated with both z196 and z114.

It is accurate to state that the zEC12 announcement is weighted more heavily toward the new zEC12 processor, but there are new enhancements to the zBX frames and the Unified Resource Manager, as well. First, a new model of the zBX, Model 003, has been introduced to support the zEC12. With the new Model 003, the bandwidth of the switch connection from the BladeCenter chassis, located within the zBX frame, with the Top-of-Rack (TOR) switch has been doubled to 20 GBe. This enhancement will provide improved band-width to those with larger POWER7 or System x blade counts. No changes have been made to the *Intranode Management Network (INMN)* or the *Intraensemble Data Network (IEDN).* Also, no changes have been made to the zBX blade options or configurations. There is no charge for upgrades of the zBX Model 002 (introduced with the z196 in July of 2010) to the Model 003 or for transferring existing blades to the zEC12 for management by zManager, including the zManager entitlements.[6]

Second, in the management firmware and software arena, new IBM *System Director* capabilities are planned to be added through the zManager's APIs. They include the ability to support virtual server provisioning and image management, energy management functions, and power capping of blades.

Third, IBM has announced its intent to deliver workload-aware optimization for System x blades through functional enhancements to zManager. This allows virtual CPU capacity to be adjusted automatically across virtual servers within a hypervisor, helping to insure that System x resources in the zBX are executing to the defined SLAs. This capability currently is provided for Power System blades, in support of the z196/zBX Model 002.

zEC12 Pricing Promotes Growth

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.

Hardware Pricing

In the case of the zEC12, when compared to the z196, customers can expect to be offered a significantly lower $/MIPS. However, as IBM does not publish list prices for zEnterprise hardware (except for certain features), no specific percentage improvement has been provided, but likely something in the range of 10% is realistic. One aspect of IBM's continued dedication toward protecting their customer's prior investment is exemplified in their pricing practices, which particularly is well illustrated with processor upgrades. Additional costs are based primarily on the incremental capacity being purchased, effectively giving credit for the all the capacity of the prior processor. And, as mentioned earlier, there is no charge for zBX Model 002 upgrades to the Model 003, or for transferring existing blades to zManager on the zEC12.

Specialty engine prices, which are unit priced – not capacity priced like the standard CP engine – remain the same and are quoted as

[6] For additional information on the zBX and zEnterprise hybrid environments see **The Clipper Group Navigator** entitled *The IBM zEnterprise Stretches Its Boundaries – New Windows are Opened*, dated October 12, 2011, and available at http://www.clipper.com/research/TCG2011034.pdf and **The Clipper Group Navigator** entitled *The zEnterprise Hybrid Model – Where Goest Thou?, dated August 28, 2012*, and available at http://www.clipper.com/research/TCG20120nn.pdf.

Exhibit 6 — Specialty Engines Make System z Special

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

The Processor Units (PU), or each of the 101 zEC12 user cores, may be configured as:

- **ICF = Internal Coupling Facility**; used for z/OS clustering.
- **IFL = Integrated Facility for Linux**; exploited by Linux and for z/VM processing in support of Linux.
- **zAAP = zApplication Assist Processor**; exploited under z/OS for designated workloads such as the IBM JVM and some XML System Services functions.
- **zIIP = zIntegrated Information Processor**; exploited under z/OS for designated workloads, such as various XML System Services, IPSec offload, certain parts of DB2 DRDA, star schema, *HiperSockets* for large messages, and the IBM GBS Scalable Architecture for Financial Reporting.
- **SAP = System Assist Processor**; offloads and manages I/O operations. Several SAPs are standard with the zEC12; more may be configured if additional I/O processing capacity is needed.

These options provide superior flexibility to meet user workload requirements, and because they are priced significantly lower than standard mainframe CPs and are not included in the MIPS capacity for software monthly licensing charges, might significantly lower the cost of computing by reducing the number CPs required. For all of the above, some configuration restrictions may apply.

Several other zIIP and zAPP happenings are worth noting.

- IBM continues to support running zAAP workloads on zIIP processors ("zAAP on zIIP"). IBM plans to provide a PTF in September 2012 to remove the restriction that prevents zAAP-eligible workloads from running on zIIP processors when a zAAP is installed on the server. This is designed to help facilitate migration and testing of zAAP workloads on zIIP processors.
- zEC12 is planned to be the last high-end System z server to offer support for zAAP specialty engine processors. IBM plans to converge zAAP and zIIP specialty engines onto zIIPs. This is intended to help simplify capacity planning and performance management, while still providing the capability to offload many workloads, which can help control software costs.

Source: IBM

$55K for the IFL and $100K for zAAP and zIIP. By virtue of the increased performance of zEC12 processor cores, specialty engines will get a 20% price-performance improvement with no increase in price per engine. Specialty engines may be carried forward to the new processor, but there is a 30% carry-forward charge rather than having to re-buy at full price, *effectively giving the customer a 70% credit for the old engine*[7]. The charges are $16.5K for an upgrade from a z196 and $33K if from a z10 EC. The same practice applies for zIIP and zAAP engines, $30K if coming from a z196 and $60K if coming from a z10EC.[8]

[7] *What a deal!* On which other platform do you get a very sizeable return on the money sunk into prior generations?

[8] For additional details on zEnterprise prices and pricing practices, 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.

Software Pricing

MLC prices for the zEC12 have been improved 2% to 7% compared to the existing *Advance Workload License Charge (AWLC)* for the z196. No additional price tiers have been added to accommodate the higher capacity of the zEC12.

Maintenance Pricing

System z maintenance is priced by capacity, not by core count. Price-performance improvement will be no less than 2% for standard CPs and 20% for specialty engines.

New Features Pricing

The IBM zAware feature is quoted as $40K for up to 10 CPs; additional features must be purchased to cover additional CPs (in multiples of 10). Flash Express is quoted as $125K per 1.6TB mirrored card pair; a maximum of four may be installed.

Evaluating Options for Your Data Center

Let's be up front here, IBM would like all of

its mainframe customers to sign up for a new zEC12 or an upgrade to the zEC12 – as soon as possible. However, as mentioned in my opening remarks, for those who have installed z196s during the last year or so, this may be present some practical and financial constraints. The exceptions, of course, are those customers that are experiencing exceptionally high processing growth rates and are running out of headroom on z196. Others for whom this offering would be compelling are those that require the security extensions provided by the Crypto Express4S coprocessor, e.g., the EMV feature (as previously discussed).

The most compelling argument and where IBM mainframe marketers likely will be focusing substantial attention are those customers whose largest mainframe is the z10 EC, particularly those who are experiencing capacity constraints or soon will be. The zEC12 will immediately deliver up to 160% additional total capacity. (See, again, Exhibit 2 on page 3). For z10 EC data centers, I recommend hedging down on this improvement for the short term on the assumption that the middleware and applications would be older versions that have not (yet) been leveraged with the zEC12's architectural enhancements. It is likely that current capacity requirements (as measured by the number of CPs on z10 EC) could be met on zEC12 with fewer CPs and therefore fewer cores, which could provide an opportunity for software cost savings.

Specialty engines gain similar performance improvement so that those workloads exploiting these engines, i.e., IFL, zAAP, and zIIP would greatly benefit. (See Exhibit 6, at the top of the previous page.)

And if z10 EC or z196 data centers also are considering more integration with their distributed systems, x86 servers under Linux or Windows, or Power Systems under AIX, they could avail themselves of the latest versions of the zBX and the Unified Resource Manager. Based on the prices proposed for the zEC12 and its features, it seems clear that the price-performance equation would tip the balance significantly in favor of the zEC12.

Conclusion

Once again, IBM has demonstrated why the mainframe is the ultimate platform for enterprise computing. On a total capacity basis, the zEC12 is 50% larger than its predecessor, which is especially important for the largest enterprises. With a clock speed of 5.5 GHz, it is the world's fastest commercial processor, allowing much more work to be done than on competitive x86-based servers. With its new Flash Express, availability of key workloads will be improved automatically by reducing latency. With IBM zAware, also new, data center staff will be alerted more quickly to difficult-to-detect operating anomalies and can act sooner to mitigate them. In summary, zEC12 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. Check it out!

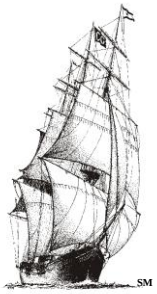

About The Clipper Group, Inc.

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.