

The Value of Virtualization for Web Environment

Jean-Pierre Kuilboer, University of Massachusetts Boston, 100 Morrissey Blvd., Boston, MA 02125, USA

Noushin Ashrafi, University of Massachusetts Boston, 100 Morrissey Blvd., Boston, MA 02125, USA; E-mail: noushin.ashrafi@umb.edu

ABSTRACT

The success of many e-commerce applications depends on their reliability, robustness and security. Designing a web server architecture that advances these properties without compromising performance is a challenging task. Effective and efficient web applications demand a virtual operating environment which is automatic, integrated and based on open standards. A virtualized environment makes the most efficient use of its resources by sharing resources and providing what is needed only when it is needed. This paper, examines the role of virtualization as a booster of IT utilization and as a driver of IT operational agility.

Keywords: virtual system, operational agility, web-server, integration, flexibility, operational access

1. INTRODUCTION

While the business world has recognized the need to be agile in order to survive a tumultuous and unpredictable environment, the IT industry has been busy through the 1980s and 1990s investing hundreds of billions of dollars into creating complex, underused IT environments. This costly infrastructure has not been able to cope with constantly changing needs and a perpetually evolving business climate. In 1993 Brynjolfsson (1) examined the productivity paradox and questioned the value of increasing IT investment without corresponding benefits. The debate is still ongoing and has produced inconclusive results (2) Following drastic cost cutting measures of the early 2000s, the market settled for flat budgets, increased workloads, and intense pressure to achieve more with less. Inputs have been affected by less extravagant resources requests and helped by the tumbling hardware prices while outputs are increasingly raised through better utilization of existing computing assets. For example IDC estimates that the overall server average system selling price fell by about 16% from 2001 to 2002 (3). Yet, with business increasingly conducted online and with advertising, sales, and support highly dependent on the internet technologies, high volume web servers are now subject to intense scrutiny. The upgrading of existing workstations, servers, networks, or storage by adding additional hardware to accommodate new requirements often competes with other business priorities and faces reluctance from the other business functions to fund perceived cost centers. Better integration and consolidations have been sought as a way to improve IT effectiveness and efficiency.

Moving forward involves the coordinated efforts to integrate technologies, processes, and people while satisfying business priorities. Innovative solutions are percolating through industries and reaching the end-users. One of the most discussed topics of the last two years has been virtualization. Virtualization is not a new concept; after going through a golden age in the mainframe environment in the 1970, a dark age brought by the underpowered PC era, and a long incubation period as computing power followed the Moore's law since the 1980s, it is ready for a comeback. Experts in the field prescribe a resource virtualization renaissance to address problems such as security, performance, and reliability (4).

The old utopia of integration and better utilization across departments, firms, and the business eco-system at large is now enabled, in part by a number of technologies under the umbrella of virtualization. Virtualization is a technology about workload consolidation to drive up utilization by subdividing the resources of modern Information Technology. It divides one piece of real hardware into numerous virtual ones, where each is managed separately and is more resource constrained than the original machine. The technology can be deployed as both

hardware virtualizations through partitioning and through software virtualization where each application brings its own set of configurations on-demand.

2. DEFINITION

Defining a concept is always open to different interpretations. For example agility, reliability, and flexibility have many facets and are defined differently by diverse professional community. Singh (5) defines virtualization as a framework or methodology of dividing the computer resources into multiple execution environments, by applying one or more concepts or technologies such as hardware and software partitioning, time sharing, partial or complete machine simulation, emulation, quality of service, and many others.

It is easy to perceive how this definition applies equally well to the concept of time-sharing adopted on mainframe in the early years of computing and the latest trends on operating system virtualization allowing for example a single laptop/desktop to run Apple OS X, Linux or Microsoft Windows Vista.


Figure 1 illustrates a generic e-business architecture with virtual servers.

3. DEPLOYMENT

While the hardware cost by itself was the main driver of virtualization in the 1960s, its recent incarnation is motivated by different factors. Standard IT budget could let both hardware and software stand at less than 10% of yearly IT expenditures. Thus the largest potential for added value is through better processes, people and utilization of existing resources. To achieve higher utilization, keep pace of changes in configuration, and maintain a high level of availability of systems (e.g. on a 24/7 e-business environment) virtualization should provide means to provision, deploy, and maintaining system using off band capabilities.

Chris Wolf (6) contends that virtualization is going mainstream. Virtualization is now considered a serious tool for production environment. The use of virtu-

Figure 1. Virtualization architecture



alization expands from its traditional use as a development platform to critical production environments.

Leading IT industry software companies are jousting to have a role in this evolution. Big players such as Microsoft, IBM, EMC, and HP are considering virtualization either as a hardware, software, or storage advantage and are acquiring smaller companies that have a head start in virtual systems. Microsoft, once reluctant to the idea has recently adopted a more flexible stance with its purchase of Connectix and the release of *virtual server*. Anticipating the market trend Microsoft acquired Softricity, a leader in virtualization and application streaming solutions, and promoted Windows Server Virtualization as a part of the new Longhorn platform using software code-named Viridian as a hypervisor by 2008. It will only be available on 64-bit versions of the server operating system, and similarly only available if the physical hardware supports either IVT (Intel's Hardware Assistance for Virtualization) or AMD-V (AMD's Hardware Assistance for Virtualization). The open source has also been working on evangelizing virtualization and the Xen project is the most widely distributed hypervisor coming through this avenue.

IBM (7) claims that today CIOs are making significant headway in *On Demand Business* and aims its virtualization to the idea of utility computing. They are leveraging their mainframe experience (i.e. They introduced the concept in the 1960s) and present virtualization and other on demand technologies to deploy services and to lower computing infrastructure and labor costs. They are supplementing their portfolio with acquisitions such as Rembo that provides the ability to maintain and deploy software images, automated installation and customization across multiple locations. The product has security features that can protect workstations used by multiple people in a virtualized environment by automatically "wiping away" operating systems and personal data after each use and re-installing clean software thereby addressing security concerns.

HP has similar ambitions. HP virtualization solutions will help an organization pool and share IT resources, lowering its costs by optimizing utilization while increasing agility, enabling the rapid response to changes in the marketplace. HP's offerings encompass three levels: element virtualization, integrated virtualization, and complete IT utility, which represent increasing business value and strategic importance to an organization.

EMC has been the major player with virtualization products it obtained from its 2004 acquisition of VMWare that provided a range of products including VMWare Workstation, GSX server software, and ESX server with its hypervisor directly installed on bare metal machine without additional host operating system (6).

In the enterprise storage domain, There is a great deal of excitement in the storage industry about the potentials of virtualization technology to reduce total cost of ownership (TCO) and increase utilization of existing enterprise storage systems. Virtualization leverages the connectivity that is provided by storage area networks (SANs), by creating an abstraction layer between the servers and the SAN. This abstraction enables the servers in the SAN to view the physical storage as a common pool of capacity. While forthcoming, this view has not yet been fulfilled for lack of complete solution across heterogeneous vendor products (9).

As the most established IT industry contenders push for its deployment, virtualization is deemed to become one of the dominant trends in the coming year and is one of the most exiting areas of IT. Virtualization is commonly quoted when associated with main aspects of IT. Primarily organizations use virtualization in conjunction with servers, operating systems, applications, storage, data center, and networks (10).

3.1 Web Server and Virtualization

Increasingly business is conducted online. With advertising, sales, and support depending on the internet technologies, high volume web server are critical to business continuity.

Efficient web applications demand a virtual operating environment which is resilient, autonomic, integrated and based on open standards.

A virtualized environment makes the most efficient use of its resources by sharing resources and providing what is needed only when it is needed. Not only are resources highly used, excess capacity can easily be used for new or unexpected needs.

Nowadays the success of many e-commerce applications, such as online banking, depends on their reliability, robustness and security. Designing a web server architecture that keeps these properties under high loads is a challenging task because they may conflict with performance.

3.1.1 Value of Virtualization for Web Server

Virtualization is directly relevant to a web server life cycle. From inception to replacement a web server (both physical and logical) is a prime candidate for virtualization. Web servers with dynamic content have particular attributes. They have an accelerated life cycle and need to be developed, provisioned, deployed, and maintained on a short timetable. In addition, in e-commerce, loads are difficult to predict and web servers located in the insecure Internet or a demilitarized zone (DMZ) are subject to intense security risks.

3.1.2 Consistent Platform

The virtual server software hides the physical hardware and creates a generic hardware platform that is consistent regardless of the physical server used to host virtual servers (11). This virtual platform makes the transition from development to deployment to operation seamless and alleviates problem of compatibility among platforms. Load testing can also be achieved by varying the resources allocated to the VMM.

For example VMware ESX allows you to control CPU time, memory space, network bandwidth, and storage bandwidth. Similarly in the Sun architecture, the Solaris Resource Manager software gives its administrators almost unlimited flexibility to assign and isolate resources to specific containers. Resource pools can be changed either manually or automatically on a rule basis.

3.1.3 Isolation of the Development Environment to the Production Environment

During development it is not rare to have systems crash unexpectedly. The partitioning of resources into secure containers will avoid side effect of the crash on the host system and concurrent VM. Virtualization will also provide almost instant recovery of an instance through a pre-staged process.

With the many generations of computer hardware reaching the market and constituting a complex mix of basically incompatible resources, virtualization in the server space hides the physical hardware from the virtual servers (12). The virtual server presents a generic consistent interface making restoring a virtual server easier.

3.1.4 Increased Automation

Manually tuning is often infeasible due to unpredictable loads and requests arrival rate. Computer systems, Web servers included, could use analytic performance models as a guide to dynamically adjust configuration parameters such as the number of active processes (13,14,15). Virtualization technology is well equipped to solve the problem by automating resources allocation initially and adjusting them based on simple rules sparing human intervention.

Automation of the provisioning, maintenance or recovery process is a problem derived from shifting cost from hardware/software to servicing systems and the unfeasibility to tend to individual server on a timely basis. For years firms such as Landesk, Softricity, and Rembo have been dedicated to such endeavors. Layering virtual solutions above the bare hardware makes the deployment of a number of virtual machines painless and will make the upgrade of virtual machine feasible with techniques such as in VMWare vmotion allowing to move a running virtual machine from one physical server to another.

3.1.5 Higher Hardware Utilization

One of the leading detriments of distributed powerful computers is their low utilization and stretched low ratio of server/administrator. Hardware costs, while not the burden it once was, still are remarkably ineffective. With rapidly advancing technology organizations want to spare unnecessary hardware purchase, knowing that the next generation will bring more power at less cost. Running multiple virtual servers per physical host adds significant savings that can be better invested in other aspects of the systems.

While potentially using hardware resources to their limit could raise problems with power consumption, its planning is enhanced by better measurement of the existing assets.

Security. As distributed architecture and increasing dependence on the Internet has grown, security has become a critical preoccupation for most organization. Risks and threats are being addressed by virtualization by isolating environment. Before virtualization, control was often a tradeoff for agility. The only way to restrict exposure was to limit the number of applications running on a server,

which lead to reducing flexibility and increasing administrative operating costs. For example for a Internet service provider hosting traditionally was procured by either offering expensive dedicated physical servers or sharing through software running on a unified platform. In the first case customers could have managed service or simply take advantage of the co-location of a data center. In both cases administration still was difficult to control. In the second case restrictions were more stringent, limiting installation or use of customized software packages. With virtual servers this is no longer the case, allowing more freedom in a more robust environment.

4. PERFORMANCE

Performance has always been an issue in IT. In fact database, storage or network experts are quite obsessed with benchmarks and measurements. Processor manufacturers have measurements by MHz for particular architectures and recently they are moving to multi-core CPU in order to take full advantage of multithreading. In networking flows are monitored with bandwidth speed. In such competitive environment any overhead is looked at unfavorably. The concept of Virtualization may be a solution to performance problem, but while it is showcased as decreasing complexity and increasing flexibility, complaints still abound about machines running slowly when running emulation.

One solution is IBM's virtual machine model where each virtual machine is an exact copy of a real machine with limited set of resources, but often optimized by hardware support. High security is achieved through filtering by a control program of instruction potentially damaging to concurrent virtual machine (e.g. memory changes, storage modification). In the Java virtual machine the software provides an abstract layer sandboxing the real hardware from harmful access.

4.1 Sources of Overhead

With full virtualization or emulation the system is simulating a completely different computer chip. Virtual applications running on top of an emulated system incur a level of overhead which almost always lead to disappointment. Even with increasingly powerful processors and faster computer memories the interpretation of programs will be slower than running in their native form.

The additional sources of overhead include the reflection of exceptions and I/O interrupts to the virtual machines, support of virtual timers and clocks, and the translation of I/O channel programs before the VMM initiates I/O. When the application depends on time synchronization with actual time, emulation often leads to incorrect results or excruciating complex solutions (16).

The virtualization control software, called a hypervisor or virtual machine manager, imposes a performance penalty as it manages resources such as memory or input-output. Traditional CPUs were not designed to run a variety of operating systems simultaneously. With new set of processors the major CPU developers are trying to fix some of the performance penalties.

4.2 Improving Performance

Hardware facilitation of virtual machine monitors is seen as one avenue out of the performance bottleneck. Intel has recently launched its VPro line of microprocessor derived from the Vanderpool aimed at boosting both performance and security. AMD has competing line of Rev F processor with Pacifica and Presidio also emphasizing virtualization and security. IBM POWER5 systems combine enhancements to the processor architecture with greatly better firmware to increase its virtualization. The most recent microprocessors architecture with multiple cores is the response to demand for virtualization. This new approach has superseded the long time race for increase in clock speed. The first generation caught software vendors and organization by surprise and subsequently added the non-technical uncertainty in operating systems and applications licensing. Similarly to obtain full benefit from multiple core systems the operating system and applications should be multithreaded. Virtualization can also assign core to alternative operating environment, making full use of the processing capabilities. Since the market place has adopted compromises to the licensing issue ranging from a licensing per processor to reduced licensing cost per core.

5. PRACTICES

IT practitioners and researchers are by nature open to adopt latest technologies. Often burned by the initial hype of new technologies but also helped by subse-

quent ramping up of emerging techniques that make practice simpler or better. Although the market response to virtualization has been quite remarkable, as a practice, it is still in its infancy.

5.1 Future Trend for Virtualization

Currently processors such as IA-32 do not facilitate virtualization and the VMM has to work hard to get the benefit of full system virtualization. Without any hardware assistance virtualization running unmodified guest OS do not get the performance expected. Para-virtualization offer a remedy to performance at the detriment of requiring either special hardware or running a limited set of guest OS customized to run over the particular VMM. The most common Para-virtualization systems are Xen (17) and Denali. A side effect of the customized OS is that they support and provide a limited set of generic devices to the guest OS. This limits the use of most innovative new I/O devices for network or storage support.

The best hope in the near future is the upcoming generation of processors from Intel, AMD, or IBM. Providing the VMM with the possibility to run at ring -1 or have the necessary hardware assist for the common privileged instructions which would otherwise have to be trapped and simulated by the VMM. Similarly experience with virtualization and a healthy competition among future OS providers will lead to more robust guest operating systems, thereby alleviating some of the risks associated with potentially errand commands.

5.2 Security Impact

Virtualization is often perceived as having a positive impact on computing security. The isolation of applications within their own space, transparency, and segmentation of resources under a tight centralized control are potentially confusing to the attacker and make the system more reliable. The network fingerprinting by hacker is more difficult as real resources are opaque to a scan. In an environment with non expert end-user, open Internet access and traditionally less robust platform, virtualization add one layer of indirection which make virus or rootkit infection more difficult. As previously stated recovery from disaster is rapid.

On the other hand integration and aggregation of the pool of resources in a data center offered by virtualization(18) is attractive from an operation point of view but will makes the target more attractive to hackers. In a scenario where the Virtual Machine Monitor is infected by a rootkit, the attacker would be able to take control of a vast array of organizational resources and would go undetected by conventional methods. A VMM running on ring -1 infected by a rootkit would present a virtual virgin interface to virus/malware checker and applications file integrity would also be duped. Similarly virtualization will have to be more coordinated with security tool providers to solve real questions. For example of IDS/IPS rely on physical host data to check for traffic patterns when not relying on signature based detection. With virtualization and knowing where the data is coming from is more indeterminate. Other security issues will arise when n layer system could be based on MAC addresses or IP addresses for authorization. Example could be a Web Server talking to an Application Server talking to a Database Server for dynamic queries. Aside from performance issues of a virtual system, the transition from virtual to physical and back to virtual could become quite confusing especially if virtual machines are dynamically moved from server to server.

The scenario could either lead to consolidation in the market or lengthy alliance process making adoption of best of breed system more problematic.

6. CONCLUSION

Almost after 80 years since virtualization was introduced by IBM (19) around 1927, IT people entertain the same concern that machine utilization percentages are an important measure of data processing management competence. Virtualization across a number of information technologies has made a comeback with a transition from mainframes to distributed systems. Virtualization has rapidly been adopted by small to large organizations. Particularly when mass customization is to be applied to systems, virtualization offers the benefit of central control and the flexibility to have stakeholders have a system adapted to their needs.

REFERENCES

- [1] Brynjolfsson, E. 1993. The productivity paradox of information technology. *Commun. ACM* 36, 12 (Dec. 1993), 66-77.

- [2] Eastwood, M. and Yang, J. 2003. Worldwide Server Workloads 2003: The continuously evolving relationship between applications and servers, IDC #30686, Vol. 1
- [3] Oz, E. 2005. Information technology productivity: in search of a definitive observation. *Information & Management*, Vol. 42, pp 789-798.
- [4] Figueiredo, R. Dinda P.A. and Fortes J. Resource Virtualization Renaissance, *IEEE Computer*, Vol. 38, No. 5, 2005, pp. 28-31.
- [5] Singh, A. An Introduction to Virtualization, Retrieved May 2006 from: <http://www.kernelthread.com>.
- [6] Wolf, C. Virtualization goes Mainstream. Retrieved May 2006 from: <http://redmondmag.com/features/article.asp?EditorialsID=548>
- [7] Popek, G. J. and Goldberg, R. P. 1974. Formal requirements for virtualizable third generation architectures. *Commun. ACM*, Vol. 17, No. 7, 1974, pp. 412-421.
- [8] VMWare, ESX Administrative Guide, 2006.
- [9] Yoshida, H. Virtualization: the Promise and the Reality, *Hitachi Whitepaper*, April 2002.

0 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage: www.igi-global.com/proceeding-paper/value-virtualization-web-environment/33202

Related Content

Improved Cross-Layer Detection and Prevention of Sinkhole Attack in WSN

Ambika N. (2021). *Encyclopedia of Information Science and Technology, Fifth Edition* (pp. 514-527).

www.irma-international.org/chapter/improved-cross-layer-detection-and-prevention-of-sinkhole-attack-in-wsn/260210

On the Study of Complexity in Information Systems

James Courtney, Yasmin Merali, David Paradise and Eleanor Wynn (2008). *International Journal of Information Technologies and Systems Approach* (pp. 37-48).

www.irma-international.org/article/study-complexity-information-systems/2532

Securing Stored Biometric Template Using Cryptographic Algorithm

Manmohan Lakhera and Manmohan Singh Rauthan (2018). *International Journal of Rough Sets and Data Analysis* (pp. 48-60).

www.irma-international.org/article/securing-stored-biometric-template-using-cryptographic-algorithm/214968

Noise Trader

Po-Keng Cheng (2018). *Encyclopedia of Information Science and Technology, Fourth Edition* (pp. 71-76).

www.irma-international.org/chapter/noise-trader/183721

Rough Set Based Similarity Measures for Data Analytics in Spatial Epidemiology

Sharmila Banu K. and B.K. Tripathy (2016). *International Journal of Rough Sets and Data Analysis* (pp. 114-123).

www.irma-international.org/article/rough-set-based-similarity-measures-for-data-analytics-in-spatial-epidemiology/144709