

Having Your Cake and Eating It, Too — Doing More with Less with IBM's z10 BC

Analyst: Anne MacFarland

Management Summary

We've all heard of the philosophical debate surrounding the half-occupied glass: is it *half-full* or *half-empty*? The optimist sees it as *half-full* and the pessimist see it as *half-empty*. As an *IT infrastructuralist*, how do you see it? Does the glass seem *half-used* (and thus able to hold a lot more) or *oversized* (i.e., you paid for more capacity than you really needed)? This is more than a philosophical exercise. If you have to pay something approaching twice as much for the two times the capacity, then this is also an economics discussion for IT decision makers.

Would you be pleased to know that you had paid for twice the capacity that you needed, because you couldn't really anticipate how much you might need on a really busy day (possibly, because you really don't know how much you are using today)? Or, that you bought that much because you couldn't figure out how to manage your provisioning processes fast enough to meet the growth that might happen over the next year or two, or maybe the next day or two? Not if you were paying the bills. Paying twice as much "just in case" almost always is a bad thing to do. It's like paying twice as much for twice the coverage on your home insurance, *just in case* you have so undervalued the replacement costs and you want to know that you are covered.¹

Just-in-case (JIC) thinking is the opposite of *just-in-time (JIT) thinking*.² For an assembly line, with just-in-case thinking, you keep lots of parts inventory on the premises so that you will never run out. With just-in-time thinking, you provision smartly and force your component suppliers to deliver small quantities just before you need them. Of course, JIT requires more competency on your part. You need to know what and how much you are going to require before you need it. JIC is easier, of course, but costs more! In tough economic times like we're in, you may not be able to afford the luxury of JIC procurement of IT infrastructure.

None of this is new, of course, especially to most IT organizations. While the JIC method applied, in spades, to the procurement practices that resulted in enterprises having thousands to tens of thousands of probably underutilized servers in their data centers, there has been a widespread transition to consolidate the single or few applications typically running on those old servers onto many fewer new servers via a variety of approaches classified as *server consolidation* or *server virtualization*. That all seems very obvious until the half-empty glass question is raised: Is your collection of new, high-powered, multi-cored servers (virtualized or not) any more fully occupied (used) than the much larger collection that you have replaced? Or, are you still *half-full* (or *half empty*, take your pick) for all of the reasons discussed previously (i.e., don't know, can't manage, or feel safer with JIC over-provisioning)?

What if you could have your cake and eat it too? That is, you could have the capacity but not pay (much or anything) for what you were not using? This is the interesting saga that we have for you today. A pre-installed 10-core³ server that is available in many less-than-a-core capacities, but scalable up to the full capacity without opening the box. Interested? Read on to learn more about the new economics for new workloads on IBM's just-announced *z10 Business Class server (z10 BC)*.

IN THIS ISSUE

➤ You Need to Read This.....	2
➤ Introducing the z10BC.....	2
➤ Increasing the Gregariousness of z.....	6
➤ Mainframe Ways to Do More with Less.....	6
➤ Conclusion	7

¹ Usually, an insurance company doesn't mind selling you more insurance than on which they may have to pay. It's like insuring a 5-year-old car for its cost when new, even though the insurance company won't ever pay you more to repair your car than its book value.

² *Just-In-Time* initially was used to reduce costs and improve productivity in the automobile industry, but is commonplace today in both manufacturing and other sectors.

³ Actually 12, explained later.

You Need to Read This

Sometimes, things just are not as obvious as they appear. It took more than a century of scientific inquiry to recognize that dinosaurs were more like birds than reptiles. So, don't jump to any seemingly-obvious conclusions when I tell you that IBM's *z10 BC* is a *Mainframe*.⁴ Unless you've been involved with a 21st Century Mainframe, your recollections or second-hand impressions may be as outdated and wrong as dinosaurs being characterized as *terrible lizards*. This is especially true for the new *business class* *z10 BC*, targeted at:

- *New applications* (think Linux workloads and Java applications – including those from SAP and Oracle)
- *Database management systems* (think *DB2* and *Oracle*), and
- *Owners of older mainframes* (running back versions of *z/OS*, *z/VM*, or *z/VSE*) and also wanting to run the “open” workloads listed above.⁵



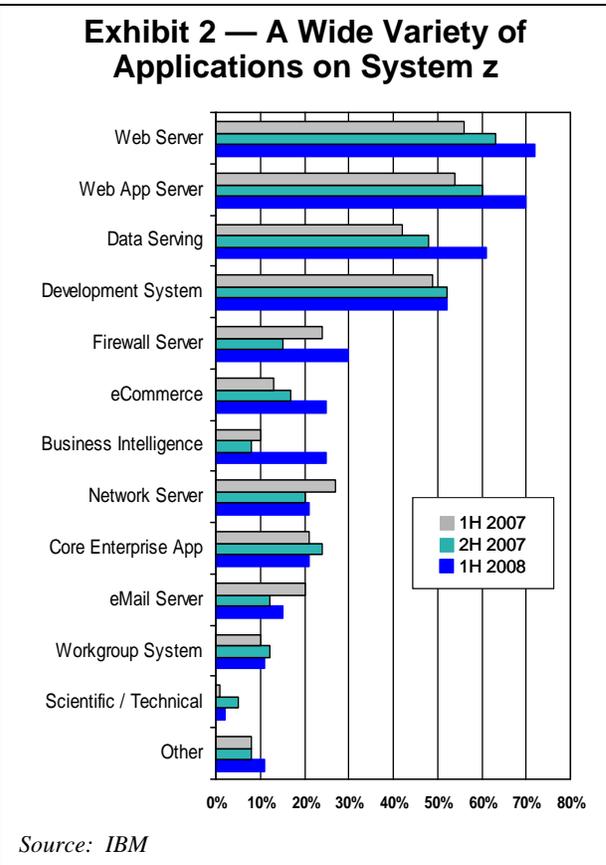
⁴ While the word *Mainframe*, to some, is emotion-laden, or even pejorative, and is often mistakenly linked to dinosaurs, it is better to think of it as a first-rate *scale-up server*. To learn more about what that means, see *Perceiving the Dark Side of the Moon - Knowing When Scale-up Computing Makes Sense*, in the September 23, 2008, issue of *Clipper Notes*, which is available at <http://www.clipper.com/research/TCG2008048.pdf>.

⁵ IBM found that the Mainframe was a most cost-effective consolidation tool at its own data centers, where thousands of x86 Linux servers were replaced by a couple of handfuls of z9 EC and, now many fewer, z10 EC Mainframes. IBM sees the Mainframe as the mission-critical hub at the center of its *New Enterprise Data Center (NEDC)* strategies. See Exhibit 3, at

While IBM may prefer to describe this first as an offering for its faithful, smaller mainframe customers, it is much more. (For a slightly irreverent discussion on this topic, see Mike Kahn's *Mainframe as Mild-Mannered Superhero?* in the issue of *Clipper Notes* also published today⁶.) This paper will focus on the *z10 BC*'s features.

Introducing the z10 BC

All of the single-book capabilities of IBM's *z10 Enterprise Class (z10 EC)* high-end server⁷ have been transferred to the *z10 BC*. It comes as a single model with all of the variations in capacity being delivered through the firmware. And, there are plenty of variations for this 10-engine⁸ workhorse. In fact, there are 130 possible capacity settings (using up to five of the engines) for traditional

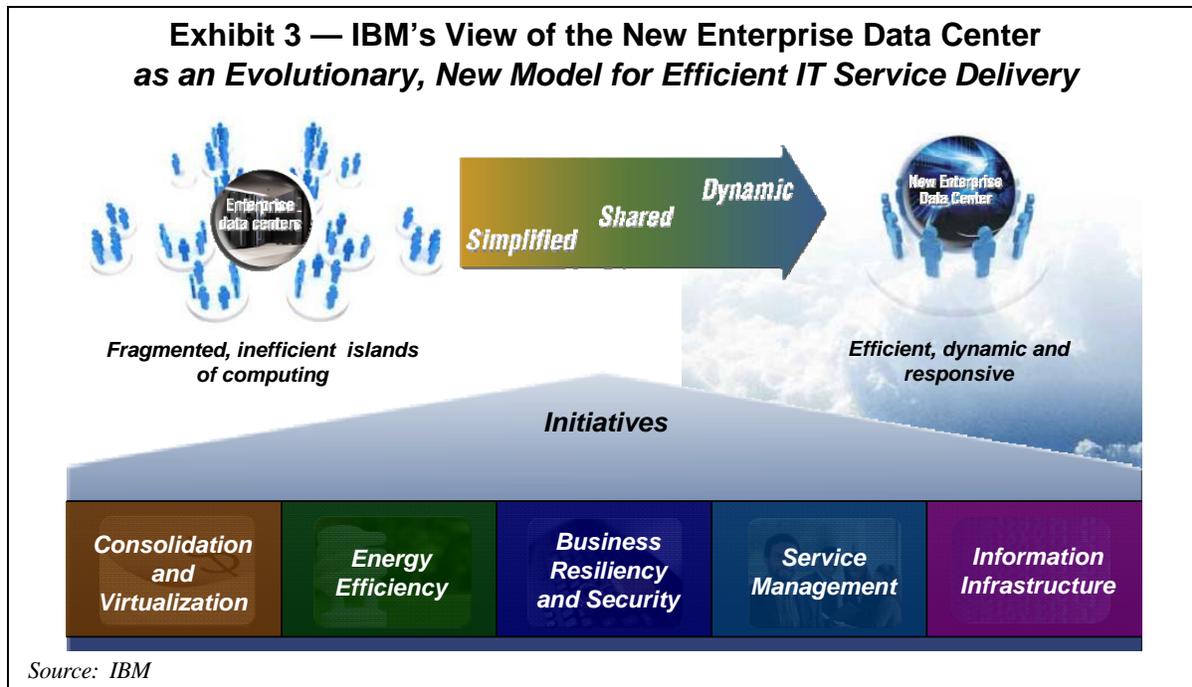


the top of the next page. (For more on NEDC, see the bulletin referenced in Footnote #7, in the next column.)

⁶ See <http://www.clipper.com/research/TCG200855.pdf>.

⁷ For details of the *z10 EC*'s capabilities, see **The Clipper Group Navigator** entitled *IBM's System z10 EC Meets the 21st Century Infrastructure Challenge*, dated February 26, 2008, and available at <http://www.clipper.com/research/TCG2008014.pdf>.

⁸ For those of you not familiar with the sometimes unusual dialect of *mainframe-speak*, an *engine* is a Mainframe processor core.



mainframe applications⁹ and up to all ten engines can be deployed for new workloads.

Looking Under the Hood of the z10 BC

Now, the issue of “10 versus 12” engines needs to be explained. (This is one of those poorly understood, yet important details from which IBM’s Mainframes get a “hidden” boost in their throughput.) There are four cores on each z10 BC processor chip. There are three of these processors in the z10 BC, for a total of 12 engines (cores). Up to ten of these are available to run the enterprise’s applications. The other two engines are called *Service Processors*¹⁰ and are used to initiate and accelerate I/O for the other five cores (and perform other “overhead” functions). If you are using all five of the user-accessible cores of the processor, you might say that you have 20% additional “horsepower” than meets the eye.¹¹

Existing small Mainframe customers can upgrade from the z9 BC and the z890. Plus, the z10 BC can be upgraded to a z10 EC, should even more operating capacity be needed.

The z10 BC has several enhancements over its predecessor, the z9 BC that will make it even more

attractive to smaller enterprises.

- In the z10 BC, there are no *books*¹². *Sockets* are now the *replaceable unit*. This makes replacements and upgrades more straightforward and faster.
- I/O now comes in drawers that can be added or removed without disruption to ongoing operations. This means that customers can start small and add more capacity as is needed. z10 BC will support multiple form factors within a footprint, which adds another level of investment protection.
- The clock of the z10 BC engines can run at 3.5 GHz¹³, which lessens the heat generated and allows the z10 BC to be installed (optionally) without a raised-floor.¹⁴ In addition, this engine, like its z10 EC counterpart, contains hardware for doing decimal floating-point calculations, a noticeable performance enhancement over the z9 offerings. (See Exhibit 4, at the top of the next page.)
- The minimum usable memory is 4GB, a lower entry point than before.
- The maximum usable memory is 120 GB (at

⁹ From a sub-capacity (i.e., less than a full engine) of 26 MIPS to 5 standard engines delivering 2700 MIPS (with the remaining 5 available for specialty engines for Linux (IFLs), Java (ZAAPs), and database (ZIIPs) workloads and interconnection to other mainframes (ICFs).

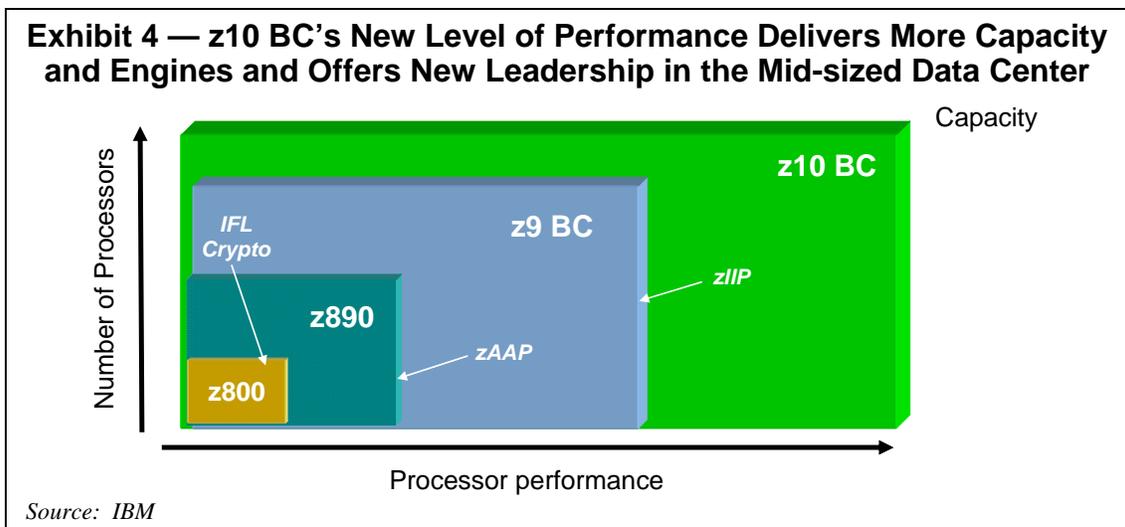
¹⁰ Also called *SAP*, for *System Assist Processor*.

¹¹ Another, maybe less orthodox, way to look at this is to consider that other processor architectures typically take some of each cores “horsepower” to do the housekeeping (i.e., the “overhead”), leaving less than the full capacity of the main processor available to do the users’ workloads.

¹² The “book” is a physical grouping of processors used in larger Mainframes since the inception of the *zSeries*. You might think of it as a physical cluster of engines. Larger Mainframes scaled physically by adding one or more books to the initial book.

¹³ The z10 EC’s engines run at 4.4 GHz.

¹⁴ I know it’s like comparing apples to oranges, but consider that Intel’s recently-announced 6-core *Xeon 7400* runs at 2.6 GHz. So the z10 BC engine is a hummer at 3.5 GHz, even though it is about 25% slower than the z10 EC’s very high-speed engine.



launch, with 4GB-chip modules) and 248 GB next June, when 8GB-chip modules are expected to be available).

- There is an additional 8GB of memory that is used for internal systems purposes and not an extra charge item (i.e., the base model really has 12 GB of memory).
- There are a number of new networking options. In connectivity with external networks, z10 BC supports the latest *OSA-Express3* communications adapter, with options for two or four ports of Gigabit Ethernet (GbE) or one or two ports of 10-Gigabit Ethernet (10GbE). *OSA Express 2* is still available on the z10 BC.
- Extending the range of ESCON and FICON data connectivity support, for both z10 EC and z10 BC, IBM offers a new high-performance FICON, called *zHPF*¹⁵, that doubles the start rate of the adapters, so fewer are required; they have more capacity for expansion and, generally, they get more data faster. At launch, the storage target for this capability is the *DS 8000* storage array (Release 4.1). As with previous enhancements of this type, third-party arrays will be supported, in time.
- The z10 BC includes, at no additional charge, full-speed cryptographic processors (the same chip as in the z10 EC), enabling high-speed encrypting, and decryption for secure sessions (such as SSL transactions). This is another way that System z delivers more bang for the buck than x86-based servers.
- For clustering, the use of an *Infiniband* bus for *Parallel Sysplex* is new – both short-range (to 150 M) and the long-range (to 10

KM) are supported – gives more than just connectivity. Infiniband's new intrinsic ability to trunk its multiple channels and manage the aggregate bandwidth opportunistically (the partitions used to be fixed) gives z10 BC owners just the capabilities they need to deal with information-intensive processes that often need a very-high priority required to support real-time operations – without having to add more boxes to get the job done.

A Special Deal on z10 BC Specialty Engines

For the z10 BC only, IBM is reducing the price by 50%¹⁶ for its *zAAP*, *zIIP*, and *IFL* specialty engines. The offer does not include *ICFs*¹⁷, which are used to link Mainframes.

IBM is doing this because it wants to encourage the use of specialty engines. The use of these engines is a key enterprise strategy both to increase price performance of running Linux, Java, XML, and database workloads on Mainframes, with a processor that has been optimized by firmware to the specific task and also as a matter of licensing cost avoidance. *z/OS* is what gives the mainframe its wealth of controls. While capacity-based pricing drives (based on MSUs) for *z/OS* and middleware has become more reasonable and offers better discounts than Mainframe pricing of the past, workloads that run on System z specialty engines avoid MSU-based pricing altogether. Thus, by transferring work from System z's standard engines to full-speed-enabled specialty engines, the savings can be sizeable. (See Exhibit 5 on the next page, for a cost comparison.)

One must remember the scope of these advan-

¹⁵ zHPF runs only on *FICON Express 4* and *FICON Express 2*.

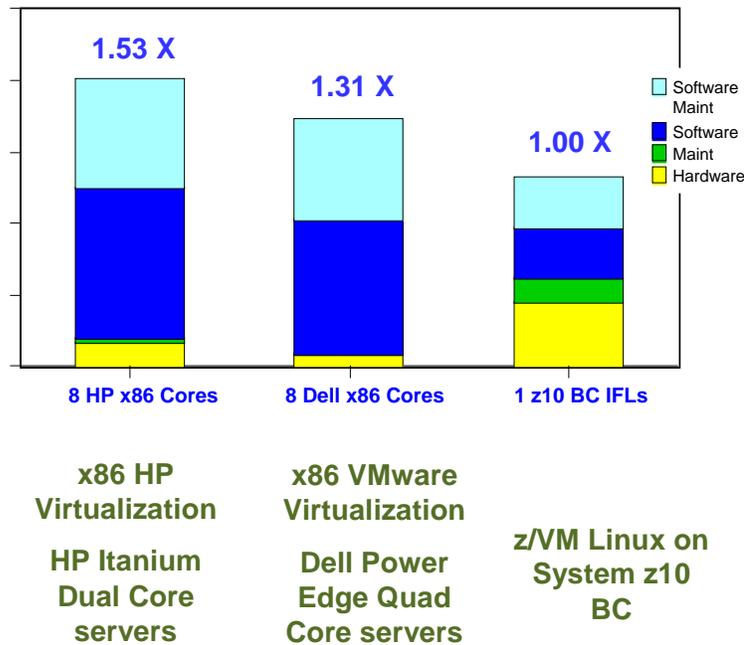
¹⁶ Over the price of a specialty engine running on a z10 EC. Given that the a z10 BC engine runs about 75% of the speed of the z10 EC engine, this works out to be about a 35% discount, a sizeable incentive, nonetheless.

¹⁷ *Integrated Coupling Facility*.

**Exhibit 5 —
Comparing the Costs of a z10BC IFL to Dell and HP x86**

All performance information was determined by IBM in a controlled environment. Actual results and IT costs may vary.

Oracle DB + WebSphere ND Workload
3-Year Total Acquisition Cost (TCA)



Source: IBM

tages. IBM claims that one IFL engine can support 7500 Domino end users at an energy cost of 100 watts. Think of what kind of energy you are spending to support employee productivity with your scale-out servers. Clearly, this is a different mode of thinking, but well worth considering.

First 16 GB Memory Less Than Half Price on z10 BC

Memory in z10 BC starts at a low entry point of 4GB, grows at 4GB increments to 32GB, and thence at 8GB increments to 120GB.¹⁸ With z10 BC¹⁹, IBM is slashing the costs of the first 16GB of memory on each processor by more than 62%, with the purchase of an IFL, zIIP, or zAAP.

New on both z10 BC and z10 EC, a plan-ahead memory option lets you pre-install additional memory and pay for it only when it is activated. This makes upgrades even more transparent.

These memory advantages are more than just a good thing. In conjunction with the specialty engines (also cost-reduced), z10 BC extends the potential continuum of what you do where – just as, at the far end of the distributed spectrum, cloud computing is extending it. There is no magic right place

¹⁸ This will increase next year to an entry point of 8 GB and a maximum of 256 GB.

¹⁹ Only the z10 BC; not the z10 EC.

for everything or even for a particular application – it depends on the business situation and, in some cases, the business model. However, the advantages of System z are unmistakably clear, once you start looking at the details.

Capacity on Demand Options

You may think that the above scenarios are well and good, but any box *is what it is* and still needs to be cost effective, no matter how well it might scale beyond your needs. However, with the Mainframe, you can choose what number of MIPS to use on standard engines and pay only for what you select as the capacity setting. Additionally, System z's (z/OS's) industry-leading temporary capacity-on-demand (CoD) options have grown from one or two variants to three combinable options, all of which can be invoked and augmented by the customers as needed – when needed – without on-site involvement of IBM.

- The classic On-Off Capacity on Demand
- Back-up capacity upgrades
- Capacity for a planned event

System z10 BC adds two more CoD options, targeted at organizations where demonstrable control over expenses is required. One option is the procurement of a fixed amount (expressed as a finan-

cial amount) of capacity to be installed on demand. The other, like a phone card, is a replenishable supply.

Storage

IBM also introduced some new Mainframe storage offerings concurrent with the z10 BC launch.

- As mentioned early, the 4.2 release of the *DS 8000* supports zHPF. It also supports *z/OS Metro/Global Mirror Incremental Resync*. This system software offers speedier recovery after a disaster or other outage by focusing on the data that are incremental changes. This is the “completer” for full restore – and by focusing on only what must be available, it reduces time to recovery from hours to minutes
- The z10 BC launch was also the occasion of the launch of two new models of IBM’s virtual tape systems for use with the Mainframe. Both new models expand the capacity and performance of the previously-announced *TS7700 Virtualization Engine*. The *TS7740* works in *D2D2T* solutions, while the *TS7720* has no tape interface and works in *D2D* solutions.²⁰ In addition, the *TS7740* now includes an integrated controller for the *TS-3500 Tape Library*, eliminating the need for an additional – floorspace-consuming – piece of hardware.

Migration Paths

Both the z890 and the z9 BC can be upgraded to a z10 BC. Both will get their specialty engines upgraded, at no charge. If you have purchased specialty engines for a prior model of System z, you will get the same number in z10 for no additional charge, representing a significant performance enhancement. Where else, beside System z, do you get free and faster replacement processors when you replace an existing server?

Lifecycle extension of z/OS 1.7

z/OS 1.7 is the oldest version of z/OS that can get current with a single upgrade, but it just went out of service on September 30. Realizing that double-hop upgrades can be extraordinarily painful, IBM has extended support for z/OS 1.7 to give those enterprises still using older versions of z/OS an easier upgrade path.²¹

²⁰ See *Virtualizing the High-End Tape Environment - IBM Enhances VTL Offering* in **The Clipper Group Navigator**, available at <http://www.clipper.com/research/TCG200858.pdf>

²¹ This is a fee-based service, available through September 30, 2010.

Increasing the Gregariousness of z

Rational Element for z

IBM Rational has made two of its most exciting elements available on System z.

- *Rational Asset Analyzer* maps application interdependencies across Mainframe and distributed platforms.
- *Rational Team Concert for z* takes Rational’s hip, new collaborative platform that supports remote teams of developers in application development, deployment, and management and enables it for System z developers and development on System z. Never before has the image of green screens been more obsolete.

Independent Software Vendors

The influx of new software for System z has accelerated. Over 130 new ISV partners have contributed over 600 new applications and tools to the platform. In addition, there are almost 2500 new applications for Linux on z.

Tivoli Business Service Manager

This product provides health and performance monitoring across Mainframe and distributed platforms. With the z10 BC release, it further adds support for Linux in z/VMs on System z.

Academic Initiative

The Mainframe Academic Initiative continues to gain momentum. Almost 500 schools participate, including many beyond U.S. borders. In this tight economy, students are gravitating to the rewarding opportunities of working in Mainframe environments

Competitive Strategy

Gregariousness can include competition as well as congeniality. With z10 BC, IBM will address more directly some of the opportunities it used to leave to other IBM platforms. The *IBM Migration Factory* is being extended from *System p* to System z. IBM will work with partners, such as *ACI*, to address opportunities in industries, like banking, where z10 BC offers significant competitive advantages (like high-availability and very-high security).

Mainframe Ways to Do More with Less

Denser Server Virtualization

This approach leverages the extraordinary high granularity and low overhead of z/VM. There are two basic scenarios.²² One involves using the z10 BC as a Linux-only box. This is a classic scale-out scenario, but instead of tens or even a hundred virtual machines per processor, several thousand can be supported. If you are paying for

²² Both modes benefit from System z’s CryptoProcessor.

your middleware or applications using per processor pricing, this granularity may affect your bottom line significantly. Furthermore, for running Oracle on Linux on z, only basic knowledge of z/VM is needed.

Add an instance of z/OS and your options are very different. z/OS handles the offloads to specialty engines – zAAPs for Java processing, zIIPs for XML and information-rich processes, and IFLs for Linux. Offloaded processes avoid additional z/OS licensing charges and benefit from the in-box efficiencies supported by System z internal communications (via *Hipersockets*, where traditionally-networked communications (from one server to another) are done, securely, by memory-based transfers at much higher speeds. Just as important, z/OS arbitrates use of resources, as needed, with full auditability and unmatched security.

The Potency of Co-Location

The advantages of co-location are many. Systemic concerns, like security, can be addressed more simply. Resource sharing (particularly with z's sophisticated methods of arbitration) can be very effective. The data is already there. There is no need for external appliances or remote data warehouses.

The Greenest Server?

We think so, but it really depends on your workloads, so you need to run your own numbers, or call us for help. For many smaller enterprises, z10 BC is the server that is *just right!* This, by itself, is a topic worthy of a separate bulletin. For now, think of the obvious energy and cooling savings offered by the z10 BC (in one rack), compared to the many (often dozens or more) of racks of servers that they will be replacing.

z10 BC-Specific Solutions

And there is still more, for some leading software solutions:

- **SAP:** For the z10 BC, IBM offers a package of *SAP* elements, software, hardware, and maintenance that is competitive with distributed platforms.
- **Cognos:** *Cognos 8.4 for z* supports real-time analytics of operational data – and the more of it that resides locally, the faster the response will be.
- **IMS:** *IMS* is one of the most venerable and most optimized database management systems on the market. *IMS* continues to be a key application because, as a hierarchical database, it is better at some tasks than a relational database. Think of the simplicity of triple stores (i.e., three-column tables) that underlie rapid, comprehensive

search.²³ *Sometimes, less is more.* A significant solution in this area is expected by the end of the year.

A Reason to Act Now

IBM has a fourth quarter financing incentive, if you order this year – a 90-day deferral on payments and interest. This spoonful of sugar will help the medicine go down, medicine that will better position your enterprise for 2009.

Conclusion

Never have the demands on technology been more diverse. Organizations want to use their assets more effectively and efficiently while limiting the opportunity for things to go wrong. They want to limit their burgeoning IT management costs while assuring auditability, findability, and forensics.

Government mandates for control and auditability sent folks in the “open” world scrambling for distributed and often labor-intensive solutions to systemic challenges. Virtualization helps, but all virtualization is not the same.

IBM's Mainframe always has had a different formula – one based on assumptions of *multi-tenancy*. These assumptions have generated and improved highly-integrated control functions that may have seemed to be overkill in fast, loose and ore trusting times. What many have forgotten is that the Mainframe also supports large-domain and multi-tasking strategies for *doing more with less* – and that it does so better and better the more diverse workloads that it is asked to handle. Unlike other platforms, it has, not a linear formula for fullness, but several ways, from z/VM, to memory controls, to HiperSockets, plus the opportunity to optimize the formula to your benefit. Only the Mainframe can use – fully – the many dimensions of capacity, at full capacity, i.e., *to do more with less.*

If you would like a change from the linear approach to infrastructure that brought you sprawl, consider the benefits of Mainframe infrastructure. It may behoove you to do more – with the smart, cool, efficient, and cost-effective *Mainframe for the Rest of Us* – the z10 BC.



²³ Check back for a forthcoming bulletin on *IMS* at <http://www.clipper.com/publications.htm#Catalog>.

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

Anne MacFarland is Director of Data Strategies and Information Solutions for The Clipper Group. Ms. MacFarland specializes in strategic business solutions offered by enterprise systems, software, and storage vendors, in trends in enterprise systems and networks, and in explaining these trends and the underlying technologies in simple business terms. She joined The Clipper Group after a long career in library systems, business archives, consulting, research, and freelance writing. Ms. MacFarland earned a Bachelor of Arts degree from Cornell University, where she was a College Scholar, and a Masters of Library Science from Southern Connecticut State University.

- ***Reach Anne MacFarland via e-mail at Anne.MacFarland@clipper.com or at 781-235-0085 Ext. 28. (Please dial “1-28” when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, 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.

Disclosure

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.

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.