CLABBY ANALYTICS

Advisory

Consolidation and Virtualization Are NOT Enough: Optimizing Application Performance in the New Enterprise Data Center

Executive Summary

Information technology (IT) executives consistently make two big mistakes when it comes to redesigning data centers to serve the next generation of service-oriented applications and business process flows. These mistakes are:

1. Many IT executives have embarked on strategies to consolidate and virtualize their information systems architectures (a good thing). But, after consolidating and virtualizing their systems and storage environments, *few of these IT executives are taking the time to properly tune and optimize their application workloads* to exploit their newly virtualized information systems resources. *And this failure results in continued waste of valuable processing power.*
2. Most IT executives fail to evaluate other systems platforms when consolidating and virtualizing information systems. (This is important because the virtualization feature/functions/management capabilities on RISC [reduced instruction set] and mainframe platforms are far superior to offerings on Itanium EPIC [explicitly parallel instruction computing] and x86 servers).

> ***To build highly-optimized data centers, systems need to be tuned to deliver the utmost performance. Further, to achieve maximum performance, applications need to be able to exploit advanced hardware-level virtualization, as well as deep software-level virtualization capabilities. Accordingly, to achieve maximum performance on virtualized systems, alternative-to-x86 and Itanium architectures should be evaluated.***

In this *Advisory*, *Clabby Analytics* (that’s me) describes some of the advantages that can be achieved by deploying and tuning applications for high-performance on *non-x86-based servers*. What my research shows is that alternative RISC and mainframe architectures can run many of the same enterprise applications far more efficiently than those currently being run on other architectures. Accordingly, *IT executives who are looking to maximize information systems performance for next generation data centers should strongly consider RISC and mainframe options.*

Background: The New Enterprise Data Center

Existing data center designs are rife with inefficiencies (inefficient power use, inefficient systems/storage utilization, inefficient business process flow designs, etc.). Myriads of distributed systems (servers, desktops, laptops, notebooks, mobile devices, etc) add to data center complexity — and distributed networks, with their plethora of entry points make data centers difficult to secure and manage. Adding to these woes, applications and data bases are inflexible when it comes to program-to-program communications and data sharing. To address these problems, a new model for data center design is emerging: the *new enterprise data center* (NEDC).

The NEDC design calls for the transformation of information systems, with a strong focus on improving overall systems/storage/network utilization performance while lowering management costs. From a program-to-program communications perspective, the NEDC model calls for applications to be written in a loosely-coupled, service-oriented form — rather than the rigid, hard-to-manage-and-maintain tightly-coupled form of the past. And NEDC calls for application "services" to be managed as contiguous business process flows that, when managed properly, help enterprises achieve the utmost in operational process flow efficiency. This new model promises to help significantly reduce operating costs (especially those related to systems management, security, and power consumption).

To reduce operating costs, information systems need to be better utilized than today's average 20% utilization rate (utilization rates are being increased by virtualizing unused resources) — and applications and databases need to be tuned to exploit more powerful, scaled-up/scaled-out dense server and storage designs.

Building Highly-optimized Applications for the New Enterprise Data Center

In order to improve data center efficiency, many enterprises have focused on consolidating existing information systems, and then virtualizing (logically pooling) unused resources on those systems to increase overall utilization while lowering management costs. And many have achieved startling results — increasing systems utilization to 60%, 80%, or higher — while also lowering systems/storage management costs by centralizing resource management. Additionally, power consumption is decreased as fewer servers that operate at high utilization rates do the work of tens or hundreds of distributed servers that had been operating at very low utilization rates.

But consolidation and virtualization in-and-of-themselves do little to improve application performance:

To maximize application performance in virtualized environments, applications and databases need to be shown how to exploit virtualized resources (how to efficiently use memory; how to efficiently use virtualized input/output [I/O]; how to acquire more processing power by sharing CPU resources [micropartitioning]; how to isolate and work around faults; etc.). And in this area (exploiting virtualization) major performance and manageability differences start to become evident when comparing x86-based architectures to RISC/EPIC/mainframe architectures.

The management of virtualized application environments is also very important. Automated workload management can also lead to performance improvements while greatly reducing management cost. Further efficiencies can be gained through automated workload balancing; heterogeneous device management; and by reliability, availability and security (RAS) features.

Why IT Executives Need to Focus on Applications Tuning and Optimization

Consolidation and virtualization deliver more computing resources to the IT organization, but failure to tune applications to run on virtualized resources means that *un-tuned applications are wasting processing cycles* — and, accordingly, virtualized servers are not being used to their full performance potential.

Failure to tune applications that are running on virtualized resources wastes precious processor cycles. Highly-optimized NEDC environments eschew processor cycle waste.

An IBM Customer Example: Energen

One IBM customer, Energen (an independent oil and gas exploration and production company) provides a great example of the kind of return-on-investment that can be recouped by taking the extra step of tuning and managing applications and databases — as well as by paying close attention to license management.

Many enterprises use multiple systems architectures to process varied workloads — from infrastructure on x86 Windows to database and applications on Unix — and Energen is no exception. Energen recently underwent a major server consolidation and virtualization trans-formation, moving from 20 physical Sun Unix servers to two large IBM Power Systems (formerly System p servers). Further, Energen consolidated and virtualized its x86 environment, moving from 60 physical servers to just 4 servers By consolidating and virtualizing these environments, Energen has been able to reduce operating costs (Energen has been able to save $500,000 per annum on its Unix systems and $600,000 per annum on its Windows [x86] environment).

A closer look at these cost savings finds that Energen focused strongly on license management and application tuning to exploit underlying virtualized resources. A big contributor to these cost savings came from reductions in Oracle licensing fees (the company eliminated redundant license costs and then tightened and tuned its database environment). But it is also noteworthy that Energen took the time to tune its applications to exploit the micro-partitioning feature of IBM's POWER processor (micropartitioning allows as little as one-tenth of a single processor to be exploited — keeping all available cycles busy). Further, Energen also takes advantage of the Virtual I/O Server feature of IBM Advanced POWER Virtualization on System p.

> ***As a result of these practices, Energen has been able to boost its systems performance by 92 percent, and cut its batch runs from 24 hours to 2 hours. This Energen example illustrates how taking extra steps to tune applications — as well as monitoring application/database license usage — can provide extra-large paybacks when consolidating and virtualizing systems environments. Failure to tune application environments can leave easy-to-recoup money on the table — something few IT departments can really afford to do.***

Evaluate RISC and EPIC Virtualization Environments As x86 Alternatives

RISC architectures (such as IBM's POWER and Sun's UltraSPARC) have been around for twenty-five years. EPIC architecture (such as HP's Itanium microprocessor) has been around for almost a decade. Each of these architectures was designed from their very beginnings to accommodate virtualization at the hardware level. X86 architecture, however, was not designed to provide hardware-level virtualization services.

> ***Executing instructions at the hardware level can greatly improve performance. It is logical, therefore, to expect that microprocessors that can virtualize instructions in hardware will perform better than those that do not.***

When comparing x86-based virtualization solutions to RISC-, EPIC, and mainframe-based virtualization, several architectural/virtualization differences become readily apparent. This section describes some of these differences.

RISC and EPIC: Comparative Advantages Over x86

RISC/mainframe/EPIC architectures offer a variety of performance-enhancing benefits in both hardware and software when compared to x86 architectures. For instance:

- At the hardware level, some RISC processors offer "micro-partitioning" (for instance, IBM's POWER6 microprocessor can be micro-partitioned into ten segments, each capable of driving an independent workload; Sun offers like functionality in "zones"). The net result is that all processor cycles can be used all of the time — ensuring maximum resource utilization). X86-based processors (and EPIC processors) do not offer micro-partitioning;
- From a hardware and software perspective, RISC and EPIC microprocessors support virtual memory management and virtual input/output (I/O). Advanced memory features ensure that applications can acquire the amount of memory needed to execute tasks most expeditiously. And virtual I/0 provides applications with access to virtualized I/O network resources — ensuring that applications can gain access to the right amount of bandwidth needed to most efficiently acquire and/or share data.
- RISC architectures offer software that allows workloads to be moved transparently between processors or servers (mobile partitioning), allowing applications to exploit unused computing resources regardless of whether those resources are located on local or remote servers. (Note: Itanium EPIC architecture with HP/UX does not provide partition mobility at present).In addition, IT administrators can move workloads off of partially used servers to other servers — enabling administrators to shut-down partially used servers. This, in turn, creates less wear-and-tear on existing systems resources while also cutting energy consumption costs; and,
- IBM's AIX (Unix) and z/OS (mainframe) operating environments — as well as Sun's Solaris and HP's HP/UX (Unix) operating environments all offer advanced automated workload management and provisioning facilities that help balance application access to underlying resources. All of these offerings are superior to what can be found in Unix/Windows/Linux environments on x86 processors.

All of these advanced virtualization features lead to improved performance advantages when compared to x86-based architecture.

EPIC Architecture: Not as Strong in Hardware/Software Virtualization as RISC

For ten reasons, *Clabby Analytics* is not a fan of Intel's Itanium architecture (see the video at http://www.accelacommunications.com/programs_ibm_sol.htm for a deeper discussion of why *Clabby Analytics* is somewhat less than enthralled with Itanium).

Three of the top reasons that my firm is "less-than-excited" about EPIC and Itanium include: 1) this processor has consistently been late to market with product releases; 2) when Itanium processors are delivered, functionality is sometimes dropped to meet release deadlines; and, 3) this processor has, at times, been years behind competing RISC processors in terms of functionality (meaning that RISC buyers have had the ability to exploit advanced microprocessor functionality sometimes years longer than their Itanium brethren).

To its credit, however, Itanium has finally started to catch-up with RISC chips in terms offering the ability to provide virtualization services at the hardware level when Montecito was released a few years ago (though micro-partitioning and other advanced functions are still not available). Further, Itanium does support Microsoft Windows as well as Linux (RISC does not support Windows natively). But, it should also be noted that Microsoft

(Virtual Server) and VMware (the market's leading x86-based virtualization vendor) have both withdrawn support for Itanium virtualization. Further, Windows Server 2008 for Itanium-based systems does not support Microsoft's newly released Hyper-V virtualization scheme (Hyper-V runs on x64-based chips only).

X86 Architecture

As with Itanium, x86/x64 hybrid architectures can run Windows and Linux. But when x86 architecture was designed, it was not designed to be virtualized at the hardware level (where virtualization services can run the fastest — thereby positively impacting the amount of performance that can be achieved from a given microprocessor).

To overcome this problem, x86 desktop and server vendors have fashioned two basic approaches:

1. Pre-processing virtualization services on an ASIC (application specific integrated circuits) such as those provided by IBM with its x4 architecture and Hitachi with its Virtage. Pre-processed, prepared and managed virtualized workloads are then handed off to x86 hybrids for expedited execution; and,
2. Solving virtualization at the software level by building x86 hypervisors and related infrastructure/management services. (A hypervisor is a software layer that sits between an operating environment and underlying hardware. Hypervisors enable several different operating environments or instances to use an underlying processor). Vendors such as EMC with VMware, Microsoft with Hyper-V, and Citrix with XenServer — as well as dozens of other companies — offer x86 hypervisors and related infrastructure/management solutions.

In addition to these two approaches, it should also be noted that some vendors now offer, or are planning to offer pre-loaded, embedded virtualization infrastructure that resides on solid state disk drives located close to an x86 processor. By doing this, virtualization services can be pre-configured and loaded quickly — enabling rapid scalability (because multiple virtual engines can be rapidly deployed), and improved performance (as applications can then exploit these new, rapidly-deployed engines). IBM, as part of an agreement with its close x86 virtualization partner (VMware), is the first to provide this type of solid-state-based, embedded VMware infrastructure service solution to the market (part of IBM's x4 architectural solution).

There are various approaches to improving virtualized application performance on x86-based architecture. The question IT buyers need to ask themselves is can RISC-based platforms provide faster virtualized application performance, richer virtualization services, and better virtualization management capabilities. For Linux and Unix environments, Clabby Analytics believes that the answer is that RISC native hardware virtualization services and deep software virtualization services are far superior to what is offered today in the x86-based virtualization.

Decision Making Criteria

IT executives who want to deploy virtualized servers — and maximize application and database performance on those servers — will find differences in platform vendor offerings by closely examining:

1. Virtualization capabilities at the hardware level;
2. The level of sophistication of virtualization software offerings; and in
3. Virtualization management capabilities.

Hardware Differences at the Microprocessor Level

The primary difference between RISC, EPIC, and x86 processors is that RISC processors can be virtualized at the micro-partition level (many processes can run at the same time on a given RISC processors, thus not wasting CPU cycles).

> ***The key point in micro-partitioning is not to waste CPU cycles. Both the Sun and IBM solutions provide a means to partition processors to run numerous workloads at the sub-processor level — enabling IT managers to use CPU cycles that go unused without hardware micro-partitioning.***

Virtualization Software Differences

When examining virtualization software, differences can be found in how many virtual machines can be supported (scalability); related performance; operational flexibility; and security/integrity. Pay particularly close attention to differences in product depth and sophistication in the following areas:

- Memory management and utilization;
- CPU sharing;
- Resource over-commitment support; and, in
- Security controls/fault isolation.

Virtualization Management Capabilities

Comparative/competitive differences between x86-based virtualization solutions and RISC become even more pronounced when examining virtualization management solutions. The primary differences can be found in the ability of virtualization software vendors to support:

- Virtual machine mobility (the ability to move live sessions transparently between virtual machines that may actually reside on different physical servers);
- Heterogeneous management (cross platform systems management at both the physical and logical [virtual] level; and especially in the area of heterogeneous storage management);
- Command and control capabilities;
- Monitoring; and,
- The level and extent of automated management functions.

> ***The bottom line in this section is that very significant differences exist in hardware virtualization, software virtualization, and virtualization management between x86 architecture and RISC/EPIC alternatives. These differences have huge implications on the number of virtual machines a given architecture can support, the kind of performance that can be delivered, the cost of each virtual machine, the level of availability and security that can be expected, and the ease of management that can be expected. All of these reasons — and more — are why IT executives should consider x86 alternative architectures when transforming existing systems and storage architectures to fit the new enterprise data center model!***

Vendor Commentary

The top tier IT systems makers are IBM, Hewlett-Packard (HP), Sun, and Dell. Each of these systems makers sells Intel/AMD x86 solutions. IBM and Sun sell scale-up RISC solutions. HP sells Itanium EPIC solutions. And Dell sells neither EPIC nor RISC solutions (hence, Dell is not included in the following discussion).

Hewlett-Packard: General Impressions

As described earlier (in the "EPIC Architecture: Not as Strong in Hardware/Software Virtualization as RISC" subsection), HP's Itanium-based line has several limitations when compared to RISC processor-based systems. On the x86 side of the house, HP can argue that it has some advantages in x86 management — but, for the most part, HP and IBM are on a par when it comes to Intel/AMD management applications, tools and utilities.

Sun: General Impressions

Sun's virtualization hardware, software, and management products are competitive with IBM AIX/POWER platforms. And Sun's forthcoming ROCK processor looks to be especially interesting from a scalability perspective. Unfortunately, I could find little evidence that Sun's virtualization software and management products have been widely adopted. Sun buyers, it appears, have opted for 3rd party software solutions (OpsWare, Symantec/-Veritas, et al) rather than widely adopting Sun virtualization and provisioning solutions.

IBM: General Impressions

IBM's POWER RISC processor-based servers have several distinct competitive advantages in the areas of hardware, virtualization infrastructure, and virtualization management. These advantages can be readily exploited to tune applications for better performance than other RISC, EPIC, and x86 solutions. Further, IBM's x86-based product lines feature outstanding virtualization acceleration (part of IBM's x4 architecture), as well as numerous advanced features in terms of x86-based virtualization management. IBM's close partnership with EMC's VMware helps complete IBM's System x virtualization environment by providing a well received x86 hypervisor and solid x86 virtual machine infrastructure and management software.

A closer look at IBM virtualization offerings shows distinct differentiation in:

- *Systems hardware* — IBM is always amongst the top players (if not the category leader) in almost every industry system benchmark comparison. Note that:
    - IBM's front-end x4 architecture helps accelerate high-end x86 server performance, enabling these servers to scale up to 64 cores with less latency and with a higher speed pipe for I/O than other all other leading x86 commercial server offerings. And, further,
    - IBM's POWER6 processor can run at 5 Ghz, providing screamingly high application performance.
- *Applications Tuning* — IBM POWER servers are on very strong footing when it comes to applications tuning for virtualized environments. At the hardware level, POWER can be micro-partitioned. At the virtualization software level, POWER stands out in memory management and CPU sharing — as well as in the areas of security controls and fault isolation.
- *Virtualization Management* — One of the ways to increase applications performance is to move applications closer to the resources that they need. IBM application mobility (a feature of IBM's AIX operating environment) enables seamless application portability across partitions within the same server to ensure availability. Live partition mobility, a Power VM (POWER processor) feature enables the portability of running (executing) partitions. Applications running in those partitions across servers *without disruption*. This provides systems administrators

with more flexible systems administration, enables faster response to changing workload conditions, and enables better resource utilization.

- *Heterogeneous Management* — IBM can run and manage multiple operating systems and associated workloads within a single platform environment (for example: IBM x86 servers can run both Windows and Linux applications; different applications can run on Linux, IBM i or AIX operating environments on the same POWER server; and IBM SAN Volume Controller (SVC) can provide for a single point of administration for multiple storage systems, even in multi-vendor environments, allowing for live data migration without application disruption, cross platform replication, storage tiering and simplified management.*Server Design Flexibility*— Customer can implement applicationservers as clusters or in SMP (symmetrical multiprocessing) configurations. IBM's System x3950, as well as its POWER 570, have 4 core building blocks that can easily facilitate cluster or in a pure SMP (NUMA) design solutions.

Summary Observations

Two of the hottest trends in the information systems industry are consolidation and virtualization. IT buyers are consolidating server/storage/networks in order to lower operational costs as well as to reduce management costs through centralized management. And IT managers are also "virtualizing" their computing environments (creating logical views of unutilized systems and storage resources) in order to more easily find and exploit unused computing resources. Each of these practices yields almost immediate return-on-investment by enabling IT executives to get more utilization out of existing information systems investments – while helping to reduce data center management and operational costs.

> ***But what good are consolidated, virtualized environments if applications are not tuned to exploit them?** IT managers may be able to achieve huge increases (in the range of 60-80%) in terms of system utilization by consolidating and virtualizing servers – but if applications run poorly in those environments, computing and storage resources will continue to be wasted. Powerful, fast, highly-optimized systems architectures that run slow, untuned applications simply cannot and will not perform to their highest potential. To maximize ROI, applications need to be tuned to exploit consolidated, virtualized resources.*

IT managers can continue to virtualize and scale-out their computing environments by buying low cost, x86-based hardware – and then by spending vast amounts of IT budget correcting resulting design and performance deficiencies. Or IT buyers can spend a little more at acquisition time to buy highly-optimized, scale-up SMP or cluster designs – and save tons of money in the long run through improved application performance and lower management/operational costs. Each approach works – but the former is highly-inefficient, wastes IT resources, is more difficult to manage, and is more difficult to tune to maximize application and database performance in virtualized environments.

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

Clabby Analytics is an independent technology research and analysis organization that specializes in information infrastructure and business process integration/management. Other research and analysis conducted by Clabby Analytics can be found at: www.valleyviewventures.com.