CLABBY ANALYTICS

Advisory

Why the *x86 Platform* That You Choose Is Crucial When Deploying Virtualization Software

Executive Summary

Virtualization is all about the logical pooling of physical systems, storage, and network resources. By pooling these resources, enterprises can:

- Increase the utilization rate of the information systems they already own;
- Reduce expenditure for redundant servers (because malfunctioning servers can fail-over to a resource pool instead of having to fail over to a redundant server);
- Reduce human-related management costs (through centralization and automation);
- Reduce energy costs (because fewer servers are needed as existing and future servers do more work); and,
- Lessen costs related to software development and testing (virtualization enables rapid set-up/provisioning of test environments).

Accordingly, enterprises that virtualize their information systems environments can expect to see significantly reduced system acquisition and operational costs.

But, information technology (IT) executives who choose to virtualize their information systems environments need to also consider that virtualization places additional demands on underlying hardware platforms and related manageability. How? Consider these points:

- *Reliability* — by increasing the utilization rate of servers (from, for instance 10% to 60%), the number of applications that can be run on a given server increases. And should one of these heavily utilized, virtualized servers fail, the impact of that failure can be huge in terms of lost productivity, lost revenue, and lost opportunity.
- *Scalability* — application characteristics dictate which scaling approach should be used (for instance, individual applications that need access to very large amounts of memory and that can exploit 4 processor or greater symmetrical multi-processing are best served by scale-up platform architectures; whereas applications that process small threads of information are best served by scale-out architecture).
- *Performance* — x86 microprocessors were not initially designed for virtualization. Accordingly, virtualization creates more work for x86 processors. To reduce this additional strain, some vendors have built "off-load engines" that handle virtualization pre-processing (IBM and Hitachi, for instance). By doing this, these vendors are able to support more processor cores per platform — as well as larger virtual machines — enabling support for a broader application set than competitors.
- *Energy Efficiency* — virtualization helps reduce energy costs because fewer servers can be used to do more work. But systems designs also play an important role in energy use reduction. For instance, IBM's iDataPlex design (a turnkey x86 large-

scale integrated rack environment) can process massive x86 workloads while drawing comparatively less energy than other x86 server designs (iDataPlex uses fewer components and advanced airflow techniques to reduce energy draw).

- *System Management* — virtualization adds the management of logical "virtual machines" (as opposed to the management of physical machines) to the tasks that IT managers/administrators must perform.

In this *Advisory*, *Clabby Analytics* describes why your choice of an x86 platform makes a huge difference when it comes to fully-optimizing your virtualization environment. We start by examining x86 application characteristics (these characteristics usually dictate whether you should buy a tower, blade, or rack x86 platform design). We then turn our attention to x86 design points such as memory management, energy efficiency, reliability extensions, and the like — and illustrate how these design points make a big difference when operating a virtualized system environment. We conclude with our *Summary Observations* — a recap of why your x86 platform choice is crucial to building an optimized x86 virtualization environment.

Picking the Right x86 Server for Your Environment

Picking the right x86 server involves understanding application requirements as well as x86 systems designs. The remainder of this section examines both of these evaluation criteria.

Understanding x86 Application Designs

In a report entitled "*Buying Criteria for x86 Servers (Blades, Rack, Towers)*", *Clabby Analytics* observes that the first step in choosing a server should be dictated by application considerations. For instance, IT buyers need to consider:

1. Whether the application is 32-bit or 64-bit;
2. What the application's memory requirements are;
3. What the application "threading" requirements are;
4. Where applications will be located (central, branch office, Web-tier, …); and,
5. What quality of service (QoS) is required.

Why are these criteria important? Consider the following:

- *The 32-/64-bit decision* — the reason that the "bittedness" (32- or 64-bit) of an application is important has to do with the amount of addressable memory it can exploit. 64-bit applications can exploit "very large memory" (VLM) ─ enabling hundreds of gigabytes (or even terabytes) of data to be addressed in main memory where it is closer to the CPU. By being close to the CPU, data can thus be processed much more rapidly ─ sometimes even 100X faster).
- *Memory management* — some x86 designs isolate memory — so memory is difficult to share; other designs allow for large amounts of memory to be shared — and at higher speeds (memory bandwidth). There is great variability in how vendors allocate memory, so it is particularly important to pay close attention to your prospective vendor's memory management strategy.
- *Scalability/threading requirements* — applications that can exploit very large memory and that can exploit 4 processor or greater symmetrical multi-processing are best served by scale-up x86 architectures. Applications that process small threads and require little memory are best deployed in scale-out environments.

- *Location* — if an application and its data are located in a branch office, a tower or blade server may be the best platform choice (because the amount of traffic sent over a network can be reduced — so performance improves while network latency is reduced). By comparison, if the application is in at the Web-tier constantly caching images — and if this is being done on a massive scale — a massively scalable x86 Internet server design may be the best option.
- *Level-of-Service* — service level requirements differ by application. Office applications, for instance, are usually considered low priority — and hence, the level of RAS (reliability, availability, and security) required to support office applications is usually low. The same is generally true about page serving Web tier applications (if a server fails, then that server is simply failed-over to another server — and the broken server is then fixed at a later point). Mission-critical enterprise resource planning applications, on the other hand, require high degrees of RAS — and hence require highly-available, advanced scale-up x86 designs. It should be emphasized that when IT managers and administrators virtualize *ANY* server, the RAS feature set of the base platform becomes increasingly important.

Understanding x86 System Designs and Balance Points

Once an enterprise has a handle on application requirements, the next step should be to choose a server platform system design that can best meet application runtime and QoS requirements while addressing other data center considerations such as power usage, heat dissipation, footprint requirements, etc. From an x86 platform architectural perspective, you will be choosing between towers, blades, racks, and large scale integrated racks. Each of these architectures has different "balance points".

Differences can be found in:

- Reliability/availability characteristics;
- Performance characteristics;
- Scalability/capacity;
- Memory management and memory capacity;
- Virtualization handling (including virtualization off-loading); and in
- Power management and heat dissipation.

Ultimately, your goal should be to pick the x86 platform that best meets your application processing needs, your application service level requirements, and that addresses "other" data center requirements (such as reduced energy usage, floor space, etc.).

A Closer Look at x86 Server Design/Balance Points: IBM's Four x86 Architectures

Note: in the following section, Clabby Analytics has chosen to base its discussion of x86 system architectures on IBM x86 platforms. Our reasons are that IBM has the broadest x86 product set in the industry — as well as greater depth in the x86 server space than any of its competitors. For example, neither Hewlett-Packard (HP) nor Dell offer a fully integrated rack equivalent to IBM's iDataPlex; HP offers two blade chassis designs as opposed to IBM's five optimized blade chassis; IBM offers virtualization offload technology as part of its enterprise eX4 architecture to speed-up virtualization and reduce memory latency (its leading competitors do not have functionally equivalent capability); and IBM offers far better x86 memory management/reliability, greater scalability, as well as far greater memory capacity, when compared to any of its leading competitors. Further, IBM's System x and BladeCenter servers offer distinct differentiation in energy efficiency and virtualization.

An Overview of IBM's x86 Server Line

IBM offers four different x86 system designs — each tuned and optimized for different QoS, application and data center requirements (see Figure 1).

Figure 1 — Differing Design Points to Meet Varying Data Center Needs

	Infrastructure simplification	General Purpose Computing	Large Scale Computing	Server consolidation
Platform	IBM BladeCenter®	IBM System x™ two-processor server & tower	IBM System x iDataPlex two processor servers in a full rack	IBM System x four + processor server
Description	Integrating server, networking and storage resources	Combining a few applications on a single server for greater utilization	Integrated Rack solution for large scale virtualized workloads	Consolidating large numbers of underutilized servers for greatest TCO
Platform Key Attributes	• Hardware usability • Packaging density • Unified management • Power/cooling savings	• Industry-standard design • Price/performance • Compatibility	• Power/Cooling Efficiencies • Footprint Density • Optimized Integrated Rack • Serviceability	• Performance • Scalability • Strong reliability features

Source: IBM Corporation — October, 2009

A closer look at each of these architectures shows the following:

1. *IBM BladeCenter servers and chassis* — the focus of IBM's BladeCenter servers is on infrastructure simplification. Accordingly, IBM has focused on building an environment that tightly integrates server, storage, and network facilities. IBM BladeCenter blade environments are optimized around hardware usability, packaging density, unified management, and power/cooling savings. Additionally, the IBM BladeCenter platform has one of the only designs that offer dual I/O and power connections from every server in every chassis, a feature that is particularly important in a virtualized environment because it increases the reliability and availability of the underlying platform.
2. *IBM System x two processor rack and tower servers* — the focus of these server environments is on general purpose computing — allowing two or more applications to be combined on a single server to increase server utilization. The key attributes of this design are industry standard design, compatibility, and price/performance, and affordable RAS features.
3. *IBM's iDataPlex environment* — the focus of IBM's iDataPlex systems (a large scale, highly integrated rack environment) is high performance in large scale application environments (such as high-performance computing environments and/or Web 2.0 environments). This environment does not provide as much component redundancy as other IBM server environments — but, as a result, it draws significantly less energy that other equivalently-powered servers (from a processing capacity point-of-view), and it requires significantly less cooling. These servers are used in data center environments where power/cooling efficiency and footprint density are crucial; and where on-the-fly serviceability is important.
4. *IBM System x enterprise 4-16 processor rack servers* — IBM System x four-processor and greater servers have been designed and optimized for server

consolidation — enabling large numbers of under-utilized servers to be consolidated and virtualized within a single, scale-up environment. This architecture emphasizes performance, scalability, and strong reliability features such as Memory ProteXion. The memory and processor scalability features allow a broad set of applications to be virtualized without impacting application performance. When larger virtual machines are needed for databases and mission-critical applications, these enterprise servers are ideal.

x86 Design Extensions That Virtualization Software Can Exploit

As stated at the outset of this *Advisory*, virtualization places additional demands on hardware platforms in the areas of reliability, scalability, performance, energy use, and human-related management costs. This section examines each of these points more closely.

Reliability

Skimping on reliability features on x86-based servers has the potential to backfire on thousands of IT executives who are currently deploying virtualized servers (because virtualized servers that fail are typically running more applications or serving more people than non-virtualized servers).

IBM x86 reliability extensions include:

- *Hot-add/hot-swap I/O and hard disks;*
- *Hot-swap/redundant fans and power supplies;*
- *IBM Memory ProteXion* — a product that improves redundancy, making it possible to correct multiple single-chip errors and protect data integrity;
- *IBM Chipkill memory* — a product that helps correct multiple, single-bit errors, while enabling IT buyers to use less expensive, off-the-shelf DIMMs;
- *Memory mirroring* — simultaneous writes to independent redundant memory cards;
- *Advanced Buffer eXecution — ABX* helps resist chip failure and improve availability/reliability, increasing performance and reducing power consumption;
- *Hot-add/hot-swap memory* — should a DIMM fail, a memory card can be swapped out (or added in) without taking the server down;
- *Memory scrubbing* — an automated test of all system memory. Memory scrubbing helps detect memory errors that might be developing before they cause an outage;
- *Snoop Filtering* — reduces latency and improves performance*;*
- *Redundant I*/O *and power connectivity for blades* to reduce potential failures;
- *Predictive failure analysis;* and,
- *I/O virtualization.*

The reason that these features and innovations are important is that some vendors don't offer this depth and breadth of reliability in their designs. When you buy from a reputable vendor, these kinds of enterprise-level reliability features and function are included as part of standard server offerings.

Performance and Scalability

From a hardware perspective, IBM has used its deep microprocessor expertise to build a specialized controller that off-loads communications, memory management, I/O management, and virtualization tasks from x86 processors. This controller enables IBM x86 microprocessors in some IBM servers to focus on processing rather than virtualization

resource handling. This advanced virtualization/memory management controller is a feature of IBM's eX4 architecture — and its use leads to significantly improved x86 processor performance.

Processing scalability is also important in virtualization. Virtual machines have processor and memory requirements that must be accommodated in order to allow them to function properly. IBM System x servers support large scale virtualization, including scalability to 64 processor cores with VMware and Microsoft Windows Hyper-V R2, and to 96 processor cores with Linux delivered virtualization.

Energy Efficiency

IBM offers a wealth of energy management products that help monitor and manage energy use by servers and within a data center. One solution is IBM Systems Director Active Energy Manager (http://www-03.ibm.com/systems/management/director/plugins/actengmgr/index.html) Active Energy Manager helps customers quantify and control energy-related costs by

- Monitoring actual power consumed by servers;
- Trending power and thermal data;
- Setting energy and thermal thresholds which can trigger automated responses
- Monitoring power consumed by any equipment connected through select Power Distribution Units (PDU);
- Collecting information from select facility providers including Liebert SiteScan from Emerson Network Power and SynapSense;
- Allocating less power and cooling infrastructure to IBM Systems by setting caps on energy consumption of servers.

Human-related Management Costs

Clabby Analytics has long argued that the best approach to virtualization management is one that allows IT managers and administrators to easily manage both physical and virtual resources in an integrated manner (what you want to avoid is having one management environment for physical resources and another, disjointed environment for logical, virtualized resources).

IBM helps IT managers and administrators by providing an integrated hardware and software management environment. From a hardware perspective, all IBM System x and BladeCenter servers are designed to provide outstanding hardware manageability starting with on-board hardware for in-band and out-of-band (sometimes optional) control, predictive failure analysis and alerting. From a software perspective, IBM's System Director management environment can be used to provide virtual and physical server management in a single console with clear views of the relationship of resources.

A closer look at IBM's Systems Director 6.1 shows that it can be used to:

- Link to the management hardware on x86 servers and allow the customer to manage System x and BladeCenter servers, as well as other IBM (Power Systems and System z) and non-IBM (x86) servers from a single console. (Note: Systems Director is free for use with IBM servers);

- Support a variety of hypervisors from the leading providers for x86 virtualization — VMware, Microsoft, Red Hat and Novell. In addition, the same Systems Director interfaces support management of virtual machines running on IBM POWER Systems and System z mainframes;
- Link Systems Director to VMware vCenter and VirtualCenter. This gives System Director users the flexibility to leverage the power of VMware vSphere from the same Systems Director console from which they can manage physical servers. Users can also use the powerful event automation facilities of Systems Director to drive VMware functionality (e.g., Predictive Failure Analysis [PFA] may indicate that a host server is about to experience a hardware failure. Systems Director can respond automatically to the PFA and call vCenter to evacuate all virtual machines from that server so that it can be shut down — all automatically and without any downtime for the applications running on the virtual servers);
- Provide rich hardware information to Microsoft Windows Server 2008 with Hyper-V users. IBM provides a PRO Pack that pushes hardware information from IBM System x and BladeCenter servers up to Microsoft Systems Center Virtual Machine Manager (SC-VMM). This IBM hardware health information enables customers managing virtual environments to see the hardware health condition of physical hosts and provides advisory PRO tips to follow should error conditions occur; and,
- Provide physical and select virtual management for virtualization delivered with Red Hat and Novell Linux OS offerings.

Summary Observations

The bottom line in this *Advisory* is that, to optimize your virtualization environment, you need to pay close attention to the design characteristics and balance points of your x86 platforms. Virtualization software can be far more effective and efficient when deployed on underlying platforms that have been optimized for reliability, performance, scalability, and energy efficiency — and in environments where close attention has been paid to simplifying the management of physical and virtualized systems.

When evaluating server platforms, focus first on the processing requirements of your applications. Some x86 applications run best on scale-up hardware; others run best in scale-out environments. Pay particularly close attention to how your applications use memory, and to application quality of service requirements (as QoS requirements dictate the RAS characteristics required on your platform choice).

After evaluating your applications, evaluate platform design points, balance points and extensions that facilitate and accelerate virtualization integration and performance. Memory management, off-load controllers, reliability extensions and energy management facilities should feature prominently in your decision making process.

Finally, pay close attention to your vendor's management environment. The last thing you want is siloed virtualization infrastructure and management environments running in your data center. Look for vendors who integrate various virtualization products under one common umbrella physical and virtual system manager.

Clabby Analytics
http://www.clabbyanalytics.com
Telephone: 001 (207) 846-0498