Harmonizing the Twin Trends of Open Source and Virtualization: How Kernel Based Virtual Machine (KVM) Drives Enterprise Business Value

Sponsored by IBM
Srini Chari, Ph.D., MBA
November, 2010
mailto:chari@cabotpartners.com

Cabot Partners Group, Inc. 100 Woodcrest Lane, Danbury CT 06810, www.cabotpartners.com

Cabot Partners
Optimizing Business Value

Executive Summary

In the 1960s, computers housed in data centers were few in number and this resource was very expensive. As a result utilization of those compute resources was kept very high. In order to optimize the utilization of those resources for different business purposes and to facilitate the ever increasing number of users, IBM invented and engineered virtualization and built these capabilities into the very essence of mainframes, and has continued to deliver innovative mainframe virtualization solutions for decades. More recently, some of these capabilities were extended to IBM's UNIX Systems (Power Systems) using PowerVM for virtualization of an entire system or within Logical Partitions (LPARs).

Today's enterprise data center crisis is largely caused by the sprawl of under-utilized x86 systems, ever escalating electricity costs, and increasing staffing costs. Using virtualization to centralize and consolidate IT workloads, many organizations have significantly reduced their IT capital costs, reduced operational expenses, improved IT infrastructure availability, and achieved better performance and utilization.

While proprietary x86-based virtualization solutions have been very successful, open source virtualization solutions such as KVM (Kernel-based Virtual Machine) have delivered significant enterprise grade benefits beyond current x86 proprietary solutions. These include the best performance on industry standard benchmarks, unmatched security, powerful memory management, best in class reliability, and the broadest support for hardware devices including storage. Further, since KVM is part of Linux, clients can benefit from the numerous advantages of Linux including lower TCO, more versatility, and support for the widest range of architectures and hardware devices. Moreover, Linux performs, scales, is modular and energy-efficient, is easy-to-manage, and supports an extensive and growing ecosystem of ISV applications.

IBM System x is built on a philosophy of working in close collaboration with many partners to jointly deliver the best solutions that are fully optimized for current and emerging workloads. IBM uses the X-Architecture approach to continually innovate and optimize systems design to improve the performance, reliability, energy efficiency, and the TCO of the System x family. The recent eX5 and M3 servers in the IBM System x family, for example, comprehensively outperform the competition. IBM works with leading distributors of Linux to help make Linux more scalable and robust for the enterprise computing environment. IBM also provides migration services that make the switch over to Linux and KVM straightforward. Here we showcase the IBM System x-Linux-KVM environment that has been extensively used in real world customer environments to run enterprise workloads. This degree of collaboration and innovation on all fronts -- semiconductor technology, hardware, software, middleware and operating systems, and enterprise workloads -- ensures that the IBM System x-Linux-KVM platform has a clear roadmap and has ongoing support. For customers this means that their investment is protected against changing market conditions.

Introduction - Smart Management of Enterprise IT Assets is Critical

Over the past decade, the cumulative effects of Internet technologies, the rise of open source software, increasingly economical computing performance, and pervasive high-speed digital communications have spawned entirely new businesses, agile organizational structures, and new ways for enterprises to maximize stakeholder value. Enterprise IT has become a strategic asset enabling businesses to innovate with flexibility and speed in response to ever changing customer demand, market opportunity, regulatory changes, or a competitor's move.

However, the current economic downturn compounded by escalating energy and staffing costs are forcing these companies to reevaluate how they can maximize their return on their IT investments. They need smarter approaches to reduce costs, manage complexity, improve productivity, reduce time to market, and enable innovation. Simply put, companies must carefully examine the business (value and costs) case of IT investments. They must maximize their Return on Enterprise IT Assets.

But the ways Enterprise Information Technology Assets are deployed are changing too. New developments in x86-based energy-efficient, computationally dense servers, data center facilities, open source software, virtualization, service-oriented architectures, and cloud computing reallocate IT costs and alter usage patterns in unprecedented ways. This rapidly shifting technology environment raises serious questions for IT managers and business executives about how to help their organizations capitalize on these transformative trends and make the corresponding strategic technology choices.

The x86 architecture continues to deliver improved qualities of service for enterprise computing; increasingly supporting advanced mission-critical features. Likewise, Linux and subsequently open source software and web solutions adoption grew rapidly as customers benefited from broad access to attractively priced solutions that worked on multi-vendor hardware spanning multiple architectures while preserving traditional rich Unix-like features of ease-of-configurability, management, reliability, and resilience.

More recently, over the past five years, proprietary x86 virtualization solutions have gained traction in the marketplace; gradually including features from more established and proven mainframe and high end proprietary UNIX virtualization solutions such as PowerVM. And just like Linux did with operating systems, open source virtualization solutions such as Kernel-based Virtual Machine (KVM) promise significant benefits well beyond today's proprietary x86 virtualization solutions.

KVM is a full virtualization solution for Linux on x86 hardware containing virtualization extensions. KVM enables users to run multiple virtual machines running unmodified Linux or Windows images -- in effect, KVM transforms the Linux kernel into a lean, efficient and secure hypervisor. IBM and other providers of open source solutions are working to help x86 customers reap the benefits of advanced data center virtualization while enjoying the economies and flexibility of customization offered by open source solutions. Many cloud service providers have implemented highly scalable and elastic cloud infrastructures using the latest IBM System x + KVM.

Through in-depth research and interviews, this paper describes how KVM as the critical technology component for virtualization on IBM System x servers deliver unique value for enterprise clients. This harmonized open source solution has been deployed to:

- ***power a large production cloud computing service*** for a global telecommunications company
- ***reduce cost, increase performance and scale growth*** for an offshore oil and gas company
- ***transform IT infrastructure through scalability and cost savings*** for a telecommunications firm
- ***launch a scalable development and test cloud*** at IBM
- ***power and maintain a continuous mission-critical ERP solution suite*** for a pharmaceutical firm.

Shifting Economics of Enterprise Computing and Data Center Challenges

Over the last several years, we have interviewed several IT executives from large and small organizations at various phases of their enterprise data center life cycle. Some data centers have been operating for well over a decade, while others are earlier in their life cycle; operating for just a few years. While some challenges are unique, many business and technical challenges are common across many of these organizations regardless of size or maturity. These challenges along with some current enterprise computing trends are:

The rise of x86 systems for enterprise computing: The continuing growth in the performance of x86 microprocessors, advances in parallel multi-core processor architectures, high-speed interconnects, coupled with innovations in open source software efforts (e.g. Linux, Internet, the w3, etc.) have enabled users in many companies to invest in building computing capability by clustering local resources that were largely in their own control; often without "buy-in" from corporate IT. Even the hardware manufacturers began shipping low-priced, affordable "cluster in a box" and blade systems that were very attractive to many business and end-user workloads. However, over time, many enterprises realized the problems with ever increasing server sprawl and the associated increase in the total cost of computing, and began consolidate their computing infrastructure.

Escalating computing costs are driving IT consolidation and centralization: With the current economic downturn, many enterprise computing decisions are increasingly being centralized driven by equipment procurement costs, budgetary pressures, and administrative issues prevalent in distributed models. Decision makers are taking a hard look at the total cost of ownership (TCO) for enterprise computing along with the total value of ownership[1]. Direct software, support, maintenance costs, and other escalating operational costs such as energy, power and cooling, people, and facilities are expected to continue to rise in the future. The following figure qualitatively depicts the changing costs of research computing over the last four decades. While the IT hardware costs as measured by $/performance have come down significantly with the adoption of new x86 based hardware technologies, the associated software and operational costs continue to rise, and in fact, today, these costs dominate the TCO.

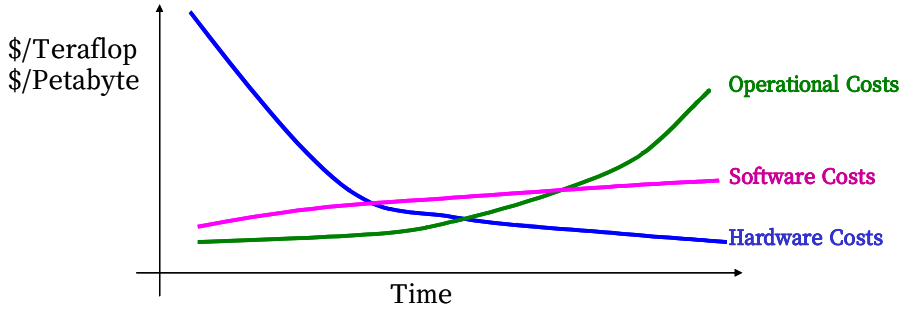

Figure 1: The Shifting Costs of Enterprise Computing Motivate IT Centralization and Consolidation

Focus on green data centers driven by corporate sustainability mandates and economics: According to IDC[2], the electricity costs in data centers to power and cool hardware are expected to increase by 11.2 percent while new server spend is expected to remain almost flat. The issue of power and cooling has become a top priority for IT executives. The IT industry is defining additional metrics such as gigaflops/watt that rate systems. The Uptime Institute has also defined additional metrics[3] that define data center "greenness". The Green500 list[4] is becoming as important as the Top500[5] list of supercomputers as vendors compete for bragging rights. Bringing sound environmental and management principles to bear in capacity planning and operating a data center can become a competitive advantage and a source of operational efficiency and increased reliability for many information-intensive industries. It's estimated that more than

[1] Dempsey, Jed, Robert E. Dvorak, Endre Holen, David Mark, and William F. Meehan III, "A Hard and Soft Look at IT Investments." McKinsey Quarterly 1: 126-137, 1998.
[2] Jed Scaramella, "Worldwide Server Power and Cooling Expense 2006-2010 Forecast", September 2006.
[3] John R. Stanley, Kenneth Brill, and Jonathan Koomey, "Four Metrics Define Data Center "Greenness"", White Paper, Uptime Institute.
[4] The Green 500 List, www.green500.org/Lists.html
[5] The Top 500 list at www.top500.org

70% of the world's Global 1000 organizations will have to modify their data center facilities significantly during the next five years.[6]

The growth of enterprise grade Linux on x86 systems: Linux and subsequently open source software and web solutions adoption grew rapidly as customers benefited from broad access to attractively priced solutions that worked on multi-vendor hardware spanning multiple architectures while preserving traditional rich Unix-like features of ease-of-configurability, management, reliability, and resilience. Moreover, customers were able to rapidly deploy, customize, and manage their IT environments. Today, Linux is ubiquitous and is expected to grow faster than alternative operating systems[7]. Linux supports the widest range of architectures and hardware devices and it performs, scales, is modular and energy-efficient, is easy-to-manage, and supports an extensive, growing ecosystem of ISV applications[8].

Enterprise customers demanding mission-critical environments have further benefited from the advent of enterprise class Linux distributions on enterprise grade hardware including the IBM mainframe, Power, and System x. Today, the current generation of x86 CPUs includes a few enterprise grade features that were once only available on mainframes and high end UNIX systems. These features include large-scale multithreading with 64 or more virtual CPUs, support for large memory systems, and chipsets with advances reliability, availability, and serviceability (RAS). These enterprise-class features are today supported by Linux on IBM mainframes, Power systems, and System x.

Addressing the data center crisis; motivating the need for centralization and virtualization: Data centers typically account for 25 percent of total corporate IT budgets when facilities, storage devices, servers, and staffing are included. The costs of operating a data center facility is growing by about 20 percent a year, far outpacing overall IT spending[9]. Further, the investment required to launch a large-enterprise data center has risen to $500M, from $150M, over the past 5 years and larger data centers take 2 years or more to design and build and are expected to last for 12-15 years. Yet, server utilization typically tops out at 5 to 15 percent[9], wasting both energy and employed capital. Many IT managers indicate that excess servers exist to provide for extreme situations e.g. holiday seasons. However, this assertion may not be completely true.

The mismatch could be because many companies have difficulties in accurately forecasting workload demands. The long-term (12 years or more) lifecycle of a data center investment coupled with the lack of a holistic and integrated view of future workload needs make enterprise-wide capacity planning a daunting task. Many companies would have difficulty forecasting whether a 50% increase in demand would require 20% or 80% more server and data center capacity. In the extreme, companies that rely entirely on in-house computing capacity may be stuck with excess wasted space in a datacenter, or may have to undertake the large expense of building a new data center almost immediately.

What's needed to address this data center crisis is a comprehensive approach with a centralized governance model in which the CIO is empowered by the CEO and is accountable and responsible for data center management. Enterprise CIOs can reform data center operations to further increase enterprise value by:

- managing IT assets aggressively through virtualization, consolidation, and standardization
- implementing new metrics for data center efficiencies that account for energy, utilization, and floor space
- improving forecasting and minimizing deviations from real demand through better capacity planning and management processes
- treating all IT resources as scarce resources or "services" that can be quickly scaled up or down
- ensuring that business units implement a TCO or ROI model for new systems and applications.

Implementing these recommendations with a harmonized open source virtualization solution and complementing this with optional cloud services can further enhance enterprise business value.

---

[6] https://www-927.ibm.com/servers/eserver/storageplaza/BERT.nsf/files/India2008CSIIPresentations1/$File/MP01%20-%20Supporting%20Business%20Innovation%20with%20NEDC-Debra%20Thompson.pdf

[7] Al Gillen, "The Opportunity for Linux in a New Economy", April, 2009, http://www.linuxfoundation.org/sites/main/files/publications/Linux_in_New_Economy.pdf

[8] Amanda McPherson, "Linux: The Operating System of the Cloud", http://www.linuxfoundation.org/sites/main/files/publications/linuxincloud.html

[9] McKinsey on Business Technology, Innovations in IT Management, Number 14, Winter 2008.

The Benefits of Virtualization

Virtualization decouples operating systems and their applications from server hardware. This is done through the hypervisor, a layer of software between the server hardware and the operating system. The hypervisor converts a single server into multiple independent virtual servers. There are many reasons why virtualization is hugely beneficial to companies:

Cost savings: Consolidating servers using virtualization not only provides savings in terms of how many servers must be bought and maintained, it also reduces the amount of physical space that a company needs for its servers or data center.

Application virtualization: Since more than one operating system can now be run on a single server, applications are no longer tied up to one server. Applications can be easily ported from one server to another or shared widely through a cloud.

Virtualization support: Companies can easily and economically outsource all or most of their computing requirements to third party providers who can now run systems for several different companies on a single server while keeping the applications securely separated from each other.

Product testing: The testing part of a computing project is often a high-overhead effort that consumes relatively expensive resources and a lot of time. Virtualization tools allow software testers to quickly and easily set up and maintain testing environments and to rapidly restore testing environments to their production state when done.

Application deployment: Computing projects are often resource intensive in terms of deployment. Virtualization allows for the easy creation of base computing environments and largely eliminates hardware compatibility issues. This can save costs and time in rolling out new or changed computing systems and applications.

x86 Virtualization Approaches

As enterprise computing demands continuously escalate with ever higher demands on servers and datacenters – and consequent runaway hardware and energy costs – virtualization has become critical in optimizing x86 server investments.

There are different approaches to x86 server virtualization and these have evolved with the x86 architecture. In conventional architecture, a single operating system manages the hardware and applications sit on top of the operating system (see Fig. 2).

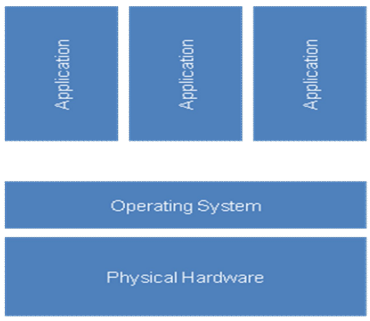

Fig 2: Conventional architecture

The operating system has full access to the hardware. Indeed, operating systems are designed to have full access to the hardware. Applications have no direct access to the hardware. In a virtualized environment however, it is the hypervisor alone that needs to have full access to the hardware to control hardware and system functions. In contrast, the operating system must be isolated from the hardware and must only communicate with the hypervisor (see Fig. 3).

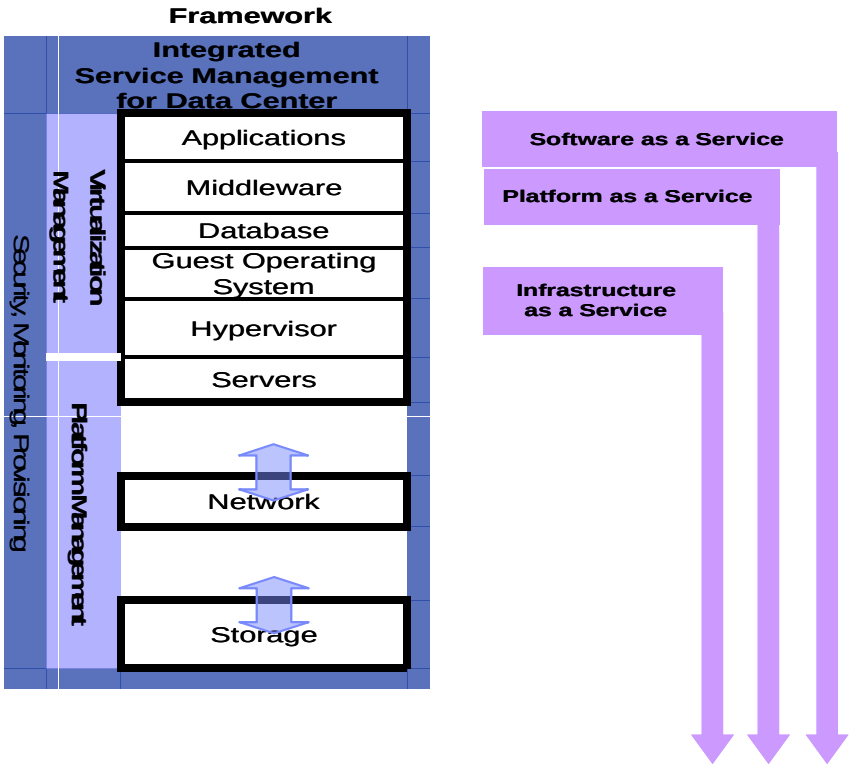

Fig 3: Virtualization architecture

In a virtualized architecture, each 'virtual machine' has its own 'guest' operating system and emulated hardware that runs on top of the hypervisor layer that sits on the 'host' system. The challenge in x86 virtualization is centered on creating a manageable configuration that takes away the privileges of the operating system to communicate directly with the host and yet retain consistency and efficiency. There are four ways to do this.

- Binary translation
- Paravirtualization
- Hardware assisted virtualization
- KVM

Binary translation: Instead of emulating the processor, the virtual machine runs directly on the CPU. To trap and manage privileged instructions, the hypervisor scans the virtual machine memory and intercepts privileged calls before they are executed and then dynamically rewrites the code in memory. The operating system kernel is unaware of the change and operates normally. This combination of trap-and-execute and binary translation allows any x86 operating system to run on the hypervisor. This approach was pioneered by VMware.

Paravirtualization: Paravirtualization is a technique that presents a software interface to virtual machines that is similar, but not identical to, that of the underlying hardware. Instead of handling a privileged instruction directly (as in Binary Translation) the approach in paravirtualization is to modify the guest operating system running in the virtual machine and replace all the privileged instructions with direct calls to the hypervisor. This approach requires modification of the guest operating system. Linux kernel 2.6.23 and above have been modified and made suitable for paravirtualization.

Hardware assisted virtualization: Hardware-assisted virtualization enables efficient full virtualization using hardware capabilities primarily from the host processors. Full virtualization is used to simulate a complete hardware environment, or virtual machine, in which an unmodified guest operating system (using the same instruction set as the host machine) executes in complete isolation.

Hardware-assisted virtualization has been recently (in 2006) added to the Intel VT-X and AMD-V x86 processors. A new operating mode is added to the CPU which can now operate in host mode or guest mode. A hypervisor can request that a process operates in guest mode, in which it will see still the conventional architecture (that is, operating system communication directly with the host) but the CPU is instructed to trap privileged instructions and then return control to the hypervisor. Using these new hardware features, a hypervisor does not need to implement the binary translation that was previously required to virtualize privileged instructions.

KVM - Architected to Maximize Enterprise Business Value

Kernel-based virtual machine (KVM)[10] is the most recent step in the evolution of open source x86 virtualization. It builds on the insights gained from previous approaches and leverages present day hardware capabilities, the broadest open-source ecosystem, and all of the outstanding features of Linux. *Further, KVM provides the best performance on industry standard benchmarks, unmatched security, powerful memory management, best in class reliability, and the broadest support for hardware devices including storage:*

Broadest open source community support and robust implementation: KVM is implemented as a loadable kernel module that converts the Linux kernel into a bare metal hypervisor. Two key design principles have helped KVM to mature rapidly into becoming a leading high performance open hypervisor. One, it takes as a given the hardware modifications at the x86 processor level featured in Intel VT-x and AMD-V. Two, it takes full advantage of several Linux features such as memory manager, a process scheduler, an I/O stack, device drivers, a security manager, a network stack and so on. This means that KVM does not attempt to build everything ground-up and presumes the x86-Linux environment that is now a standard in enterprise computing. It does not waste any resources in accommodating other 'legacy' environments. In addition to the broad Linux community, KVM is supported by some of the leading vendors in the software industry including Red Hat, AMD, HP, IBM, Intel, Novell, Siemens, SGI and others.

Leverages all outstanding Linux capabilities: In the KVM architecture the virtual machine is implemented as regular Linux process, scheduled by the standard Linux scheduler. In fact each virtual CPU appears as a regular Linux process. This allows KVM to benefit from all the features of the Linux kernel including supporting all devices that Linux supports and Linux power management. Device emulation is handled by a modified version of QEMU that provides emulated BIOS, PCI bus, USB bus and a standard set of devices such as IDE and SCSI disk controllers, network cards, etc.

Best performance and scalability for industry standard benchmarks: KVM is designed to perform and scale very well for server consolidation and leverages all of Linux scalability and performance optimizations. These include performance optimizations for the latest processor, continued multi-processor scalability optimizations, large memory performance, NUMA topology support, and SR-IOV[11] support. In fact, RHEL 6 has greater scalability limits than ESX 4.1, and support for more devices, including SR-IOV.

Several key criteria to model and measure typical enterprise server consolidation workload performance and scalability include the ability to: run lots of virtual machines (VMs) of different sizes running many server types, monitor response times, handle variability in each of the VM's workload with substantial I/O, and reproduce the results consistently.

[10] KVM – KERNEL BASED VIRTUAL MACHINE, http://www.redhat.com/f/pdf/rhev/DOC-KVM.pdf
[11] http://www.pcisig.com/specifications/iov/single_root/

SPECvirt_sc2010[12] is SPEC's first industry-standard benchmark addressing performance evaluation of datacenter servers used in virtualized server consolidation. SPECvirt_sc2010 measures the end-to-end performance of all system components including the hardware, virtualization platform, and the virtualized guest operating system and application software. The benchmark supports hardware virtualization, operating system virtualization, and hardware partitioning schemes. Currently, KVM is the only virtualization platform for any reported results. These demonstrate the best performance and scalability.

Unmatched security[13]: Because a virtual machine is implemented as a Linux process, it leverages the standard Linux security model to provide isolation and resource controls. The Linux kernel includes SELinux (Security-Enhanced Linux) to add Mandatory Access Control. SELinux policies loaded into the kernel by default in RHEL implement Type Enforcement and Multi-Category Security. SELinux provides fined-grained access control for processes running on top of the Linux kernel. The sVirt project builds on SELinux to provide an infrastructure to allow an administrator to define policies for virtual machine isolation. Out of the box sVirt ensures that a virtual machine's resources cannot be accessed by unauthorized processes (or virtual machines). SELinux and sVirt comprise an infrastructure that provides a level of security and isolation unmatched by other x86 hypervisors.

Powerful memory management: KVM inherits the powerful memory management features of Linux. The memory of a virtual machine is stored as memory is for any other Linux process and can be swapped, backed by large pages for better performance, shared or backed by a disk file. NUMA support allows virtual machines to efficiently access large amounts of memory. KVM supports the latest memory virtualization features from CPU vendors with support for Intel's Extended Page Table (EPT) and AMD's Rapid Virtualization Indexing (RVI) to deliver reduced CPU utilization and higher throughput. Memory page sharing is supported through a kernel feature called Kernel Same-page Merging (KSM). With KSM more virtual machines can be consolidated on each host, reducing hardware costs and improving server utilization.

Broadest range of hardware support: Since KVM is a part of Linux it leverages the entire hardware ecosystem, so any hardware device supported by Linux can be used by KVM.

Most efficient storage and redundancy support: KVM is able to use any storage supported by Linux to store virtual machine images, including local disks with IDE, SCSI and SATA, Network Attached Storage (NAS) including NFS and SAMBA/CIFS or SAN with support for iSCSI and Fiber Channel. Multipath I/O may be used to improve storage throughput and to provide redundancy.

In addition, in the last 4 years since KVM was added to Linux, a lot of work has been done to improve and optimize Linux for KVM to be a better hypervisor. For example the Linux memory manager has been extended to support memory page sharing to dynamically scan memory of virtual machines to "de-dupe" memory or to support transparent huge pages to optimize performance for large memory systems. Further, the Linux scheduler supports resource prioritization and thresholds using cgroups.

Why is Linux the Ideal Operating System for Enterprise Computing in the Future?

The TCO advantages of Linux as an applications server operating system over other operating systems including Solaris and Windows has been well-documented[14]. Linux is versatile; even Windows desktop users can transparently use Linux clouds on the back-end. Moreover, Linux supports the widest range of architectures and hardware devices and it performs, scales, is modular and energy-efficient, and supports an extensive and growing ecosystem of ISV applications[2]. Using Linux for Enterprise IT, customers can:

[12] http://www.spec.org/virt_sc2010/
[13] http://www.atsec.com/downloads/white-papers/kvm_security_comparison.pdf
[14] Robert Francis Group, "TCO for Application Servers: Comparing Linux with Windows and Solaris, http://www-03.ibm.com/linux/whitepapers/robertFrancesGroupLinuxTCOAnalysis05.pdf

Build a flexible scalable architecture: The Linux kernel supports a degree of componentization and customization that is unrivaled among general purpose operating systems, and is remarkably adaptable to a broad range of computing environments. It's configurable and powers everything from a hand-held device to a laptop to a supercomputer. Linux today supports more hardware devices than any other operating system in history[15].

Optimize workloads to drive efficiencies and innovation across the business: Linux supports the broadest range of workloads from traditional transaction processing to high-performance computing (HPC). With over 78% penetration[16], Linux dominates the Top500 list of supercomputers – yet another testament to the affinity of Linux to large-scale data centers.

Drive standardization, interconnectedness, and collaboration: As virtualization and cloud offerings continue to grow, there's evidence of incompatibility between these nascent offerings[17]. The lack of standards and the resulting potential for lock-in into a specific vendor's proprietary islands of insecure and hard to manage offerings, are key concerns. By standardizing on Linux workloads, customers will have the unique flexibility to deploy private, public, or hybrid clouds to drive an unprecedented level of interconnectedness and collaboration.

Leverage an extensive application and ISV ecosystem: With thousands of Linux compatible and certified applications available, customers have many options for their specific workloads. By standardizing on Linux in their departments and data centers, these users will be able to transparently extend these capabilities to a virtualized environment or private, public, and hybrid clouds.

Benefit from attractive licensing costs and terms: While there still exists some misconception that Linux is always free, the majority of enterprise production deployments are commercially licensed and supported. Despite this, a comparative application server TCO study[12] found that Linux could reduce on-going licensing and maintenance costs especially when per-processor licensing models are used. The study participants who used Linux drove their systems to higher utilization levels and achieved greater consolidation benefits.

Operate energy-efficient IT infrastructures: Beyond its advantages in licensing, Linux is a very cost-effective system to enhance data center energy efficiency. With increasing usage in small, power sensitive devices such as netbooks, embedded devices, and other mobile devices, Linux has benefited from a great deal of research in lowering total energy consumption. With features like the tickless kernel and PowerTop[18] that avoid or mitigate traditional high-alert processor time-keeping in somnolent or quiescent states, Linux has become more energy-efficient relative to competitors. When combined with other energy reduction strategies[19] within virtualized or cloud data centers, Linux helps significantly lower the TCO.

Access and use a large existing skill pool and IT resources: Enterprise IT organizations are very often challenged to hire and train the best skills to run and operate their environments. They seriously consider and evaluate this large expense during the purchase and deployment of technology. Fortunately, because managing and developing for Linux are common skills, customers can easily access a large existing skill base or hire new personnel easily in the market.

Improve patch management, security, and availability: Keeping the IT systems up-to-date and compliant with the corporate security requirements is complex, because some vendors release a large number of

---

[15] "The Linux Driver Model: A Better Way to Support Devices", https://www.linuxfoundation.org/en/Linux_Device_Driver_Model#Executive_Summary

[16] Sean Michael Kerner, "Linux dominates top 500 supercomputer list, November, 2009, http://blog.internetnews.com/skerner/2009/11/linux-dominates-top-500-superc.html

[17] Bob Sutor, "To Deliver on the Promise of Cloud Computing, Follow an Open Path", Linux + DVD, April, 2009, pp. 62-63.

[18] Stephen Shankland, "Linux Coders Tackle Power Efficiency", http://news.cnet.com/Linux-coders-tackle-power-efficiency---page-2/2100-1007_3-6192865-2.html?tag=mncol

[19] Data Center Efficiency, http://www.ibm.com/itsolutions/energy-efficiency/

updates on a frequent basis, the tracking and evaluation of the appropriate software updates is time-consuming, the patch deployment must be managed to reduce downtime, and auditing and reporting is required. In a study[12] comparing Linux deployments with Windows and Solaris, it was found that it is easier to lock down a Linux system and deploy patches, and that Linux patch deployments reduced downtime.

Enhance utilization and TCO through virtualization and consolidation: Today, virtualization is a mainstream technology in most data centers, enterprises, and cloud platforms. Simply put, virtualization enables the ability to abstract operating system or application instances from the underlying platform. With virtualization, Windows or Solaris images or applications, can be hosted and run on top of a Linux platform. At the core of a broad array of virtualization technologies are hypervisors that actually make virtualization possible. With complementary "self-service" systems management tooling, virtualized IT resources can be dynamically assembled and provisioned in a matter of minutes instead of months.

Build very scalable and efficient cloud computing environments: Cloud computing promises to provide dynamically scalable and often virtualized IT (hardware, software, and applications) resources as a service transparently to a large set of users who may possess a broad but differing range of knowledge, expertise in, or control over the technology infrastructure.

Today, many large public cloud-service providers such as Google and Amazon use Linux[20] in their large-scale, high-demand, extremely-elastic cloud environments – another testament to the value of open source for cloud computing. Many private cloud deployments at enterprises are combining Linux with KVM; furthering the momentum around open source virtualization.

IBM Delivers a Broad Portfolio of x86 Systems with Red Hat Enterprise Virtualization

IBM's history of close collaboration and joint innovation with Intel, Red Hat, and other ecosystem development partners has consistently resulted in systems that leverage the full power of x86 processor performance built for enterprise grade IT environments. IBM-led initiatives in Linux, the Trusted Computing Group, and Blade.org, the industry consortium driving open innovation in blade-based solutions that today has over 200 members, are some key examples of this open joint innovation and development efforts. Recent joint achievements that provide IBM differentiation with System x include:

An outstanding portfolio of x86 systems: The latest series of eX5 and M3 servers based on Intel Xeon 7500 and 5600 series processors continue that collaboration to deliver performance that no other manufacturer can match.

eX5 servers: IBM's eX5 line of servers, which encompasses servers, flash technology, a way to expand the memory footprint considerably and greater server flexibility, significantly improves the economics of large-scale, memory-intensive x86 workloads.

The System x3850 X5 comes in four- and eight-processor configurations using the Intel Xeon 7500 processor. This 4U box can hold two-, four- or eight-core Xeon 7500s. The eX5s have outperformed anything the competition can offer. The x3850 X5 scored the highest number of transactions-per-second ever achieved by a four-processor system on the TPC-E benchmark -- the top spot for four-socket and 32-core systems using Nehalem EX on VMware virtualization benchmarks. The BladeCenter HX5 comes in 2- and 4-processor configurations. It also features the Xeon 7500 or the Xeon 6500. It has eight I/O ports (two sockets). IBM has engineered it to be able to hold up to a half TB of memory on one blade.

M3 servers: This new line of System x products includes the x3650 M3 and x3550 M3 rack servers, the x3500 M3 and x3400 M3 towers, the BladeCenter HS22, and the virtualization-optimized BladeCenter HS22V. These provide 50% more cores and a performance hike of about 50% compared to the previous generation. IBM lists consolidation ratios of 20:1 compared to models from three or four years ago[21]. In addition, IBM has released the two-socket dx360 M3 iDataPlex server. Aimed at power-intensive computing, this form factor lets blades and racks play together.

[20] WebHostingUnleashed, "Cloud-Computing Services Comparison Guide", http://www.webhostingunleashed.com/whitepaper/cloud-computing-comparison/

[21] IBM benchmark tests

In all cases, IBM emphasizes its virtualization credentials as well as lower power and more processing capability. The HS22V, for instance, is claimed to enable 30% to 50% more virtual machines on a single blade, has memory that consumes 15 percent less power, and runs Java applications up to 43 percent faster than IBM's prior-generation, two-socket blades. All of these servers use the Xeon 5600 processor.

Collaborative innovation drives performance, energy-efficiency, reliability, and enhanced security: Such demonstrable performance and reliability gains are a result of the innovative IBM X-Architecture that encompasses IBM System x and IBM BladeCenter servers, IBM System x iDataPlex solutions, energy-efficiency and cooling via IBM Cool Blue technology and proactive management. For virtualized environments, these systems when harmonized with open source virtualization such as KVM deliver added enterprise class benefits.

Best performance and scalability for server consolidation with IBM System x + Red Hat Enterprise Virtualization: SPECvirt_sc2010[22] is SPEC's first industry-standard benchmark addressing performance evaluation of datacenter servers used in virtualized server consolidation. KVM is the first and only hypervisor used to date for a published result with a score of 1169 on 72 VMs[23] on a 2 socket, 12 core Intel 5600 processors @3.33 GHz (IBM x3650M3) with RHEL5.5 host and guests. Likewise, on an IBM System x3690 X5, the published result is 1369.23 on 84 VMs. IBM continues to work collaboratively with several industry partners to define procedures and techniques to optimize systems with KVM for better performance and scalability.

Reliability, availability, and serviceability: IBM has made sustained technology investments to enhance functionality, performance and reliability across its portfolio of systems and technology offerings. Investments in semiconductor processor technology and architecture, and RAS (reliability, availability, serviceability) features that are part of the mission-critical System z mainframes have been systematically implemented across the IBM portfolio and next generation energy-efficient data centers. A recent survey[24]- which we believe to be the most comprehensive recently published study of high performance computing application workload performance, system utilization and system availability using a Linpack peak performance benchmark -- found that the x86 based IBM HS-21 based BladeCenter returned availability figures of 100% compared to a 95% availability for a SuperMicro based x86 system. This is an illustration of how IBM's approach of bringing mainframe-inspired RAS features into x86 based servers helps enterprise computing at all levels.

Similarly, Linux surveys by industry analysts show that CIOs, IT managers, and system administrators generally consider Linux to deliver the reliability needed for business-critical workloads. Industry-standard hardware running Linux has reached a level of maturity where one can configure fault-tolerant systems that are strong enough to match UNIX systems on proprietary hardware. For example, the Yankee Group 2007-2008 Global Server Operating System Reliability Survey[25] found that the top Linux distributions Red Hat Enterprise Linux and Novell SUSE Linux notched the biggest reliability improvements in their most recent survey. Each of them decreased per server per annum downtime by an average of 75%.

Enhancing energy efficiency: Another example of how IBM innovates to achieve leadership in systems design is in its use of complex CFD (computational fluid dynamics) to optimize the thermo-electrical-mechanical properties when laying out components (processors, memory systems, and interconnect hardware) and cooling systems within servers. This helps achieve optimal energy efficiency within the power/cooling envelope for datacenters. The resulting efficiency gains are significant -- for a 1U configuration, the energy consumption has been reduced from about 19W to 6W for the IBM System x iDataPlex server. Each Watt saved translates to about $7 of customer savings/server/year. Lower energy consumption results in lower TCO (total cost of ownership) and higher reliability as components operate at lower temperatures.

---

[22] http://www.spec.org/virt_sc2010/

[23] http://www.spec.org/virt_sc2010/results/specvirt_sc2010_perf.html

[24] Alan Simpson, Mark Bull, and Jon Hill, Identification and Categorization of Applications and Initial Benchmarks Suite, PRACE Consortium Partners, 2008.

[25] Unix, Linux Uptime and Reliability Increase, Yankee Group, 18 Jan 2008

Improving systems management: IBM also innovates on server management to optimize the IT administrator experience for its System x servers. IBM Systems Director is a platform management foundation that streamlines the way physical and virtual systems are managed across a multi-system environment. Leveraging industry standards, IBM Systems Director supports multiple operating systems and virtualization technologies across IBM and non-IBM platforms. Through a single user interface, it provides consistent views for visualizing managed systems and determining how these systems relate to one another while identifying their individual status. This improves the IT staff efficiency as well and helps them to be more responsive to business needs.

Building enterprise grade security with KVM and IBM System x: The expected number of vulnerabilities for KVM and other recent x86 hypervisors (Xen, VMWare, Hyper-V, etc.) should be roughly equivalent as these hypervisors are similar in terms of program size, program complexity, and maturity of code. But, in reality - according to an analysis derived from CVE[26], a dictionary of publicly known information security vulnerabilities and exposures - KVM continues to consistently exhibit fewer common vulnerabilities and exposures compared with VMware (Fig. 4).

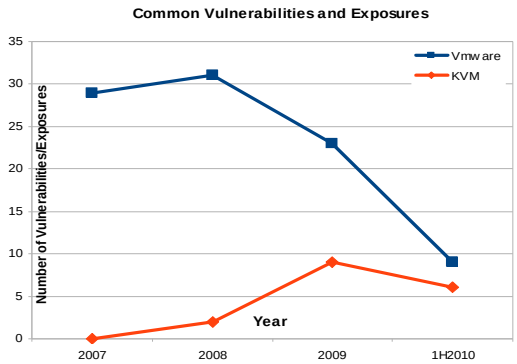

Fig 4: History of Common Vulnerabilities and Exposures

Further, IBM and Red Hat are working together for Common Criteria Certification for KVM which should bring it up to the same level as other x86 hypervisors. This certification process will bring KVM up to the same Evaluation Assurance Level (EAL) as VMware ESX and MS Hyper-V which are at EAL4+[27]. While, XenServer is still in evaluation under the UK scheme at EAL2+[28].

The IBM System x security solution is built on the Trusted Computing Group[29] (TCG) industry standard with an integrated Trusted Platform Module (TPM) hardware device that facilitates higher levels of system security[30]. This security solution is also currently implemented with the UEFI[31] (Unified Extensible Firmware Interface) specification that helps hand off control of the system for the pre-boot environment (i.e.: after the system is powered on, but before the operating system starts) to an operating system, such as Windows or Linux. UEFI provides a clean interface between operating systems and platform firmware at boot time, and supports an architecture-independent mechanism for initializing add-in cards. With the continuing integration of the IBM System Director Security solution with KVM, customers can benefit from an integrated enterprise security solution for secure multi-tenancy, integrity, threat, compliance and patch management functions using a single pane of glass. This level of integration on System x provides an unparalleled level of differentiation especially in security management for KVM.

[26] http://cve.mitre.org
[27] http://www.commoncriteriaportal.org/products/#OS
[28] http://www.cesg.gov.uk/find_a/cert_products/index.cfm?menuSelected=1&displayPage=152&id=450
[29] http://www.trustedcomputinggroup.org/
[30] http://www-03.ibm.com/systems/x/hardware/enterprise/xarchitecture.html
[31] http://www.uefi.org/about/

Comprehensive migration services to Linux and support from IBM: Customer concerns on migration are centered on the following issues:

- Technical: Can it be done?
- Costs: Can it be done within budget?
- Schedule: Can it be done on time?
- Skills and culture: Are the required resources available?
- Operational: Will it work?

IBM Linux Migration Services: IBM addresses customers' migration concerns through its (1) proven track record of migrating tens of thousands of applications and database migrations, (2) promise of quick return on investment through increased asset utilization and potentially significant lower software costs after consolidating server and workloads to achieve rapid migration payback, and (3) a well defined migration strategy for infrastructure, databases, enterprise applications, and custom migrations if needed.

IBM Linux Migration Processes and Tools: The process utilizes tools wherever possible. These are standard UNIX tools, commercial tools and custom IBM tools. In addition, IBM reuses code wherever possible. Reusing code allows for more cost effective migrations and avoids recreating the wheel. For custom applications, IBM works closely with client teams on requirements since they are the ones that know how the current system works. The IBM migration process synchronizes code at each iteration of the new developments and for bug fixes. Working with ISV apps and databases, IBM will also upgrade to new versions if the customer wishes. The IBM process believes in testing and re-testing extensively. Testing is an integral part of the methodology. This means testing on an iterative basis and also collaborative basis with customer and IBM teams. And lastly, benchmarking of the new system against the old system is done to make operational capabilities and expectations are met.

Further, with Tivoli Provisioning Manager for Images on System x, customers can manage the virtualization infrastructure and accelerate provisioning to target computers by building a virtual image once and deploying it multiple times. In addition to helping with migration, Tivoli Provisioning Manager for Images helps perform the following tasks:

- Provide hardware and hypervisor-independent imaging across x86 platforms from a single easy-to-use console
- Optimize server assets with anywhere-to-anywhere server image conversions
- Reduce the total number of images
- Automate the server consolidation operations and help to reduce energy costs.

IBM Linux Support: IBM provides telephone or online access to IBM support specialists through IBM's Linux Support Line and Linux Advanced Support offerings. This is available for all users from small to medium businesses running just a few Linux servers or mission critical enterprise customers requiring highly customized and proactive support. The offering include options for 9x5 or 24x7 support, unlimited calls and callers, client customized response times, and support for IBM and non-IBM Servers.

Customer Examples Highlighting the Business Benefits of IBM and Red Hat Enterprise Virtualization through KVM

IBM has extensive experience in working with enterprise customers to set up, optimize, or migrate their enterprise workloads from almost any computing environment to any environment. Enterprise customers looking to migrate their workloads to System x-Linux can draw inspiration from the following real-world customer cases featuring many diverse enterprise workloads on IBM System x-Red Hat Enterprise Linux environments with KVM.

IBM System x and Linux with virtualization powers NTT's cloud computing service

Business Challenge: When NTT was developing their new cloud computing and hosting service offering, BizHosting Basic, they needed a secure and scalable environment for sharing resources among multiple users and organizations and for managing them simply and flexibly.

Solution: NTT Com, an original Red Hat Enterprise Virtualization beta customer, worked with Red Hat during the development of the KVM-based virtualization platform. The companies continued collaboration today with NTT Com's selection of the Red Hat Enterprise Virtualization platform for its pioneering commercial cloud service.

Software**:** Red Hat Enterprise Virtualization

Hardware**:** Intel Xeon processor-based IBM System x servers

Benefits: IBM System x servers and Red Hat Enterprise Virtualization deliver a secure, scalable and robust foundational virtualization layer for the NTT cloud offering.

Great Offshore reduces cost, increases performance and scales growth with virtualization on Linux

Challenge**:** To design a robust, stable, high-performing, secure, and cost-effective Enterprise Resource Planning (ERP) system based on a combination of SAP and IBM DB2 technologies

Solution**:** Great Offshore selected Red Hat Enterprise Linux as the operating platform for its business-critical SAP and IBM DB2-based Enterprise Resource Planning system

Software**:** Red Hat Enterprise Linux Advanced Platform with integrated virtualization and clustering technologies, SAP Business Suite (150 users) with modules such as Financials & Controlling (FICO), Sales & Distribution (S&D), Material Management (MM), Project System (PS), and Human Resources (HR), Database DB2 (Approx. avg. size of the database is 200 GB)

Hardware**:** 8 Intel Xeon processor-based IBM System x servers

Benefits: Client implemented a robust and cost-effective ERP system on IBM System x servers with Red Hat Enterprise Linux Advanced Platform with its integrated virtualization and clustering technology. This has resulted in lower IT costs, simplified management, reduced systems maintenance, and increased scalability and performance with SAP Business Suite.

Etisalat transforms IT infrastructure through scalability and cost savings with virtualization

Business Challenge**:** To set up a scalable, flexible, and cost-effective IT infrastructure to support the organization's objective of rapidly increasing its telecommunications operations in Sri Lanka

Solution: Etisalat Sri Lanka selected Red Hat Enterprise Virtualization to resolve its business challenges by providing on-demand business scalability with significant cost savings and enhanced operational efficiencies

Software: Red Hat Enterprise Virtualization; guest operating systems, including Red Hat Enterprise Linux and Microsoft Windows 2003 and 2008 Server Edition; Microsoft SQL Database Server 2005 and 2008

Hardware: IBM and HP Intel Xeon processor-based Blade Servers (2 Quad Core and 8 GB / 24 GB memory)

Benefits**:** Achieved significant cost and time savings; gained the ability for real-time provisioning of servers; improved operational efficiencies and development of a highly consistent and scalable virtualized IT infrastructure; reduced power, space, and management overhead costs; gained significant internal customer satisfaction.

IBM launches development and test cloud using Linux virtualization

Business Challenge**:** Cloud computing, with its promise of improved ROI on IT dollars through shared resources, is an attractive proposition for enterprises. Wide adoption of cloud computing will depend on how well customers are convinced about the robustness, scalability and security of clouds.

Solution: IBM has selected Red Hat Enterprise Virtualization as the hypervisor technology for its new IBM Smart Business Development and Test on the IBM Cloud.

Software: Red Hat Enterprise Virtualization

Hardware: IBM System x servers, storage and networks

Benefits: The solution combines Red Hat's virtualization software with IBM's server, storage and network hardware, and IBM's enterprise support and services to form a robust, scalable and secure cloud infrastructure offering for enterprise customers.

Hikal powers mission-critical ERP solution with System x and Linux with virtualization

Business Challenge**:** Needed to use technology to reduce operational cost, and deliver a sturdy operating system for Hikal's Enterprise Resource Planning (ERP) and BIS applications

Solution: Hikal chose Red Hat Enterprise Linux Advanced Platform, including its built-in clustering technology, as the base platform to implement its mission-critical Oracle E-Business Suite

Migration Path**:** Red Hat Enterprise Linux 3.2 to Red Hat Enterprise Linux Advanced Platform 5.1

Software: Red Hat Enterprise Linux Advanced Platform with integrated virtualization and clustering, Oracle E-Business Suite, Oracle DB (Approx. total database size is 1 TB)

Hardware: Intel Xeon processor-based IBM System x servers

Benefits: ERP project completed in a record time of close to six months, complete freedom to add multiple applications to infrastructure, Red Hat Enterprise Linux and clustering technology delivers continuous uptime with zero maintenance issues, has helped in implementing the latest available features of Oracle E-Business Suite.

Why KVM will build Greater Momentum in x86 Virtualization in the Future

Over the last five years, many enterprises have obtained considerable business value with virtualization: reduced IT capital costs, lower operational expenses, improved infrastructure availability, and better performance. While many proprietary x86 virtualization solutions have been very successful, solutions based on Kernel-based Virtual Machine (KVM) - representing the harmony of several key recent IT trends: rise of enterprise x86 systems, virtualization, and open source and Linux, and cloud computing - will continue to advance with greater momentum for the following reasons:

1. Many enterprises have already implemented KVM and have obtained considerable benefits from these virtualized environments
2. KVM development is backed and supported by several major IT solution providers. In addition to the broad Linux community, KVM is supported by some of the leading vendors in the software industry including Red Hat, AMD, HP, IBM, Intel, Novell, Siemens, SGI, and others
3. With the planned future development of additional enterprise features and function and continuing support for KVM, customers will continue to reap the increasing benefits of fully-supported enterprise grade open source virtualization solutions
4. Proprietary virtualization solutions "lock-in" clients who may not have the appetite to continue to pay the escalating software and services costs
5. KVM performance and scalability continues to be the best in class for many enterprise workloads, including server consolidation and ERP workloads. With access to the source, enterprise clients will have the flexibility to tailor their virtualization environment optimized for their specific workloads
6. IBM System x standards-based security solutions developed in partnership with ecosystem partners will continue to support and integrate KVM as a strategic component. This will provide enterprise grade security critical in enterprise computing
7. As cloud computing adoption continues to grow and since the cloud ecosystem is large, diverse, and is evolutionary (rather than disruptive) with no dominant incumbent, customers will demand that solution providers provide the flexibility that open systems typically provide
8. Many cloud services providers including IBM, Rack Space, NTT, and others have already deployed large-scale KVM powered infrastructures. This trend will only continue to accelerate in the future.