



IBM zEnterprise Builds Business Value in an Age of Colossal Computing

Analyst: Anne MacFarland

Management Summary

Did you ever go to a small-town, one-tent circus? You know, the kind that arrives one night, and sets up its tent and infrastructure the next day, performs a few shows for a day or two, then breaks it down and moves to the next town, maybe 40 miles away, just to do it all over again? The pace is chaotic, yet familiar from repetition; you might say that it is *well orchestrated*. It has to be! However, because there is only one ring (the main and only tent), all of the animals and acts come through one after another in a serial fashion. Now, imagine a three-ring circus. Clearly, there is more infrastructure to assemble and manage and the orchestration is much more complicated because animals and performers are in three rings at the same time. Nonetheless, it is well scripted and orderly. It has to be!

Editor's Note: This is third of three Clipper bulletins focusing on IBM's new zEnterprise Mainframe. The first two focus on (a) capturing the rewards of a diverse server environment and (b) the "speeds and feeds" and details of the new Mainframe hardware. This paper deals with the software that enables this new platform, including the new Unified Resource Manager.

Don't you wish that your datacenter would be as orderly and orchestrated as a three-ring circus? Unfortunately, it isn't and can't be, at least any more. In the days of batch processing, everything was predictable, scheduled, and orderly. There were only a handful of computers, so it was akin to managing a three-ring circus. The programs were fixed and data arrived on time, did their act, and then exited, pretty much as planned.

Unfortunately, your datacenter is more like a 1000-, 10,000-, or 100,000-ring circus, where the rings are systems images, either occupying a physical server or a slice of one (i.e., a virtual partition). However, it is even worse, because much of the work being done is unpredictable (in so many ways) and the rings vary in size and capabilities. It is a massive set up, queuing, execution, and facility recycling effort; one in which there are oh-so-many actual and potential dependencies between the many *acts* (operations) performed by your *performers* (applications and the users behind them) in your thousands of rings (systems images, mostly on virtual machines). In some instances, the virtual machines exist only for a minute or two, before being absorbed back into the pool and reallocated. No wonder that you feel like you are repeatedly being shot out of a cannon!

Simply put, this is a massive resource management problem. While it might seem attractive to manage a modest collection of servers as a singular entity, this only makes the problem worse, as you end up with dozens, hundreds, or even thousands of siloed entities to manage. Control of operations and optimization of resources become a nightmare all too familiar. **What you need is one ringmaster capable of managing the whole thing and a whole lot of support from everyone sharing the goals of optimizing the assets, time, opportunities, and the ultimate customers' (users') experiences.**

IN THIS ISSUE

- The Changing Game of IT 2
- IBM Responds with zEnterprise 2
- The IBM Surround..... 5
- Conclusion 8

This is what IBM has just announced as *zEnterprise*, a multi-dimensional colossus controlled by a single ringmaster. It consists of a larger and faster System z Mainframe (the *z196*), and, optionally, can have one or more specially-attached and specially-enhanced *BladeCenter* racks (the *zBX*) that will house selected *POWER7* and *System x* blades, and a *Unified Resource Manager (zManager)* to serve as the ringmaster. But, as we so frequently hear these days, *there is more* – more than hardware and systems management software, that is. There is so much more that you will have to read on to get the whole story.

The Changing Game of IT

In the past few years, what it means to be an enterprise datacenter has evolved. The scope and variety of workloads has proliferated far beyond what was once known as *data processing*.

- A vast variety of information is needed to produce properly-targeted business insight in a useful (usually brief) time frame.
- The extent of enterprise operations now goes far beyond the traditional corporate boundaries not just through partnerships, but also by various forms of participation by customers, potential customers – and, from time to time, by competitors.
- Frequent adjustments to business models made in an attempt to capture profit from volatile markets demand many new kinds of operational and informational datacenter agility.
- At the same time, datacenter budgets are tightly constrained, and any initiative must have not just business value but quick contribution to profitability.

To meet all of the colossal needs – informational, operational, and infrastructural – enterprise IT must become more powerful, more empowering, and less needy. This is a challenge that often appears to be a matter of draconian choices. Simplify all operations to one platform (at one extreme) or outsource and hope that all your sourcing choices are good (at another extreme). Unfortunately, most enterprises need something in the middle that works for them, most likely including a transformation of how IT resources are managed. **Thus, most seek a way to accommodate diversity within the datacenter without being crushed under the weight of managing the heterogeneity.** Three strategic steps will help.

1. **Rethink *what is done where*.** This is multi-dimensional. First, are you placing applications where they really belong and not just where it seems easiest or cheapest? Second, are you treating all server resources as if they were equal?¹ Third, are you considering the proximity of data to where it is being used? Fourth, have you looked at how much data you are moving around and duplicating? Fifth, have you fully considered how all of that data being collected is going to be used by the enterprise in real-time (as opposed to how it is going to be processed)? This list could go on, but these kinds of questions frequently reveal that the datacenter may be doing what accommodates the most requirements most easily rather than meeting the greater challenges for timely and integrated information service delivery within the enterprise and beyond.
2. **Identify common processes and make them more broadly consistent, then reuse them.** This kind of refactoring is behind both SOA and Application Modernization. It applies to many components – development tools, applets and applications, and virtualization (where hypervisors add an intermediary layer to support finer-grained resource consumption).
3. **Build support for a closeness of data and process** (moving either, when appropriate, especially to accommodate high-security requirements). Too often, our assumptions start with the status quo. Unfortunately, we tend to have gotten to the *where we are* by a series of long-forgotten requirements and possibly no-longer valid-decisions – sort of a random walk through a changing swamp. To support adjacency of data and process in environments with a high rate of change, you need to manage a lot of local variables and deal with changing requirements, if you hope to succeed. This is akin to good neighborhood planning.

IBM Responds with zEnterprise

The new IBM *zEnterprise* has leveraged all three of these strategies, and more. It supports not just more processing but, using a hybrid approach and sidecar called *zBX*, it spans a new breadth of

¹ This is a hot topic. Is it better to make all server assets identical (so that they are interchangeable and can be managed more simply) or is it better to apply the right server resources where the applications and/or data are best (most-appropriately) executed? For more on this, see [The Clipper Group Navigator](http://www.clipper.com/research/TCG2010032.pdf) dated July 22, 2010, entitled *Capturing The Rewards of Server Diversity - IBM's New Approach*, which is available at <http://www.clipper.com/research/TCG2010032.pdf>.

processor capabilities, enabling work to be done where it should be done. This breadth calls for management at a new scale. This scale, in turn, can support a more complete, heterogeneous set of datacenter operations with a much lower cost of management.

zEnterprise is, in many ways, a colossal amalgamation of infrastructure – now supporting BIPS, not just MIPS – but there is more to the story than hardware.² The rest of IBM has contributed mightily to development tool, systems software, and middleware enhancements (all of which have leveraged the three transformational strategies that were discussed on the previous page). These additions have expanded and completed how zEnterprise's unified (but heterogeneous) infrastructure can support a colossal, complex, modern enterprise. Read on to learn how all of the new zEnterprise hardware comes together through software.

zEnterprise represents a giant step forward in large-scale enterprise computing. Colossal-scale computing has been held back by the management challenges. Scaling management is a matter of offloading more of the routine to automation, often by doing it in firmware. However, with very large-scale operations of any kind, the detail continues to matter, for it is at the detail level that problems can be addressed before they grow. With zEnterprise, IBM has redesigned the instrumentation and added analytics to yield actionable information sooner.

Colossal-scale computing has also been hampered by fragmented and inconsistent approaches in areas such as development, testing, process integration, and tuning. All have grown more comprehensive in the situations that they can address, but none has been what you would call *industrial strength*. Several multi-year initiatives at IBM have transformed legacy assets and have built the accessories that are needed in both software and in their development and testing. More initiatives undoubtedly will be initiated in response to what customers will do with zEnterprise.

Many software elements transform zEnterprise from merely bigger and better to the foundation for workload optimization at a massive scale. The Unified Resource Manager, systemic role of zBX, and SAO³ provide the basis, but many other

² For more details on the hardware, see [The Clipper Group Navigator](#) dated July 22, 2010, entitled *The IBM zEnterprise System Reaches Out - Higher, Wider and Deeper*, which is available at <http://www.clipper.com/research/TCG2010033.pdf>.

³ SAO is discussed on the next page.

IBM also add to the compelling value of the offering.

IBM Unified Resource Manager – a Sea Change in Server Resource Management

The new IBM Unified Resource Manager (a.k.a. zManager) is the crux of the transformation, for it reduces or even removes the traditional need to manage by platform that has inflated the head count of enterprise datacenter operational staff. With zManager, hardware management functions are provided for *System z*, *POWER* blades (later this year), *System x* blades (first half of next year), and IBM data management appliances.

Hypervisors for the zEnterprise will be treated (shipped and serviced) as firmware – something that will speed deployment, simplify testing, and accelerate time to market. The broadened domain (across diverse server platforms) of zManager simplifies hardware problem resolution and actions, such as power on and off and call home, especially for applications that cross multiple platforms. zManager also can initiate pre-emptive actions, such as failover, and choreograph guided repair⁴.

Through its *Hardware Management Console*⁵ (HMC), zManager aggregates the management of operating systems, including all those on z196 (there are many⁶), AIX on POWER7 blades in the zBX and, in the future according to a Statement of Direction, *Linux* on System x blades on zBX. zManager runs independently of z/OS.

HMC also serves as the management access point for third-party System z management tools, such as those from CA, BMC, and other independent software providers, as well as for *IBM System Director*, *IBM Tivoli* products, and other IBM products, such as service and performance management elements from *Tivoli* and *Cognos*. Through the HMC, these management elements can address partitions, operating environments, middleware, and applications on both zBX and the z196, as well as the service management of traditional System z workloads.

Critical workloads can be prioritized on the zEnterprise, just as they have been for years on

⁴ This is a term for grooming of zManager-managed elements. Their pervasive redundancy and automated failover makes “failure” an inappropriate term.

⁵ The HMC interface may be CIM, SNMP, or REST. HMC can express, via that interface, zEnterprise-specific protocols, models, and functions.

⁶ z/OS, z/VM, z/VSE, z/TPS, and z/Linux.

System z, but offloads to the blades on zBX can be included in zManager's workload-centered management. The aggregation of platform metrics by zManager, presented through HMC, supports a system-wide view of resources, as well as an ability to drill down to specific workloads, including those running on the connected platforms.

All management elements can be brought together as role-based views, and can be presented on a single-pane-of-glass on the HMC User Interface. HMC also aggregates⁷ instrumentation data feeds and events relating to resource utilization, capacity, availability, performance, problems, and energy in the context of the deployed workloads. This single, more complete⁸ point-of-application control is a game changer.

The Systemic Role of zBX

Not all computing has the same requirements or economics. At one end of the scale, harvesting sensor metrics, such as health vital signs, requires aggregation to turn values into useful information. This situation may require privacy, but not necessarily high availability – it often is the changes in value that are important (such as when a valve's state changes from closed to open, not each instance of the value itself). By contrast, the processing of our banking transactions does merit cast-iron security and high availability, thank you very much. It requires the kind of security that z/OS's inspection of each executable⁹ can provide.

In zEnterprise, three networks allow all this to be accomplished without impairing performance. One is a *Service LAN* that connects a support module in the z enclosure with an advanced management module in z/BX. The second is an *Internal Platform Management Network (IPMN)*, a single, private, secure isolated network that facilitates management of complex topologies without affecting other (business-oriented) communications. The third, as one might expect, is a capacious (10 Gigabit Ethernet) Data Network called the *IntraEnsemble Data Network (IEDN)* to move data and programs between and among the zEnterprise servers.

The addition of zBX to the traditional Main-

frame delivers more choice and a better matching of workloads to systems – a specialization that, at large scale, becomes extremely cost effective. Because the platform consists of all IBM-certified components that are tested and packaged together, it is a faster install than is found with silos of servers.

SAO analytics – Just Another Offload

System z specialty engines have proved to be a very good idea – not just for the savings on standard engines and related software, but because of its cost effectiveness for executing frequent, repetitive tasks, in an optimized environment. Analytics (business intelligence applications) have their own patterns of process that are very different from transaction processing. To optimize analytics, IBM developed a *Smart Analytics Optimizer (SAO)* has been in beta testing since the second half of 2009.

When situated on zBX on multiple blades, SAO will offload analytic workloads from the z196, much as IBM's System z specialty engines offload work from standard engines.¹⁰ In the case of SAO, DB2 query optimizer on z/OS¹¹ is used to identify queries that are likely to be lengthy or unpredictable in terms of time to completion (or when it is uncertain whether it will complete at all). These queries are offloaded to zBX, whose blades have the vector processes to accelerate the query processing. This minimizes the distance the data has to be moved, and keeps analytics in the zEnterprise security zone. It uses several forms of extreme compression that IBM has patented, which speeds it along.

The sweet spot for SAO is a query that involves scanning a large subset of data, i.e., when aggregating data from multiple sources, particularly those using a snowflake schema of a fact table surrounded by tables, hierarchies, and other data structures. These elements must be joined in a particular fashion to produce actionable results for specific business situations, such as branch store stocking, that must be specifically targeted at local customers – to prevent loyalty erosion.

¹⁰ For More information on System z zAAP and zIIP see two issues of **The Clipper Group Navigator: zSeries Zips through JAVA with zAAP**, dated April 7, 2004, and available at <http://www.clipper.com/research/TCG2004030.pdf> and *System z9 adds zIIP to Ally with DB2 on z/OS to Better Serve the Onslaught of Business Data*, dated January 4, 2006, and available at <http://www.clipper.com/research/TCG2006006.pdf>.

¹¹ Remember that DB2's current version supports most of Oracle's special quirks. Therefore, SAO can be a vehicle for addressing hard queries across a heterogeneous application domain.

⁷ The publish/subscribe model that HMC leverages to populate its dashboards is particularly appropriate, because the metrics relevant to many roles will continue to expand and change.

⁸ It crosses all relevant domains – server, network, and storage.

⁹ This capability was built into z/OS decades ago, but has found new value in preempting certain kinds of system malware.

The offload to non-z blades may seem counterintuitive to devotees of the Mainframe, but the processing of the complex queries that support analytics is best done in an environment optimized to the task – a *specialty environment*, as it were. SAO can be seen as another form factor of a z specialty engine. It is zEnterprise's first instance of this kind of offload, but there are likely to be more. In the past few years, many new workloads have been developed with particular data or processing needs and patterns of use. One can imagine that more compelling offloads will also be delivered on zEnterprise because of its tight integration of heterogeneous platforms and advanced management. It is time for enterprises to think creatively about how they can both optimize and simplify operations with zEnterprise.

zEnterprise Slashes the Costs of Management

The capabilities of the new zEnterprise are, as could be expected, bigger. More significantly, the opportunities that zEnterprise can now address are more vast, and the way this vastness can be managed is transformed by the new key elements discussed above. The multi-platform consolidation that these elements, plus the standardization of operational and management processes that they can enable, should bring about considerable management savings.

According to IBM's estimates, the greatest savings will be in deployment, release and change management, where costs might be reduced by up to 85%, in asset management – with up to 77% in savings, and in security management – with up to a 37% reduction. These do not include the reductions in physical infrastructures, like the separate firewalls that may no longer be needed. Also, with zEnterprise the incident management costs can be reduced by about a third because of better targeted, more comprehensive monitoring and reporting. In addition, zEnterprise has kept energy consumption the same as the z10, plus there are additional power access (HVDC) and cooling options (traditional, chilled water, or fresh air (suitably filtered, of course) for further reductions.

The IBM Surround

The scope of zEnterprise has ramifications that echo far beyond IBM's hardware business. For it to be properly productized, aggressive agendas across all of IBM have to be completed, amalgamated, and tested. The (admittedly rather large) IBM Village has delivered. Let's take a look at how other initiatives within IBM Software Group build additional business value in zEnterprise,

focusing on the context of the three transformational strategies introduced earlier.

Microprocessors and Compilers

Over the past several years, IBM's microprocessor engineers have worked closely with IBM's compiler folks to create compilers for both POWER and z platforms that give a significant boost to how applications run. Starting with z10, some compiler-generated operations have been moved into the zProcessor hardware. There are 100 new instructions for the z196 processor, many of which are used to accelerate execution of compiled programs.

These compilers are simple and low-risk to install. The latest versions, in part because of the offload to hardware, give significant operational benefits – IBM estimates 30% for C, C++, and JAVA workloads. The benefits are particularly high for JAVA, for it presents a special opportunity. JAVA is, by its nature, chatty, so it offers the opportunity for an additional layer of transplant. More JVM instructions have been moved into hardware, improving JAVA performance 15% and garbage collection is up to 30%. With JAVA 7, still more execution improvements can be expected. This more efficient execution of code, when considered at the scale of a 100,000-ring circus, it aggregates to considerable efficiency.

Supporting Colossal-Scope Development and Testing

Over the past five years, Rational has been reengineering the z software stack to be based on open standards in its "Phase z" project. During the past few years, this project has enabled classic Mainframe middleware like – CICS – to work as callable Web Services. Now, Rational also has transformed z development tools to support congruent development processes across multiple platforms, leveraging the widely used, open-sourced *Eclipse* development environment. This effort has made the differences in code libraries of different platforms transparent to the application developer, which not only improves developer productivity, but also gives them new freedoms of choice in both platforms and languages for development and for execution.

Another multi-year effort has emulated the native z/OS stack and core application on Intel¹²

¹² Called the RDz Unit Test, it supports a System z operating environment on an x86 PC running Linux. It includes actual middleware and compilers, plus RDz and RTCz agents and will handle entry, debug and compile as well as unit test. It includes

processors, allowing unit testing of System z-targeted applications to be done either on an Intel workstation or on the zBX.¹³

Large-scale application testing obviously is integral with the larger scope of manageability that zEnterprise will support. IBM has developed, in another multi-year initiative, a methodology for testing of complex systems and systems-of-systems, called *CALM (Collaborative Application Lifecycle Management)*. CALM offers a way to write, evolve, and test large applications with due attention to dependencies and all the irritating details that must be addressed scrupulously.

These initiatives bring a new set of choices into play. These choices allow enterprises to build and evolve not just through situational applications, but also by evolving old and building new core processes. More near-term, zEnterprise gives complex applications a new, heterogeneous environment in which to thrive. Think of the optimizations that are possible when you can leverage both the security and speed of within-zEnterprise operations, the breadth of processor specialization, and the opportunistic offloads.

In an interview with The Clipper Group¹⁴, Martin Kennedy, Managing Director for Citigroup's enterprise systems and infrastructure, said: *I don't see any existing zLinux workloads moving off [of the Mainframe]. The hybrid [zEnterprise] will allow us to integrate [workloads] more easily. I could envision developers using tooling available only on x86 technology and deploying the application on the technology that makes that best fits service level requirements of the application – z, p, or x – or any combination of technology.*

Further Optimizing Workload Execution

DB2 10

IBM has now rewritten DB2 to leverage open standards wherever possible. Now, the new compilers will shrink the DB2 footprint, leveraging a few new chip instructions and also hardware compression of indexes. *DB2 10¹⁵ for z/OS* will support ten times more concurrent users – something required by financial institutions, for example, as more customers check their accounts more fre-

quently on their smart phones.

DB2 10 also introduces the concept of *temporal data management*. Details change on each customer's master record. In the old management paradigm, these changes had to be handled by an administrator updating and interrupting the application. With temporal data support, DB2 surfaces the small changes to the DB2 engine via a *Change Data Store*, allowing the database to run unaffected until later, when the update is applied.

WebSphere

Applications remain the crown princes of IT systems (as they have become on cell phones) – they are the instruments by which an enterprise thrives (or dies). No matter how or from where they are sourced, users have expectations about what the information should look like and simple pruning to fit can expunge details that really matter. It is important to take a systemic approach to the accessories that keep quality and consistency of results intact. IBM offers a broad array of solutions for zEnterprise.

When *WebSphere Application Server (WSAS)* is co-located with data in application subsystems, nesting JAVA logic into z/OS or z/Linux with *Hipersockets'* communication to the data improves throughout, cuts CPU use, and functions more efficiently than traditional Web Services. *WebSphere Process Server (WSPS)* benefits from similar placement. WebSphere is the basis of the *zBX Application Server*.

More new WebSphere products, scheduled for release in the second half of 2010, will enrich further the hybrid space that zEnterprise will provide. What such software can do there is based on many new areas of commonality that IBM now supports. A look at these commonalities will show the richness and breadth of these opportunities.

Data Warehouse

The *System z9600 Data Warehouse* (targeted at 4-10 TB+) moves from being a back-end element to encompass the full value cycle from data acquisition to data presentation, via either z/OS or z/Linux. This unity supports both consistent access to data (which is good) and consistent answers to queries (which is even better).

IBM sees data marts, such as the *DB2 10 Change Data Store*, to be a critical part of heterogeneous hybrid hardware and software environments. Data marts can treat data that has a particular context in the hybrid with the right qualities of service but no more (essentially putting a cap on service delivery). They can keep external

z/OS, IMS, DB2, CICS, PL/1, COBOL, and C++, as well as a debugging tool.

¹³ This complexity is one reason why that the zBX's inclusion of System x blades will not take place until 2011.

¹⁴ On July 16, 2010.

¹⁵ Remember that DB2 supports XML stored natively and understands many Oracle structures.

“dirty” data secure and also isolated from more highly-cleansed data, so that the latter is not degraded. This will become more important as use of sensor data becomes a bigger part of business operations.

InfoSphere Information Management Software

In the past year, more of *IBM InfoSphere*'s software has been ported to System z, both under z/OS and under z/Linux. The many sophisticated data services that comprise InfoSphere will be of obvious use in the colossal operations that zEnterprise supports. *InfoSphere Foundation* supports wider reuse of data¹⁶ through its discovery, collection, and classification of relationships in structured and unstructured data, plus its many data handler transforms that let data meet the expectations of the applications that consume it. In supporting enterprise operational efficiencies, zEnterprise applications are going to consume a wide variety of data sources.

Of particular interest, in terms of managing colossal environments, are its *Master Data Management*, *Cognos Performance Analytics*, and *SPSS Predictive Analytics* offerings. All offer ways to manage and leverage the reality of a very large datacenter. *Cognos* supports analysis, dashboarding, and scorecards for a single, vast SOA – something that zEnterprise and zBX embody. *SPSS*'s predictive capabilities let organizations leverage the cadence of expected events.

The wealth of IBM products for everything from data mining to customer analytics to eDiscovery and fraud detection take on an extended relevance, especially when they can be applied on zEnterprise across z196 and zBX servers and collect data from multiple, related business processes. This is a vision for a single, colossal, yet unified datacenter.

Martin Kennedy (of Citigroup) commented further: *Our goal is to work toward a single secure copy of master data, thus avoiding the complications of offloading and managing multiple copies. The zEnterprise architecture facilitates moving forward toward that goal because of its high degree of integration.* He added: *We would also expect storage savings as a result of these improvements.*

z/OS 1.12

zEnterprise leverages its z/OS heritage as both

a secure control system *par excellence* and as the seat of management. Until this version, z/OS was fairly self-referential. Now, it too has been reworked to be more congruent with AIX and Linux and to be more manageable by IBM *Tivoli* software and by zManager. Other enhancements include the following.

- IBM now offers *Active-Active Sysplex* over unlimited distance. Sysplex operations benefit from automated actions that preserve availability.
- z/OS supports a new concept of *Continuous Batch Processing* – somewhat like the continuous data protection for data with which storage folks have become familiar. It deploys small pieces of batch opportunistically and tracks the whole process to completion. This removes (or reduces) the need for a batch window and assures that all batch processes are completed within a specified target. This is part of a better traceability that the extended reach of zManager supports.
- *Extended Address Volumes* for stored data results in fewer, larger volumes to manage and is now extended to more data types. This is not new, but the size of the volumes has been extended, again.
- Provisioning of resources by z/OS can be workload driven, based on CICS and IMS monitoring data.
- z/OS 1.12 has improved the management of data fragmentation, particularly *VSAM*, using a new *Control Area Reclaim Facility*.
- Finally, z/OS 1.12 supports *Elliptic Curve Cryptography (ECC)*, *Internet Key Exchange Version 3 (IKEV2)*, and *Federal Information Processing Standard (FIPS) 140-2*. These keep zEnterprise at the forefront of securing enterprise data.

Tivoli Products for Colossal Computing

At zEnterprise launch, Tivoli products have been developed and enhanced to support the heterogeneous scope of the new zEnterprise, including those listed below. Of course, many Tivoli products, such as *Tivoli Access Manager* and *Tivoli Federated Identity Manager*, always have supported a heterogeneous infrastructure. Other products, such as *Tivoli zSecure Manager for RACF z/VM*, are specific to z/OS and z/VM on the z196.

The following products comprise *IBM Service Management System for zEnterprise*. In time, these products will also apply to zLinux and to the AIX blades in the zBX.

¹⁶ This wider reuse is a critical part of greater organizational transparency. This is not just good for datacenter management, but plays a part in traceability (as in farm-to-fork) that is becoming mandated in many industries.

- *Tivoli Asset and Financial Management for zEnterprise* gives a single view across heterogeneous resources of software usage, compliance, and resource usage. It also provides a centralized view of requirements and costs.
- *Tivoli Application Management for zEnterprise* leverages TCAM (*Tivoli Composite Application Manager*) to identify bottlenecks.
- *Tivoli Application Resilience for zEnterprise* provides automated high availability and disaster recovery, and operations across multi-tier, composite application environments. It supports a business-process view.
- *Tivoli Service Request Manager* has obvious importance in orchestrating the distributed applications that zEnterprise supports.

In addition, z/OS continues to support a full array of outward-facing services, such as its *PKI Services* (a full certificate authority solution).

Conclusion

Just as multiplexing in telephony greatly increased its capacity and affordability, zEnterprise takes System z's share-everything strategy to a higher level. The colossal IT that zEnterprise can support will enable enterprises to re-imagine how and where business process can work best and most economically. The changed economics it supports is a matter of creating a larger addressable space and leveraging common services. The deft management that comes with the zEnterprise paradigm supports quick identification of exceptions and isolating their treatment, through either fast tracking or remediation.

Enterprise computing has always wanted both the economies of scale and tunable qualities of service. In the past, it has been a trade-off. zEnterprise offers the ability to do both without the consequence of high management costs. Its paradigm will suit any sourcing model, including cloud, but its business ramifications are useful particularly to the complex, interactive business services of the large enterprise. This kind of colossus gives truly a lot to appreciate and is so much more approachable than a 100,000-ring circus. Consider how zEnterprise can optimize where and how you are spending your IT budget.



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 a Senior Contributing Analyst 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 in 2001 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. 128. (Please dial “128” 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, 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.