

IBM's System z10 EC Meets the 21st Century Infrastructure Challenge

Analysts: Anne MacFarland and Mike Kahn

Management Summary

The race is on! However, make sure that you are watching the right race! This might be the greatest challenge facing data centers and IT organizations in larger enterprises. If one were to listen to the announcer calling this horse race, you might hear that *product A* just passed or even leapfrogged over *product B* and is on its way to cross the finish line first. If we were talking about horses on a dirt track, then this kind of simplicity might make sense, because the track is linear and the race is a one-time event.

Unfortunately, delivering IT to and through the larger enterprise is anything but linear and so multi-dimensional that it is unfathomable by most, if not all, humans (i.e., without computer-aided planning, measurement, and operational tools). So why, then, do we always go back to checking the status of the latest horse race? Possibly, because it is measurable, and that, by itself, delivers certain pleasures.

With IBM's announcement of its next-generation mainframe, the *IBM System z10 Enterprise Class (z10 EC)*, the first impulse is to slice and dice the specification and offering details to see exactly how much more "powerful" it is. We will satisfy some of that urge within this paper. In short, it is faster (clock speed, as if that has any real meaning) and bigger (maximum MIPS) than anything that we have imagined, short of a special-purpose computational grid. **The more important question, that must be answered first, is why do we need a monolith of this scope? Our goal is to answer that question, in several ways, and then to describe the specific details of the System z10 EC.**

The Obvious Need for "More"

Today, few of us are satisfied by the status quo. Most certainly, few businesses would say that they would be happy with *less*, be it revenue, profits, customers, or the ability to do business efficiently and effectively, etc. Even when we try to contain growth, say, as in the growth of data storage, we are talking about slowing the rate of growth and not shrinking to some quantity that is less than presently consumed. **Thus, in our IT-centered world, "more" is presumed in all dimensions, except for the available budget. What this really means is that we have to do more work (however you want to measure it) for less money per unit of work done (than in the previous year).** If we think about computers, networks, and storage in traditional ways, we might conclude that we need more processing power, more bandwidth, and more capacity; all of which may be true. On the other hand, maybe we need to be smarter about how we do processing, how we consume our networks, and how we fill our arrays. In the end, it still may come down to needing more, but the sum total likely will be much less than when you make simple, linear extrapolations. *What makes the difference? How well we manage our resources.* Please read on, for the details.

IN THIS ISSUE

- The Obvious Need for "More" 1
- The Other Side of the "Big" 2
- z10 EC's Scalability, Flexibility and Performance – All a Matter of "Virtualized Big" 3
- Managing Capacity, Availability and Resilience in a Climate of Change 5
- Diversity, Too 6
- Conclusion 8

Rationale for Decision Making is Important

How we make IT decisions often comes down to a few very important concepts and our motivations in addressing them.

- **Simplicity** – Do we do things simply, because that is the best way? Or, do we act simply because it is easier to comprehend and do?
- **Status Quo** – Do we presume that our status quo is a meaningful base point for looking to the future, because it is understood and functional, at least supposedly? Or, do we use this as a base point for next year because anything else would be too hard to consider and sell to management?
- **Open-Mindedness** – Do we presume that we are “open-minded” because we continuously consider new products and technologies to add to and refine the “stew” that constitutes our data centers? Or, are we really being xenophobic to rethinking how we do what we do? Does being open minded mean following the herd of generally-accepted platitudes and conclusions?

This list could be longer but we think that our point has been made.

If you are doing what you are doing because it allows the status quo to evolve (presumably for the better) in a simple and understandable way then you may be missing the most important possibilities for improvement (because they are either not evolutionary or because they are not simple)!

The Better Road to “More”

Sometimes, “more” is just an incremental improvement, like a chip going faster and, thus, work getting done more quickly. However, as we all have seen, “more” can sometimes mean “more cost, but no more, or possibly less, benefit”. For example, if more CPU cores means potentially more processing power, then that could be good, but only if you can take advantage of the increased processing power in an economical and meaningful way. If you have to add virtualization software to pack mixed work onto a multi-core server (in order to get sufficient utilization from your investment), if you have to manage that in real time, and if you also have to rethink how to do your backups and replications (maybe by buying additional software), then – just possibly – you haven’t made the progress that you had been seeking through consolidation (except, maybe, your energy costs did go down).

So, is the “simple path” from where you are today the best way to go? This question does raise the huge issue of **how do you go from an infra-**

structure-centric approach to provisioning IT for the enterprise (the mantra for the last 50 years) to one that is much more focused on delivering IT services to the enterprise (the mantra of the 21st Century). This requires a lot of analysis and rethinking of your “rules of the road”, something that is not the focus of this paper.¹ **Nonetheless, it is most, most important that you think of the mainframe in terms of the latter.**

The Other Side of the “Big”

Looking at IT only from the inside out will give and your enterprise a never-ending headache. To balance the tendency to see from the inside out, you also need to consider the health and welfare of the business side of your enterprise.

With global competition, “good enough” processes, timeliness, and security often are insufficient in four ways – (1) the quickness of response, (2) the breadth of self-knowledge, (3) the integration and pervasive reach of controls, and (4) the readiness for change. While these represent challenges a many, a bigger problem exists! You must address resolution by the opposite of a divide-and-conquer strategy. Stick with us and think about those four insufficiencies and we’ll explain through some examples.

- **Operational baggage alienates customers and corrodes profit quicker and more implacably than ever before.** To achieve quicker time-to-market of goods and services, providers now buy and partner for capabilities they once would have developed internally. *Thus, there are more relationships to be managed well.*
- **The cheeseparating of cost cutting, to meet the pragmatic demands of customers, requires more evaluations and decisions than old-fashioned “business-as-usual”, and drives a need for more people to know more about the organization.** *Thus, there is more information to be shared well.*
- **Self-service, one of the main strategies of cost cutting, allows retrieval of more accurate information, more quickly, at less cost,** but it also incurs additional quality control and security requirements that must be built in, not festooned at organizational edges, *for there are few edges.*
- **Markets have seldom been decorous, but now they are more rambunctious than ever.** No company, and no set of loyal and understanding customers, can set the cadence of demand. **Thus, more unexpected things will happen more frequently.**

To stay solvent, businesses have to address more markets, support a high rate of organizational

¹ Look for a forthcoming bulletin on changing priorities and perspectives for enterprise data centers.

change, and wring more value from their intellectual property than they would have found comfortable just a few years ago. In short, they have had to become more flagrantly (but carefully) opportunistic. This involves embracing many dimensions of “more” and a willingness to accept of the vastness of the challenge and the need for a *big solution* – but, preferably, one that might be described as *virtualized big*. Sitting in traffic after a game or a movie, who has not stared at the empty lanes on the other side of the street or the highway and dreamt of a virtual on-ramp and a way to reuse those lanes without having to pay overtime to hundreds of police officers?

Technology does *virtual big* very well – but even technology comes with terms and conditions. Total cost of ownership, including energy costs, limits which technology choices are possible. The new complexities of business resilience (driven by a high rate of change in most businesses) demand that recovery processes be leveraged to support business change. Process acceleration must be available where it is needed, supported by whatever IT strategy (parallelization, offload) makes sense. Enterprises are looking for their IT systems to provide the quiet capability, the elegance, and the grace under pressure, to run the business at speeds and under pressures that 20th Century businesses would find incredible. And, its only 2008!

The Need for Massive Scale

Certain things have become clear.

- **The economies of scale that have driven businesses to grow have become more central to strategies than ever, and the scale that is demanded has grown enormous.** At large scale, a share-everything architecture – with the controls to handle the trade-offs, such as the mainframe – is the most economical per unit of work, particularly when that unit of work is complex. Grids and clusters are steps in this direction and the System z mainframe uses and participates in those strategies, too.
- **Large data sets and large data blobs, such as video, have become part of more and more business processes.** High-performance computing has followed this large data into everyday processes, be it design, multi-variable analysis, or reality simulations for training or education. The need for this capability often is intermittent, however.
- **Managing IT capacity and availability – in an environment of integrated workloads – by anything other than policy and automation is unworkable.** Managing capacity, availability, and recovery as separate processes carries a lot of redundancy that, in an environment of rapid change, becomes unsupportable. Managing such

an environment by a bunch of separate controls, albeit federated to some extent, becomes onerous, and achieving end-to-end security is even more so.

z10 EC’s Scalability, Flexibility and Performance – All a Matter of “Virtualized Big”

For many large businesses, highly-integrated operational support of complex business operations, including all of the strategies outlined above, is provided by one or more IBM mainframes. For them, **the mainframe is not just an alternative to “disposable computing”; it is indispensable.**

The next generation of the IBM mainframe, called *System z10*, optimizes this integrated, virtualized multi-workload computing still further. **IBM’s “virtual big” is the System z10 EC (Enterprise Class), in a still larger overall sharable architecture, with more capabilities (on-the chip encryption, decimal floating-point calculations, and data compression), and with more flexibility to address both the unpredictability of workloads and the addition of new workloads (another mainframe design point).** There are also some revolutionary changes in the z10 EC mainframe, made possible by a new chip design, changes in System z’s z/OS, and by enhancements to IBM storage products, coupled with z/OS data protection software.

z10 EC provides the scale, speed, granularity of controls and pervasive virtualization to address the many kinds of more that allow an organization to be approachable, competent, and elegant. This is a far cry from the mainframe of a decade ago – though still backwards-compatible to support the legacy applications, honed over decades that represent the secret sauce of many businesses. The z10 EC addresses the opportunistic nature of contemporary business by adding speed, connectivity, capabilities, and new kinds of sophistication that are needed to handle the applications that underlie both very large enterprises and very demanding smaller ones.

Providing Big Capability at the System Level

z10 EC blows away the previous mainframe scale by offering over 30,000 usable MIPS. The processors (engines) come in Multi-chip Modules (MCMs), each of which has five four-core chips, which are grouped into physical collections called “books”. Two of these cores are kept as spares; more details follow. This sparing increases the resiliency of the system, since failover is transparent.

The components of the new z10 EC are significantly bigger. The memory minimum is 16 GB, and the maximum 1.5 TB (384 GB per book). The 16 GB Hardware System Area (HSA) is now separately managed, which makes enables more non-

disruptive hardware adjustments, such as partition deployment and capacity changes. (More details, below.)

The L2 cache interconnect between books is now a *star* rather than *ring architecture*. This better accommodates the heavy data use that IBM forecasts as increasingly important.

The five available Models of the 2097 z10 EC machine range from the *E12* to the *E64*. The *E12* features one Multi-Chip Module (MCM) with 17 usable processing units (PUs), including three System Assist Processors (SAPs) and two spares. The extended model, *E64*, has one 17 Engine MCM, and three 20 engine MCMs, for a maximum of 77² engines, although only 64 can be controlled under PR/SM at this time. At the top of the line, the *E64* can support up to 11 SAPs – it is a sliding scale, which makes sense given the nature of their function.³ There is an upgrade path from the *E12* to the *E64*.

All PUs are configurable as *Central Processors (CPs)*, *Integrated Facilities for Linux (IFLs)* or, within limits, as *Internal Coupling Facility (ICF)*, *System z10 Application Assist Processor (zAAP)*, *System z10 Integrated Information Processor (zIIP)*, or additional *System Assist Processor (SAP)*. How they are used, and at what capacity they run, allows a great deal of flexibility of configuration.

In addition, newly-available 6 GBps InfiniBand buses and Infiniband coupling links⁴ can connect z10 ECs into a *Parallel Sysplex*.⁵ Also, OSA-Express3 10 GbE will be available in the second quarter of 2008.

Supporting Big at the Chip Level

The System z10 chip runs at 4.4 GHz. The four cores per die each have their own 3 MB Layer 1.5 cache. There are on-chip engines for data compression, cryptographic, and decimal floating-point operations. SMP communications are achieved via a hub chip (there are two on each MCM), featuring shared cache, SMP fabric, and a bus speed of 3 GHz. These all bulk up the processing chops of the mainframe to meet the new challenges that information technology is asked to address.

There are over 50 new instructions, focused on improving compiled code efficiency. The new de-

sign increases the physical closeness of address, control, and data flow paths. This allows smaller chunking and better parallelization. Aggressive branch prediction and Millicode entry/exit predictions improve performance. Added front-end cycles resolve interoperations timing. All these are of particular benefit to Java and Linux workloads (and SOA operations), which are more CPU-centric and therefore more sensitive to CPU pipeline performance. Java workloads get an additional boost from dynamic (“just in time”) compilation, which allows existing applications to exploit new instructions and re-optimize for the new design as soon as they are moved to the new system.

The System z10’s chip supports three addressing modes and multiple arithmetic codes. It supports all z-compliant operating systems and preserves mainframe application compatibility going back to 1964.

System z’s new decimal floating-point accelerator on the chip delivers performance, precision, and function. Previously, the z9 EC had this in firmware. Now, with it in hardware, you get a performance boost of perhaps an order of magnitude over z9 EC – and even more over software-only implementations of this function on other platforms.

Even the acceleration that the z10 chip provides is massive in scope, building on the function supported on the z9 processor. It features dictionary-based data compression and expansion, supporting dictionary sizes of up to 64 KB (about 8000 entries). It also features support for a broad range of cryptographic standards⁶ in hardware. The on-chip bulk encryption rate is 290-960 MB/second.

The z10’s chip supports error correction and recovery throughout processing. ECC is used on second- and third-level cache, store buffers, and R-Unit state array. Parity error correction is used on all other memory. There are over 20,000 error checkers on each z10 chip. The chip architecture allows precise core retry for almost all hardware errors and the machine check architecture allows precise software recovery. All these features give little opportunity for the errors that cause outages. A truly dynamic, integrated system can be built with this chip. It is anti-commodity massiveness with a purpose.

Facilitating Big Scale by z/OS

z/OS v 1.9 and v.1.10 will extend their n-way limit to match the capacities of the new hardware. It will use *HiperDispatch*⁷ – an ability to align tasks with a logical processor and with a physical processor – to keep data local, because latency really mat-

² 11 SAPs and 2 spares comprise the extra engines.

³ System assist is all about the system controls built into the hardware that give the mainframe its imperturbable resilience and ability to share all resources between multiple workloads by virtualization (LPARS, z/OS and z/VM) and granular inherent resource controls.


⁴ With the new InfiniBand Coupling, clustering with other mainframes is extended to 150m.

⁵ The InfiniBand Sysplex will work only with z10 and the prior two generations of mainframes.

⁶ These include DES (DEA, TDEA2, TDEA3), SHA-1 (160 bit), SHA-2 (256, 384, and 512 bit), and AES (128, 192, and 256 bit).

⁷ Available only with System z10 servers.

Exhibit 1 – Power and Space Comparison (z10 EC vs. Linux on Intel)



The Linux on Intel servers selected in this example are functionally eligible servers considered for consolidation to a System z running at low utilization such that the composite utilization is approximately 5%. The utilization rate assumed for System z EC is 90%. This is for illustration only, actual power and space reductions, if any, will vary according to the actual servers selected for consolidation.

Source: IBM

ters in delivering real-time capabilities. This feature typifies System z10's strategy to meet the needs of a business world where massive data is critical to how well business is done. The operating system knows all the relevant information and is good at this complex optimization. HiperDispatch is part of a move to define, more clearly, where human decisions are needed. All else should be automated.

These improvements allow a single system image to match more closely the multi-image LSPR number. Such virtual congruence lets the mainframe scale to 64 engines more transparently. This n-way optimization initiative started with the *System z990*. It will benefit customers with its ability to do analytic workloads much faster and the ability to consolidate even more workloads on the energy-abstemious z10 EC. (See energy discussion, below.)

Extending the Big Scale to Storage

The large data sets and data volumes, discussed earlier, demands some extensions to mainframe data storage capabilities.

- **New Extended Address Volumes (EAV)** on the *IBM SystemStorage DS8000* with z/OS v 1.10 will take volume growth beyond a 64 KB cylinder cap, which has limited volume sizes to 54 GB. Initially exploiting VSAM access methods, EAV will enable volumes up to 223 GB in size. IBM plans for this number to grow much larger, over time, to hundreds of terabytes. Fewer, bigger volumes simplify management.
- **HyperPAV⁸** will allow the scaling of I/O against

these larger volumes. The DS8000's *Dynamic Volume Expansion* can be used to migrate to the larger volumes. In addition, the just-announced next-generation *IBM SystemStorage SAN 768B* director and fabric backbone will include 4 Gbps and 8 Gbps link speeds.

More Potent Specialty Engines on z10 EC

The performance of specialty engines, once again, has improved, while the price remains the same. Existing specialty engines will get upgraded at no additional charge when transferred to a new z10 EC.

With three times the memory of z9 EC, z10 EC can do more. In a highly-virtualized Linux consolidation environment, the performance improvement is estimated at 1.5 times, and should be even better for CPU-intensive workloads.

Doing Big with Less – z10 Energy Efficiencies

The z10 EC offers a 15% improvement in performance per KWH over the z9 EC. It uses the energy-monitoring tool that was introduced last year on the z9 EC. In addition, IBM also introduced last year the *IBM Systems Director Active Energy Manager (AEM)* v.3.1 for Linux on System z, which provides a single view of actual energy usage across multiple heterogeneous IBM platforms⁹. In addition, compared to servers based on Intel x86 archi-

⁸ Parallel Access Volumes

⁹ AEM can both monitor and manage power use. On System z, AEM will be used only to monitor, since z has its own separate, well-honed power management. (See Exhibit 1, above.)

tures, System z uses a lot less energy.¹⁰

Managing Capacity, Availability and Resilience in a Climate of Change

Addressing the business challenges outlined early in this paper with an integrated approach begs a focus not just on system availability and resilience, but also on its management. This, like the operational challenges facing business, is an expanding set of concerns. For decades, IBM has worked to increase the reliability of the mainframe, to the point where unplanned downtime is truly a rare occurrence. With z10 EC's separately managed HSA, IBM has increased the configuration changes can be pre-planned, so that IT operations that formerly required an outage can now be invoked non-disruptively. With z10 EC, you can pre-plan LPARS, so that they can be defined and non-disruptively enabled, when you need to do so. Additionally, more IT changes can be done on the fly, as needed. This is particularly important in the areas of capacity expansion, availability, and resilience.

Capacity

Providing the ability to scale up and scale back capacity, and giving the ability to allocate spare, latent capacity where and when the business will need it, have been capabilities that the mainframe has provided sooner and more granularly than other platforms. z10 EC takes these capabilities several steps further.

In z10 EC, sub-capacity operations are available for up to 12 Central Processors. The sub-capacity comes in three levels, so up to 36 sub-capacity settings are possible¹¹. All Central Processors (CPs) must be at the same capacity within one z10 EC. On a system with both CPs and specialty engines, the specialty engines run at full capacity even if the CPs are sub-capacity. All sub-capacity-governed Central Processors will have built-in cryptography and compression, both of which will run at the engine's full speed. This makes sub-capacity operations (and sub-capacity pricing) more generally available to businesses that may want to build in performance sparing for particular workloads.

On-Off Capacity on Demand is the classic mainframe offering for temporary capability. This has been joined by *Capacity Backup Upgrade (CBU)*, which adds capacity to speed back-up. With the

release of z10 EC, these options are joined by *Capacity for a Planned Event (CPE)*, which moves capacity between machines for, say, a planned event, such as an upgrade.

It's not just that you now have three tools instead of one that can be used to optimize IT operations. They can be in use simultaneously, and can be activated without any interaction with IBM¹². Moreover (and this is new), existing capacity offerings can be reconfigured or replenished dynamically and in response to policy. This improves an enterprise's ability both to react to real-time events and to use its existing capacities to solve logistics problems and otherwise pursue dynamic business ambitions.

These features, together with the reduction in occasions of outage and the ability to automate incremental system changes, lets customers coordinate the imperatives of availability and cost control into a strategy focused on cost-effective system resilience. Customers can begin to prioritize and schedule their capacity needs. This flexible provisioning approach provisioning architecture will let capacity be added for certain workloads at certain times, while staying within all terms and conditions that pertain to the physicality of the model and the policies of the business. If you think back to all the demands and uncertainties discussed on page one, these kinds of capabilities give several ways to hedge operational bets. (See Exhibit 2, on the next page.)

Availability

The need for low-latency access to rich data and large data sets means that data replication has relevance far beyond data protection and automated data recovery. System z has a well-honed set of capabilities that can be used to help business operations in this area.


A significant improvement in FICON is the ability to process more commands in flight. Now channel extenders do not have to be FICON-specific for *Global Mirror (XRC)* solutions. This offers a big opportunity for significantly reducing costs. The ability to use multi-protocol extenders brings down the cost of an availability solution and lets multi-site data redundancy be practiced as an operational tool, as well as for business continuity.

Now, *z/OS Global Mirror* allows most System Data Mover (SDM) processing to be eligible for the zIIP specialty processor. With zIIP-assisted z/OS Global Mirror, the zIIP essentially becomes a z/OS data mirroring engine that can help lower server utilization at the recovery site, or create server "white space" for other projects. This strategy will work with all z9 and z10 models with z/OS 1.8 or

¹⁰ In a consolidation, the System z10 EC may provide up to 8 times the same work in the same space and may provide up to 16 times the work for the same power consumption.

¹¹ For a single sub-capacity processor (most installations have more), the sub-capacity ratios are .69, .51 and .24. As the sub-capacity processor count grows, the percentages creep up a percentage or two per additional processor. Exact ratios are available in the zPCR or LSPR tables on the IBM Website.

¹² All capabilities will be on the z10, to be entitled when the server is ordered or upgraded.



later and with IBM SystemStorage DS8000 (or any storage controller supporting z/OS Global Mirror). z/OS Metro and Global Mirror also now feature an incremental sync using GDPS that is faster than in previous releases.

Advances in GDPS and Basic HyperSwap

By integrating GDPS¹³ with Veritas Cluster Servers, IBM provides coordinated disaster recovery across heterogeneous environments. In the first half of 2008, IBM will integrate GDPS with System Automation for Multiplatforms Application Manager, which will bring Tivoli into the mix.

IBM plans to offer Basic HyperSwap, a single-site data protection scheme targeted at planned and unplanned disk outages. Like GDPS, Basic HyperSwap has the interface to manage a remote freeze for consistency. By contrast, Basic HyperSwap is designed for a seamless switch over to secondary disk within the same site to recover from disk failure, while GDPS also allows an entire data center swap for disaster recovery across a long distance. Basic HyperSwap will be enabled by z/OS and plans to use a TotalStorage Productivity Center for Replication for System z (Tivoli) interface. It does not involve a service component. If a data center is knowledgeable, it can set it up by itself for no charge. Setting up GDPS includes service costs to help install and customize and to provide skills transfer and project management. GDPS uses a NetView interface. With these developments, IBM offers a full spectrum of data availability options for the mainframe environment. (See Exhibit 3, on page 8.)

¹³ Geographically Dispersed Parallel Sysplex.

Business Resilience

All the capabilities detailed above give new ways to foster business resilience while keeping our System z10 mainframe tuned to the needs of the moment. Not only is business recovery easier, but supporting changing business operations gets an assist from both capacity boosts and data mobility.

Managing Only Exceptional Exceptions

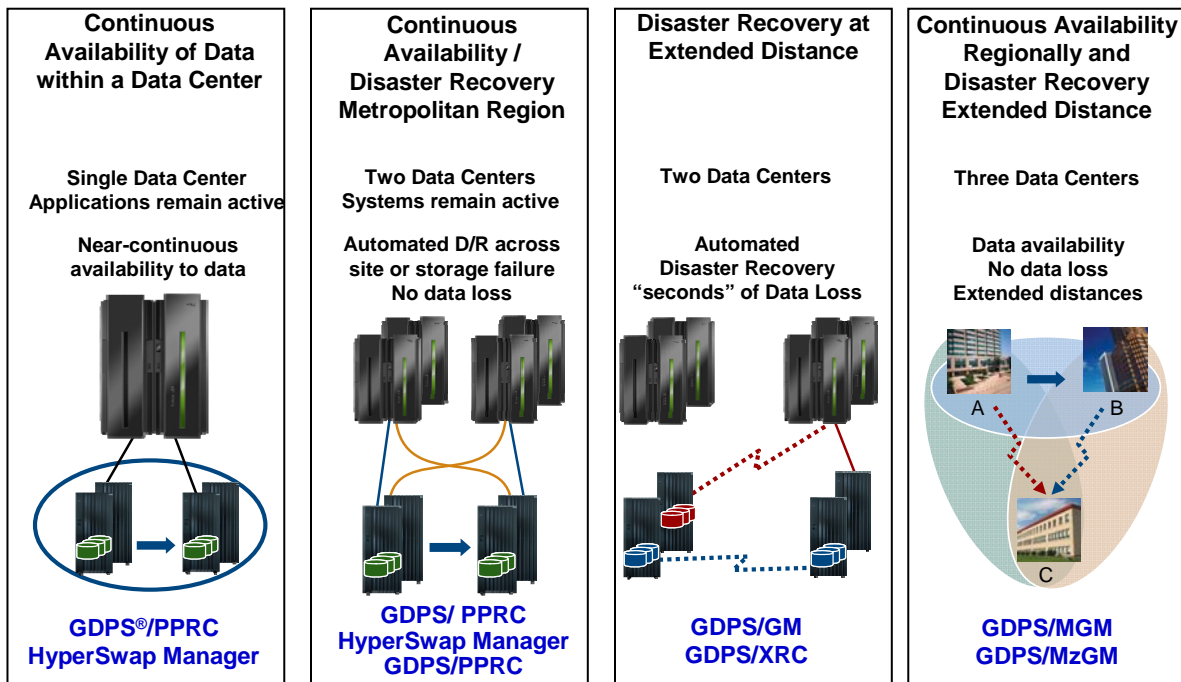
The change in the System z's approach to managing only *more exceptional exceptions* is a more subtle revolution. Previously, System z's pre-planning was far better than the crisis management found on many other platforms. Now, administrators are asked to define a worst-case scenario. System z10 will automate the reactions to anything less dire. This supports a management simplicity unmatched by other approaches.

With System z10, the exceptions alerted to human management are only when a decision has to be made (though all events are logged and available for trending and analysis). This not only improves the management of mainframe environments (which is good for IT), but increases the flexibility and responsiveness of the mainframe to unexpected business events, which is good for the business.

Diversity, Too

Special requirements often require special solutions. There's a lot of diversity with the System z10: many operating systems and many specialty engines, to start. But, sometimes more is needed beyond the more traditional hub, peer, and tier that have been evident in the past. This was true with Hoplon, a

Exhibit 3 – IBM's GDPS Offers a Broad Variety of Protection Options



Source: IBM

multi-player game provider.¹⁴ Their hybrid vision introduced combined *System z9* with *Cell BE* outboard processing. The hybrid concept has ramifications far beyond gaming. Similar capabilities are needed to support real-time market analytics within the financial industry and the medical community's desire to use real-time imaging to guide new procedures. Now, IBM's vision for *System z* and *Cell BE* places a *BladeCenter* with *Cell BE* blades adjacent to a *System z*, connected by Ethernet, with *System z* operating systems and *Cell SDKs* managing the cell blades and coordinating activity between the two environments. IBM's hybrid vision is one of linkages, not necessarily of physical integration. These outboard engines work as part of a logically-integrated, total solution, providing increased diversity and opportunities for being managed under the *System z* umbrella.

Conclusion

The new capacities and capabilities of *System z10 EC*, *z/OS*, and *DS-8000* are all very big news, but the enhanced abilities to do "more" via *System z* is the really big news. *z10 EC*'s comprehensive offering defines the future of very-large-scale, virtualized, integrated computing. **Only System z's**

pervasive and efficient virtualization and share-everything approach can give both front-side users and back-end systems the massive capacities and on-demand scalability that often is needed these days, in an energy-efficient, affordable, and secure way.

With *System z10 EC*, big iron solves big business' needs for performance, flexibility, security, and most importantly, manageability. To do more, you have to manage everything better. Being simple isn't enough, if management becomes complex and burdensome. **This big mainframe can do more and do it more efficiently and at a lower cost per unit of work – all important in the quest to do more.**



¹⁴ For more on Hoplon, see [The Clipper Group Navigator](http://www.clipper.com/research/TCG_2007065.pdf) entitled *Hybrid Solution for Xtreme Information Use - Hoplon Leverages IBM's Cell/B.C. and System z*, dated May 22, 2007, and available at http://www.clipper.com/research/TCG_2007065.pdf.

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 Authors

Anne MacFarland is Director of Enterprise Architectures and Infrastructure 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.)*

Mike Kahn is Managing Director and a founder of The Clipper Group. Mr. Kahn is a veteran of the computer industry, having spent more than three decades working on information technology, spending the last decade and a half at Clipper. For the vendor community, Mr. Kahn specializes on strategic marketing issues, especially for new and costly technologies and services, competitive analysis, and sales support. For the end-user community, he focuses on mission-critical information management decisions. Prior positions held by Mr. Kahn include: at International Data Corporation - Director of the Competitive Resource Center, Director of Consulting for the Software Research Group, and Director of the Systems Integration Program; President of Power Factor Corporation, a Boston-based electronics firm; at Honeywell Bull - Director of International Marketing and Support; at Honeywell Information Systems - Director of Marketing and Director of Strategy, Technology and Research; with Arthur D. Little, Inc. - a consultant specializing in database management systems and information resource management; and, for Intel Corporation, Mr. Kahn served in a variety of field and home office marketing management positions. Earlier, he founded and managed PRISM Associates of Ann Arbor, Michigan, a systems consulting firm specializing in data management products and applications. Mr. Kahn also managed a relational DBMS development group at The University of Michigan where he earned B.S.E. and M.S.E. degrees in industrial engineering.

- *Reach Mike Kahn via e-mail at MikeKahn@clipper.com or via phone at (781) 235-0085 Ext. 121. (Please dial "121" 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.