

Capturing The Rewards of Server Diversity — IBM's New Approach

Analyst: Mike Kahn

Management Summary

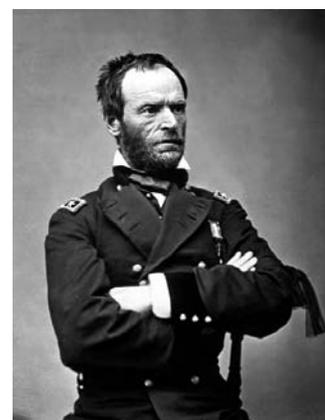
War is hell, said Civil War General William Tecumseh Sherman. **So is running an enterprise datacenter, but you already know that.** By looking at a couple of related military concepts, we can gain much insight into winning on the enterprise IT battlefield.

First, let's look at the problem of staffing a military. There are two approaches. One is for the *universal soldier*. All universal soldiers receive the same training, so that they can be deployed interchangeably. These folks are generalists, capable of being deployed anywhere (from the Arctic to the desert to a submarine) and performing standard military functions wherever they are placed. The other is for the *specialist soldier*. While they all may have similar basic training, it is only within their specialty that they excel, e.g., driving tanks, military police, first-to-arrive Marines, etc. Most of the time, cooks are not interchangeable with helicopter pilots. These folks are specialists that have been optimized (trained, equipped and otherwise prepared) to do certain tasks very well.

Second, let's consider provisioning an air force. Even in the best of situations, it is difficult to have one plane for all purposes. Almost every attempt to have a multiple-use military aircraft (like the Joint Strike Fighter) has resulted in serious real-world problems, due to too many requirements and compromises. Different assets are needed to fly reconnaissance, to provide protective cover, to transport personnel, equipment, and supplies, to drop bombs, etc.

It is rare for "one-for-all" building block (i.e., a universal solution) to work really well. Yet that is exactly what most datacenters are trying to do under the mantra of simplification by standardization, more specifically, standardization on x86 servers. As you might expect, the IT world generally is divided into *universal believers* (in this universal solution) and *diversity believers* (i.e., those who value specialization). To simplify our discussion, let's exclude from the latter category those datacenters that are "historically diverse" but wish they weren't. The opposing positions of these different believers boils down to a religious war, of sorts, with a clear dividing line between their fundamental beliefs.

For all intent and purposes, there is only one vendor whose server product lines are sufficiently diverse to be labeled as a strong proponent of diversity (a.k.a. *heterogeneity*) and that is IBM



Photograph of Gen. William Tecumseh Sherman (circa 1865) from The National Archives

Editor's Note: This is first of three Clipper bulletins focusing on IBM's new zEnterprise Mainframe. The other two focus on (a) the improvements and enhancements to the platform and (b) the software that enables this new platform, especially the Unified Resource Manager.

IN THIS ISSUE

| | |
|--|---|
| ➤ The Benefits of Diversity..... | 2 |
| ➤ The Challenges of Diversity..... | 3 |
| ➤ IBM's Approach for Optimizing the Diverse Datacenter | 4 |
| ➤ IBM's Unified Resource Manager | 6 |
| ➤ Conclusion | 8 |

(with *x86*, *POWER* and *System z* (Mainframe) chipsets providing heterogeneous server product lines). Once, this was the norm among leading computer vendors. Yes, there are other vendors (like HP) that still have a historically-diverse server product lines (principally to run *HP-UX* and non-stop applications on *Itanium*), but even HP says that it wants to deliver convergence on *x86*. New entrants into this space (like the *Virtual Computing Environment* from Cisco, EMC, and VMware) start with the assumption that *x86* architecture is the only answer needed. *So who is right?*

That's a very tough question, especially since the only correct answer is: *it depends*. It depends on many factors, the most important of which are the goals and objectives (results) that the enterprise seeks from its IT efforts and the risks that the enterprise needs to ameliorate or avoid. Let me explain by example. If you have modest goals and objectives (like getting from one city to another at reasonable expense), then you may have multiple transportation alternatives, many of which look remarkably similar (like non-stop flights on competing airlines). While you may have a personal preference (or allegiance), the flights are fundamentally equivalent, except possibly for schedule. However, if you are the CEO and you or your company has concerns for your safety and/or your productivity in flight, then a private jet may be the answer. This is why we almost always are offered a selection from multiple choices, usually spanning a range of costs and benefits. However, if you never had the option of flying anything but coach at a discount rate, you might not be able to articulate or justify the benefits that come with the private jet.

This is where we are today. **Some enterprise datacenters envision their entire world being satisfied by *x86* servers, while others have tasted diversity and recognize its benefits.** This bulletin is for the believers in diversity. It assumes that the reader values server diversity (for one or more of many reasons to be discussed) and wants to leverage what that diversity has to offer. To make this discussion even more relevant, IBM has just announced a multi-architecture workload manager, with a new Mainframe at the lead. Read on to learn more.

The Benefits of Diversity

The idiom "Different strokes for different folks" is right on target. People's needs vary. If there is a universal solution that does what everyone needs, then it becomes a killer in that category. Regarding the *Model T*, Henry Ford said, "Any

customer can have a car painted *any colour* that he wants *so long* as it is *black*."¹ For those of you old enough to remember, AT&T (the original, regulated phone company) also saw no reason to offer (lease, the only way, as well) phones in any other color than black, until they figured out that people would pay extra for a different color.

In general, choices evolve (expand) when there is a lack of choices. Alternatives tend to shrink when there is a surplus of choices, especially when the differences are trivial.

In the computer industry, enterprise computing has gone from less than two handfuls of architectural choices (in the 1960s) to an oversupply of choices (in the 1980s and 1990s) back to a handful of choices today. The real question is: *Does anybody really care anymore?*

This is a very important question in an era of cloud delivery vehicles, both within and outside the datacenter. The cloud makes transparent the architecture(s) of the servers beneath it, or at least it should. There are two ways that this can be accomplished. First, it can reduce everything to a common denominator (say a single *Linux* operating environment) or to a common virtualization layer (which masks the peculiarities of the processors being used). Let's call this approach to transparency *infrastructure commonization*. Second, the diversity of the architectures can be preserved but the management of applications, data, security, etc., can be made common; let's call this *managed commonization*.

Can you do one without the other? Yes. Can you do them both at the same time? Yes. However, this is where the universal believers differ from the diversity believers. The former see no architectural advantages from diversity beneath the cloud², while the latter see operational and other advantages³. Thus, the true value of diversity seems to depend on the eye of the beholder.

As an example, some see each of us driving the same vehicle while others advocate a diversity of choices. Maybe it is a distinction between the *rich* and the *not-so-rich*. The rich (in this case the largest of enterprises) can afford the benefits that

¹ *My Life and Work* - Henry Ford, by Henry Ford and Samuel Crowther, 1922 (Garden City Publishing Company).

² Or they are willing to abandon the potential benefits of diversity (as unaffordable or without sufficient ROI).

³ Diversity believers not only believe that special-purpose servers can be affordable but also likely believe that they can contribute to one or more forms of operational and/or competitive advantage.

diversity might bring while the not-so-rich have to accept what they can afford (to procure, use, and manage), usually through some form of simplification.

You might be asking, *Is simplification inherently bad and diversity inherently good? No!* Each has there own place in the IT toolbox. Larger enterprises also seek simplification but it doesn't necessarily mean winnowing down to a single infrastructure building block.

Here's an analogy. Every adult should have a will. (No debating on this point.) Most people need only a simple will, because they do not have the complications that wealth brings to the situation. However, some (the wealthy, to make a broad generalization) need a more complicated will that reflects personalized estate planning that incorporates many parameters. Their wealth-transfer optimization goals are more complex because their needs and situation are more complex. The diversity of estate-planning choices is worthwhile to them. A good estate attorney will optimize from the many possibilities to achieve the best possible outcomes. Simplification is not the ultimate objective for the wealthy. For them, it is all about optimizing their wealth.

Infrastructure Commonization

If you never have had a will and you ask an attorney to prepare one for you, he or she likely will dip into their bag of tricks (previously created clauses) to build a will that is appropriate for you. If you are 22, unmarried, without children, and the only thing that you own is your car, maybe the entire will is *canned* or pre-assembled. There is nothing wrong with using this preconceived offering, just like there is nothing inherently wrong with buying your server cycles from a cloud on the web. It may be the right thing for you at this time. Thus, the goal of infrastructure commonization is worthwhile. Certainly, it is easier to manage a fleet of the same vehicles that are made from the same components by the same manufacturer. This is good, as long as you can satisfy all of your needs with that same vehicle. However, if you need delivery vans for the city use and bigger trucks for use between cities, then absolute simplification does not meet your requirements. At this point, you have two choices: (1) Stop managing your own delivery fleet and turn it over to someone in that business or (2) Step up to a diverse fleet and seek to manage it well and optimize its utility.

The best rule of thumb may be to do as little as possible (i.e., to stay on the standard path), as long as it works for you and your

business. The corollary is that you need to do what will allow (increase the likelihood of) meeting your business goals, while mitigating the inherent risks and optimizing your resources.

Management Commonization

So, maybe you recognize the need for a fleet of diverse trucks dedicated to your business. While it might be nice to buy all of the various trucks from the same manufacturer, it might reduce your ability to negotiate the best deals. Or it might not even be possible. Because of your requirements, you may need what a diversity of suppliers has to offer. Yes, your life will be more complicated, but that may be what your business requires.

This is an old story, not limited to delivery fleets or even server assets. **While you do what you can do to assemble the right fleet and keep it in good repair, the more important measurement is what you accomplish with these resources, which often is measured by how well you manage and optimize these resources.** If you can manage and optimize the diverse resource as if it is one collection, then you have a good chance of sleeping at night.

The Challenges of Diversity

Obviously, none of this is without many challenges. If you can get by with a simplified solution or, even better, by tossing the problem over the fence into someone else's *ready, willing and capable arms*, then by all means do so. However, if you've read this far, this probably isn't a viable option.

While there are many challenges, some of which are exacerbated by diversity, four float to the top of the list. Each is discussed below.

1. Rationalizing the Diversity

You need to find out why you are diverse in your server infrastructure. *Is it a historical accident? Was it intentional? If so, how and why?* What you need to do is recognize the *different strokes* that your servers provide to *different folks* (i.e., users and their applications). You need to segment your applications and users according to their needs and priorities. If high, real-time transaction volumes with the ultimate in security are the requirements, then certain platforms and environments come to mind. If you have complex computations to run regularly or continuously, then other platforms come to mind. If you are just doing file or web serving to an indiscriminating audience, then the simplest of servers may be all that you need. **Simply put, you need to rationalize**

the different strokes for different folks.

2. Dealing with the Differences

Diversity has its benefits and it has its complications. This is true whether talking about assets or people. It is always simpler and thus easier if everything or everyone is the same. Some countries and cultures pride themselves on their homogeneity. Other countries pride themselves on their diversity and work to:

- (a) Identify the potential
- (b) Minimize the challenges
- (c) Maximize the synergies, and
- (d) Reap the benefits.

For servers, this is exactly what you need to do. The question is: *How?*

3. Managing the Diversity

The answer is managing the diversity in a common way. If work (a.k.a. applications) in need of a place to run can be assigned to servers (or virtual partitions thereof) by a common workload manager (preferably driven by predefined policies executed in real time), then much is possible. If you try to do this with a bevy of server administrators in real time operating from spreadsheets, then you are facing a daunting task.

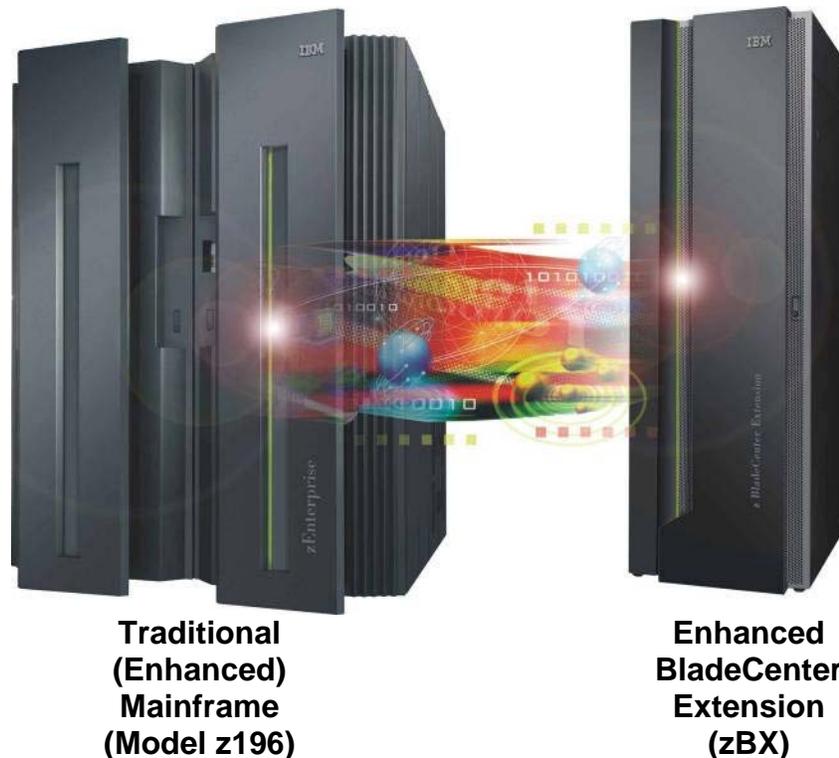
4. Optimizing to Benefit the Business

If you are able to get work done when it needs to be done and where it should be done, you have reached an important milestone, at least for the moment. **The problem is that business is proceeding relentlessly, growing rapidly, and, of course, doing so in ways that are hard to predict.** You need to optimize all of your server assets in real-time – *as things are changing!* In the past, you might have over-allocated everywhere, so that you didn't come up short when something unpredictable happened. **We now all know that stranded (i.e., reserved but not used) IT assets are wasteful in many ways, so we need to optimize intelligently on the fly, automatically, to meet the most important needs of the business.**

IBM's Approach for Optimizing the Diverse Datacenter

It's been a long introduction, but finally I am getting to IBM and their recent announcements. It might be best to begin by saying again that **IBM's answer to the resource optimization problem is centered on their new Mainframe, the zEnterprise.** Its technical specifications and new features will be discussed in detail in a separate bulletin. It

Exhibit 1 — zEnterprise Hardware



Source: IBM

Exhibit 2 – Summing the Pieces

$$\begin{aligned}
 \text{zEnterprise} = & \quad \text{“Enterprise-Class” Mainframe (z196)} \\
 & + \text{ BladeCenter Extension (zBX)} \\
 & + \text{ Unified Resource Manager (zManager)}
 \end{aligned}$$

can be summarized by saying that the z196 is “enterprise class” in its design, with 96 5.2 GHz cores, 80 of which are customer usable, with the rest being used for systems management, acceleration and spares. **While zEnterprise is a larger than the previous top-of-the-line z10 Enterprise Class, it is not its larger capacities that address the challenge of managing diverse servers in the large enterprise datacenter.**

Many things make these other differences possible, but three will be discussed in this bulletin. While zEnterprise might, to some, look like a z11 (or what was expected to be the z11, i.e., the successor to the z10), that is only part of what makes a zEnterprise something more. However, if you have a z10EC and just need a bigger or faster Mainframe, zEnterprise is what you need.

Nonetheless, there is more to zEnterprise than just being an upsized z10EC, although some will chose to use it in this way. New, and very important to the focus of this paper, is the system portrayed on the right side of Exhibit 1, at the bottom of the previous page, with the traditional Mainframe on the left. The server on the right is a special-purpose IBM zEnterprise BladeCenter Extension (zBX), capable of housing selected Power System and System x blades. These are connected to the traditional Mainframe by three networks, two for systems management and the other a private, secure 10GB Ethernet for moving data around. The whole can be managed coherently by the new Unified Resource Manager (zManager), which sits on the traditional Mainframe and runs independent of z/OS⁴. The three parts – the Mainframe server, the BladeCenter Extension server, and the zManager (system resource management and workload control software) – constitute the zEnterprise. (See Exhibit 2, above, for the key equation.)

That’s all of the zEnterprise’s hardware and operating systems details that will be presented in this bulletin⁵, as the focus turns to the managing

and optimizing a diverse set of server platforms and environments from a single point of control.

zEnterprise’s Applicability – In a Nutshell

Here’s what zEnterprise does to solve the problem of the diverse datacenter.⁶ First, let’s slip into some IBM-speak and talk about *workloads*. Application workloads are business functions that need to be run, possibly at specific times of the day, or in response to on-demand requests, or possibly, whenever the demand is light (and there are idle server assets). Think of a workload as an order that needs to be satisfied. In the real world, some orders are more important than others. For example, you might have three workload priorities: *mission-critical*, *business-critical*, and *less-critical*. In addition, there may be explicit or implicit dependencies between workloads. For example, one program may need to run before another or might create a need for another program that is to follow. You need a policy-driven workload manager to supervise all of this, at the least, and to optimize the resource utilization and scheduling, if it is really good. Traditionally, in a diverse server environment, each of the server classes (and, possibly, each of the servers) is an island unto itself, with its own workload manager. The isolation happens because the workload managers don’t cooperate with alien (other) workload managers very well, if at all. This creates an administrative nightmare and the cry for making everything the same.

With zEnterprise’s new three-part solution, one workload manager (zManager) can manage workloads on:

- The traditional Mainframe (including applications running on z/OS, z/VM, z/TPS, and on IFLs (the Mainframe specialty engine for Linux),
- Later this year, on AIX (on POWER blades) in the zBX and,
- Even later (first half of next year), on Linux (on x86 blades) in the zBX.

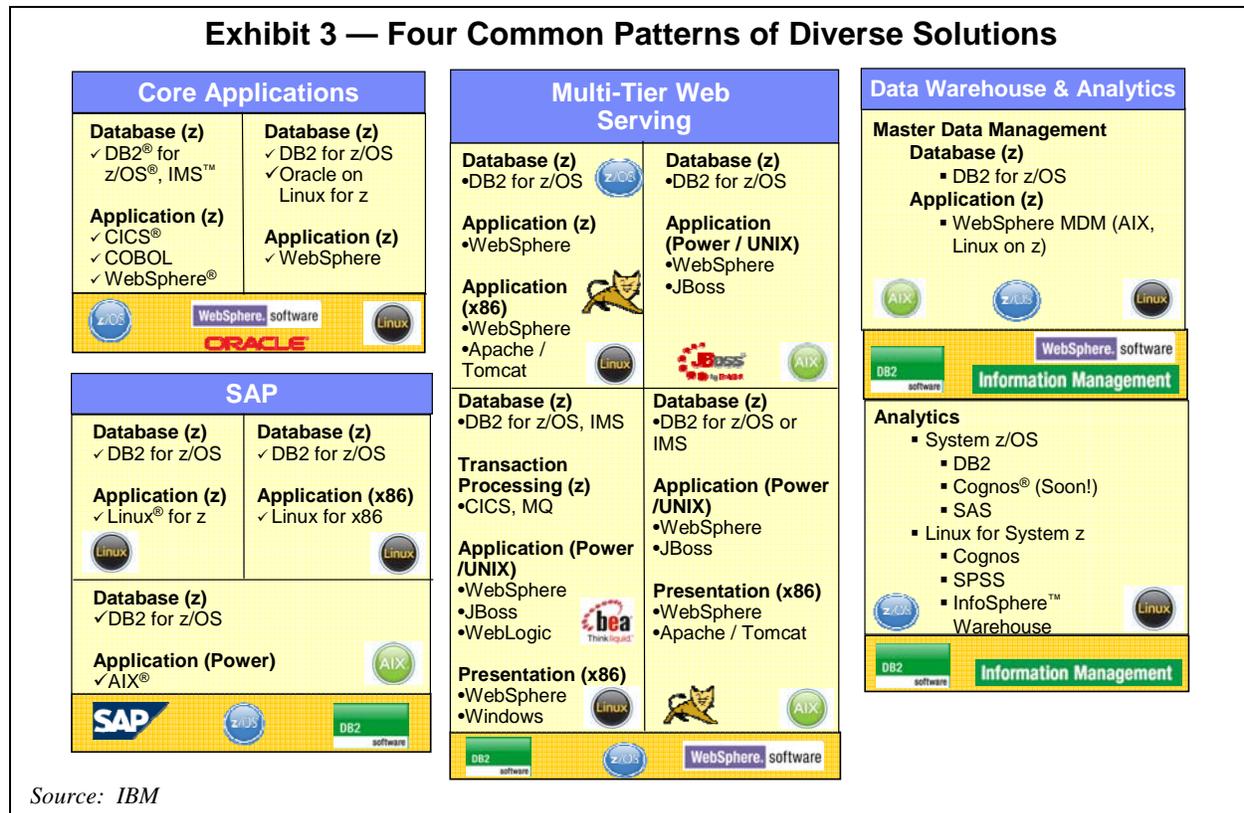
IBM is considering other operating environments,

⁴ zManager is integrated as part of the system z firmware, which includes its hypervisors, SE, and HMC appliances.

⁵ 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>.

⁶ Recognize, up front, that this is an all-IBM solution and that it will be released in stages. In addition, you need to acquire all three pieces shown in Exhibit 2 in order to obtain the benefits described.

Exhibit 3 — Four Common Patterns of Diverse Solutions



Source: IBM

but that is what has been announced for the next year.

You might imagine a number of ways that you would like to take advantage of this diverse architecture. IBM has identified four *common patterns of diverse use* that are applicable to many large enterprises.

- Core Mainframe Applications
- SAP
- Multi-Tier Web Serving
- Data Warehouse and Analytics (Business Intelligence)

These are shown, with software on different platforms and operating systems, in Exhibit 3 on the next page. **What they have in common is the use of DB2 on System z for mission-critical applications, with supporting applications running on System z IFLs, Power (AIX), and/or x86 (Linux).**

As stated earlier, if you are a believer in universal (x86) platforms, this isn't for you. If you don't use DB2, your circumstances probably don't fit one of these patterns very well.⁷ If you rely on Windows servers, IBM has not yet committed to including Windows on x86 as part of what zManager manages.

⁷ However, do note on Exhibit 3 that Oracle DBMS running on an IFL (i.e., on Linux) is one of the core scenarios.

However, if you have IBM Mainframes and DBMS (DB2 and/or IMS) at the heart of your datacenter, surrounded by applications running on AIX on POWER and/or Linux on x86, the single point of control offered by zEnterprise may be just what you seek. There are two ways to look at this offering.

1. You may already be doing this, without the benefits of centralized management and control via the zManager.
2. Your crystal ball shows you a future in which some application delivery and/or acceleration will be done on platforms other than the Mainframe, because they are better suited for the task, in terms of delivery, performance, and/or cost.⁸

As Exhibit 3 (above) shows, these tend to be complex two- and three-tier applications. What is implied by the housing of important information on the Mainframe (in DB2 and IMS) is that they are mission-critical. The databases tend to be large and deeply integrated into the business of the enterprise. **Failure to protect and deliver enter-**

⁸ AIX on Power blades may offer accelerated execution of computational applications (like financial modeling) or for middleware execution (like WebSphere or JBoss). In addition, the dynamic virtualization of PowerVM may offer superior partitioning than is available on x86. Linux on x86 may offer the lowest cost for presentation to users.

prise information (in a timely manner) is not an option. Doing it well is the only option, which now is made easier by zManager.

IBM's Unified Resource Manager

Let's take this step by step. Today's enterprise datacenter might be described as *virtualized everything*, and for good reason. Virtualization makes it easier to manage and optimize the IT infrastructure that "hides" behind the virtualization.⁹ By hiding the physicality, it is easier to manage. By being able to partition the resource into more usable, yet isolatable *chunks*, it is easier to maximize the utilization. However, we're not talking about everything today, just servers, so we need to start with server virtualization.

Server virtualization is like a hotel. Rooms are rented (occupied) for a short period, then cleaned up for the next occupant. Of course, not all rooms are equal. Some are bigger than others, with more amenities and services provided. There are some shared facilities, like exercise centers and restaurants, but the business reason for existence is the hotel rooms. To the occupant, his or her room is a private enclave, separated from all other hotel rooms. In general, terms, rooms are occupied on a daily rental basis. The goal of the hotel managers is to maximize room rentals.

The same is true of servers. The goal of IT managers is to maximize server utilization, while still satisfying its finicky customers, each with their different requirements. However, in a virtualized environment, the rooms can change size, like taking two adjoining rooms and making it into one larger room. (This is accomplished by changing the partition size, in terms of processing power, memory, bandwidth, etc.) More importantly, occupancy can be much more transient, i.e., the partition can be occupied for seconds, minutes, or hours. Thus, with server virtualization, there are a lot (of rooms) to manage, a lot more options (variables) with which to deal, and there can be many changes in occupancy, such as when a more important need arises (or event occurs). While there are many technical ways to sum this up, if you have thousands of rooms (partitions) across many hotels (servers) to manage in real time, you have a real challenge on your hands. The *flexible partition* is the unit to be managed and optimized (from the manager's perspective), although s/he must never lose touch with the qualities of services being delivered to the real customers – *applications* and their *users*, and the *authorities* that have over-

sight responsibilities.

Each flexible partition needs to be properly sized and prepared for the next occupant. Once this is done, there should be a "picture-perfect" image waiting for the next occupant. The key word here is *image*. The image can be described as a defined set of hardware, software, and networking resources that have been isolated and furnished in a specific way. It might also be called a *virtual image*, but in the virtualized datacenter, all images are virtual, so we'll just call them *images*.

What we need is a way to prepare and provision images for occupancy, to monitor and service them while they are occupied, and to return the resources behind the image back into the pool of available resources that are waiting for the next occupant. If the requirements change, we may need to adjust the resources or service levels – on the fly. We need to do this across heterogeneous (diverse) servers (hotels). If this hotel was a dormitory, where all of the rooms were the same and all of the services were equal all of the time, this would be easier to manage, because it is inherently simpler. If we have many classes of hotels (services) offering a variety of qualities of service and cost (images on z/Architecture, POWER architecture, and x86 architecture), then management and optimization is more complicated. To make this easier, we need a way manage and optimize images in real time, so that they can be occupied and the underlying resources maximized.¹⁰ That's where zEnterprise and the zManager come into play.

zEnterprise (encompassing a traditional Mainframe and a zEnterprise BladeCenter Extension) provide a diverse set of choices for configuring and using images (across many physical server topologies, up to ensembles¹¹ of many zEnterprises). **zManager provides the vehicle for defining, provisioning, managing, and reclaiming the images in real time. It makes possible what was either very difficult (because it is challenging even in a siloed or homogenous environment) or impossible (because of the diversity).** The details are important but too voluminous to include in this introductory bulletin.¹² You might

¹⁰ Stated another way, what we need is a management tool for managing diverse virtualization layers, so that we can provision and optimize virtual machines (images) on various hardware in a common way.

¹¹ Centrally managed clusters of zEnterprises.

¹² For more on zManager, see [The Clipper Group Navigator](http://www.clipper.com/research/TCG2010034.pdf) dated July 22, 2010, entitled *IBM zEnterprise Builds Business Value - in an Age of Colossal Computing*, which is available at <http://www.clipper.com/research/TCG2010034.pdf>.

⁹ The Mainframe does this better than any other server platform.

think of zManager as the virtualized air-flight control system that we wish we had in the United States. More flights (thousands) can fly closer to each other (resulting in higher utilization of air space), while reducing the risks and getting the job done more efficiently. **Restating that in data-center terms – *more applications will be running on less hardware, while maximizing the outcomes according to policies and enterprise goals. Isn't this the Holy Grail that you desperately seek?***

In an interview with The Clipper Group¹³, Martin Kennedy, Managing Director for Citigroup's enterprise systems and infrastructure, said: *We all agree we need heterogeneous environments. He added: The right way to look at it is that we try to figure out the best place to put every new app. Today, there is no perfect place. Now [with zEnterprise], we have the potential to look at a workload and put pieces on the right platform. We hope to avoid the compromises inherent in hosting an entire app on a single type of technology.*

Conclusion

If you just want more processing power than a z10EC can deliver, you need to look at the enhanced Mainframe of zEnterprise, the z196. It is full of new features and capabilities.

If you believe in server diversity (especially on Mainframes (z/OS, z/VM, Linux on IFLs, etc.), Power (AIX), and x86 (Linux)) and don't want to operate them as independent silos of computing, zEnterprise and the new BladeCenter Extension are for you. It breaks new ground for tightly-coupled diverse server infrastructure.

If you want the single point of management and control of your virtual images on zEnterprise, you need zManager. It goes where none has gone before.

Each part of zEnterprise is a significant improvement in the state-of-the-art. All are worthy of in-depth consideration, if you believe in diversity and IBM.

General Sherman also said *A battery of field artillery is worth a thousand muskets.* The same might be said of IBM's zEnterprise. Better check this out, and be sure to read Clipper's two other bulletins on zEnterprise, also published today.



¹³ On July 16, 2010.

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

Mike Kahn is Managing Director and a cofounder of The Clipper Group. Mr. Kahn is a veteran of the computer industry, having spent more than four decades working on information technology, spending the last 17 years 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; at Power Factor Corporation, a Boston-based electronics startup – President; at Honeywell Bull - Director of International Marketing and Support; at Honeywell Information Systems - Director of Marketing and Director of Strategy, Technology and Research; at Arthur D. Little, Inc. - a consultant specializing in database management systems and information resource management; and, at 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 Mike.Kahn@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, 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.