IMPLICATIONS OF STORAGE VIRTUALIZATION FOR SMB

Adriana BARNOSCHI*

Abstract
In today's datacenters, the unexpected increase of virtual and physical servers has led to exponential increases in data-transfer volumes, bandwidth bottlenecks, and connectivity costs. This growth causes major problems for IT managers as they struggle to see more cost-effective, scalable ways to manage the increasing complexity of servers and their associated LAN and SAN environments. Virtualization is a general term that could be applied to: storage systems, databases and networks. In this article, I present an overview of server virtualization architecture. I also specify the goals and benefits of virtualization for a SMB. The article points out some issues about server virtualization: definition, implementing problems and security features. In case of major calamity, you need a disaster recovery plan. The article lists a series of differences and similar parts of DR planning and business continuity plan.

Keywords: *storage management, server virtualization, disaster recovery plan, business continuity planning.*

Introduction

PC storage capacity has grown; users have stored more and more data on their machines, being the primary targets of attackers. So, companies need to backup their information in order to limit data loss and people started to think about disaster recovery. In other words, companies need to stay in business and people are ware of the value of lost data.

I achieved this conclusion by collecting data about natural and human inducted disasters correlated with business process, by studying different points of view of IT experts in storage network solutions for consolidation, by judging from the laws and standards addressed to BCP for improving an organization's information security, by analyzing disaster survival statistics, by making the choices for storage systems from SMB market.

I collected statistical data from analysts of IDC, of Forester Research Inc., of famous companies and I selected the ideas concerning the promised benefits of virtualization technology. I was thinking it's important to know about their work and performances, not just because their word does matter, but their experience facilitates us at storage management into IT departments, it helps us to well-formulate the goals of our projects and business and it contribute to state the metrics of software quality.

The paper is organized as follows:

In sections 2 and 3, I give an overview of virtualization that contains a short history of this technology [1, 9], preferred definitions of concepts [2, 6] and one simple description of host/guest paradigm from the base of virtual machines [4, 5, 6].

Section 4 presents the goals of virtualization when it is used to put into practice a wide range of applications and the benefits of desktop virtualization for SMEs, since this technology has become so popular in storage management.

The section 5 discus about how important is the impact of the virtualization process on day-to-day security management.

Section 6 tells us what disaster continuity planning (DRP) is; why virtualization in a disaster recovery environment is very high on the managers' list; what the gain of using virtualization in DR is? It is important to understand that business continuity planning (BCP) is a

* Associate Professor, Ph.D., Social and Administrative Sciences Faculty, "Nicolae Titulescu" University.

methodology and backup is a process. The logistical plan is called a Business Continuity Plan. DRP is a part of BCP. That's why, when I design a disaster recovery plan I have to go through the methodology phases.

The last section summarizes the contributions of this paper and discusses for future work.

Literature review of virtualization

The technology that has created virtual instances of operating systems has been around in one form or another for years, and has been generally approved into the industries in which top priorities are costs and mobility.

The idea dates rear to the days of mainframes and workstations, a model that required all of the computing to be done on the mainframe with the results of the calculations then displayed on the terminal. This multi-user, time-sharing model maximized utilization of the mainframe's resources simultaneously. It was the efficient standard for years, up through the beginning of minicomputers such as the VAX and PDP.

The idea of virtualization in computing systems is to add a layer of abstraction between two layers in that computer system [1]. This layer allows reducing the management reliance on complicated elements, like building new servers or deploying new applications, while also enabling transiency of the underlying virtualized elements:

Virtualization refers to the pooling of IT resources in a way that masks the physical nature and boundaries of those resources from resource users. In more concrete terms, virtualization is the decoupling of software from hardware. It is the abstracting of the software from the underlying implementation [1].

Server virtualization is the masking of server resources, including the number and identity of individual physical servers, processors, and operating systems, from server users. The server administrator uses a software application to divide one physical server into multiple isolated virtual environments. Server virtualization means the ability that allows multiple independent operating systems to run on the same hardware at the same time.

Virtualization software runs like an application on a computer, separate from the operating system and it avoids hardware and software incompatibility problems. Because the software runs separately from the OS, the different versions of operating systems or other applications can run at the same time [3].

This, along with the need to reduce the number of physical servers in data centers to save money on hardware and power costs, has led to the current attraction with server and desktop virtualization. Today, vendors are building hardware and software platforms that can deliver virtualization solutions at near-native performance.

The recent rush of interest in virtualization has meant large business for ***vendors*** like Sun, IBM and Novell, as well as a host of smaller vendors, who sell virtualization software and services. Much of the business these companies have seen so far has been in the data center as part of server consolidation projects [5].

Theoretical Background

There are three popular ***approaches*** to server virtualization [3, 6]:

1. The virtual machine model,
2. The paravirtual machine model, and
3. Virtualization at the operating system (OS) layer.

Moving up from the bottom there is the hardware layer, followed by the operating system and finally the applications.

Virtualization is being widely embraced by IT industry and smaller organizations are now looking to make use of this technology. This is a resulting in an increased number of product offering as vendors compete to capture a share of emerging SMB market. This reduces a major obstacle to deploying virtualization at the SMB level: *cost*.

Virtual machines are based on the ***host/guest paradigm.*** [3, 6] Each guest runs on a virtual imitation of the hardware layer. This approach allows the guest operating system to run without modifications. It also allows the administrator to create guests that use different operating systems. The guest has no knowledge of the host's operating system because it is not aware that it's not running on real hardware. It does, however, require real computing resources from the host -- so it uses a ***hypervisor*** (VMM) to coordinate instructions to the CPU. The hypervisor (called a virtual machine monitor - ***VMM***) is referred as a virtual manager; it is a program that allows multiple operating systems, which can include different operating systems or multiple instances of the same operating system, to share a single hardware processor. [7]

VMware and Microsoft Virtual Server both use the virtual machine model. [8, 5]

Two new major developments will have a dramatic effect on virtualization technology adoption. On the hardware side, x86 architecture- based microprocessor manufacturers have released a new generation of chips that support virtualization natively. On the software side, the emergence of the open-source Xen* hypervisor virtual machine technology has eliminated much of the performance impact associated with the mediation layer that accompanies full virtualization and software emulation. These developments could completely decouple software from the underlying physical implementation [1].

HP Virtual Connect technology is an interconnect option for HP BladeSystem environments that is used instead of standard pass-through or managed switch offerings. It provides a simpler way to connect blade servers to datacenter networks by creating pools of LAN and SAN addresses that can be assigned dynamically to server bays in software, instead of being hardwired into the servers' individual NICs.

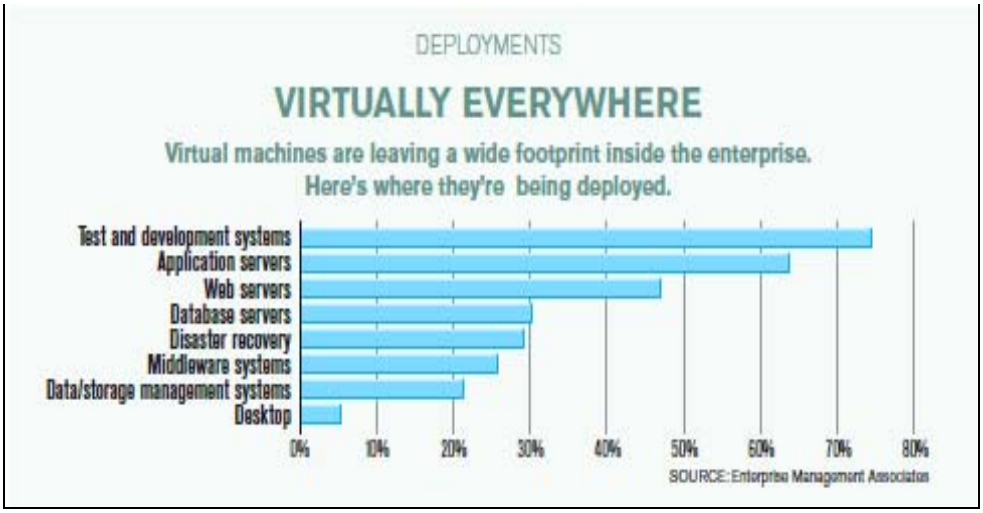

Figure 1: A new study of Enterprise Management Associates (EMA) has established that almost 75 percent of enterprises have deployed virtualization in some form.

The goals and the benefits of virtualization technology

The key ***goals of virtualization*** are:

1.To ensure independence and isolation between the applications and operating systems that run on a particular piece of hardware,

2.To provide access to as much of the original hardware system as possible,

3.To do all of this while minimizing performance overhead.

Although server virtualization can provide great benefits, it's not the only option out there for using virtualization.

- **User density increases dramatically.** The average user density grows by a factor of three on a per-server basis, while number of users per server manager goes up by a factor of between four and five times.
- **Availability improves.** System availability goes up even for basic virtualization. The real benefit comes from an advanced virtualization scenario in which downtime drops by 50%.
- **Scalability.** Once virtualized, an application that needs more scalability can be moved to a server that can fulfill that requirement with little more than a few clicks of the mouse.
- **Cost reductions.** Cost reductions occur across the board, but with future deployments, customers can move to server operating systems that offer unlimited virtualization rights, extending their savings dramatically in many cases.

To get the most from virtualization technologies, keep in mind that the answer to every consolidation or availability problem may not be a single virtualization technology, but instead a combination of complementary solutions.

Hosted virtualization solutions are good options for SMBs that need to fast and securely offer mobile and contact workers. Lambert said: "The virtual machine image has all the attributes of a file, so IT staff can blow away the PC image very quickly if they need to." [9]

Virtualization has proven its ***benefits*** to the organization. "The No. 1 reason [benefit] is efficient use of resources," says Dan Thompson, network engineer at Young America. "Secondary reasons include ease of backups and disaster recovery." With server virtualization, an IT group can run its own backup data center or hot site at a safe distance from the main data center without having to precisely match the original hardware configuration.

The benefits of a virtualized approach to server I/O in the HP BladeSystem environment include the decrease of ***management effort.*** The virtualization of I/O helps reduce the time spent on issues associated with moving, adding, or changing servers. With HP Virtual Connect, the network can be pre-provisioned and the network identity lives with the bay in the enclosure, not the server. Because HP Virtual Connect saves each server profile and the server is stateless from a network perspective, organizations can get back up and running with a new bare metal server very quickly, without loading drivers and so forth.

Life would be great if you could configure and deploy your Web server and then you are going to accomplish new tasks. At start moment, your server is connected to Internet. From time to time, server configuration requires reviewing and modernizing, as new technologies and threats continuously come up. You should enlarge and maintain a list of resources on security problems and software updates relevant to your system and applications. It’s important to set up o procedure for monitoring these sources of information. Not all updates will be applicable to the configuration of servers and security requirements, so it is necessary to evaluate these updates from the ***applicability*** point of view. The expert advice consists of installing the updates on isolated environment and run a series of trials.

Desktop virtualization eliminates the testing of multiple configurations, it increases data security and it improves system stability.

Virtualization is not a panacea. Without the right infrastructure and management tools, virtualization won't stop the complexity and inefficiency.

IT organizations will increasingly embrace virtualization in their infrastructures, allowing creation and management of compute resource pools that can be easily provisioned and changed to meet organizations' fluctuating demands and service-level commitments. Virtualization of all IT resources will play a role, including server, storage, and I/O virtualization.

Managing secure server virtualization

Similar to Web services and Wi-Fi in the past, virtualization is the present love of IT. Departments in enterprises, small- and medium sized businesses (SMB) and universities are deploying virtualization in vast numbers, looking forward to economies through server consolidation projects and reducing costs of office systems.

And like the other hot technologies of their time, virtualization is being deployed with little or no thought to ***security***. The benefits that IT vendors could realize by server virtualization are more important than the real problems the technology can explain the security and compliance.

"Security is part of our DNA at Esurance," says Marjorie Hutchings, director of Internet operations at the San Francisco-based online insurance company.

Security experts say that, although the concept of virtualization is decades old, the current usage models are still relatively new and the security implications have yet to be fully worked out.

"The security issues really depend on the usage model," says Nate Lawson, a senior security engineer at Cryptography Research in San Francisco, who has done research on the security models of virtual machines.

"In server consolidation projects, there's no firewall between the virtual machines, so if one gets compromised, it can be a platform for attacks on the others. Also, some people may be putting two different virtual machines with different security levels on the same host. No one has really done a full security analysis of VMware, so it's possible that a well-designed attack could allow a compromised virtual machine to escape from its partition."

Officials at VMware challenge this notion, saying that server virtualization increases security in most cases, especially when a *hypervisor* is used [14].

The hypervisor is the software abstraction layer between the bare metal and the operating system instances that run on top of it. The question is: *how vulnerable is this software*? The stakes are high; if the primary hypervisor is compromised, it's possible to own all of the virtual machines that run on top of it. If the hypervisor becomes vulnerable, a good analogy would be building a skyscraper on a foundation of quicksand. It is not necessary to be a structural engineer to outline how that works or not!

The other issue that tends to affect security management in virtualized environments is ***server portability***. It is a common practice for virtual machines to be moved from one host server to another. This allows organizations to group virtual machines on host servers in a way that makes the most sense from a performance standpoint.

This is important because virtual machine security works on multiple levels. The virtual machine itself must obviously be secured, but so, too, must the host operating system. As you can see, virtualization tends to complicate the subject of securing your servers. As long as you adhere to the various industry best practices for security, though, and are diligent about keeping your security up to date and consistent across the organization, virtualization should not cause any security issues [17].

Virtualization storage: the option for disaster recovery

Disaster recovery (DR) becomes more and more vital aspect of SME computing. In the past, it has been expensive to get one server to replicate to the other because those two servers had to be identical, so we needed the same hardware in both the main and backup locations. With virtualization technology, the hardware costs are cut down significantly, since the ability to host several machines on one server [6].

Disaster recovery plan (DRP) - sometimes referred to as a ***business continuity plan (BCP***) or business process contingency plan (BPCP) - describes how an organization is to deal with potential disasters. Just as a disaster is an event that makes impossible the continuation of normal functions, a disaster recovery plan consists of the precautions taken so that the effects of a disaster will be minimized and the organization will be able to either maintain or quickly resume mission-critical functions.

Disaster recovery planning always involves an analysis of business processes and an investigation of continuity needs; it may also include specifications about disaster preventions.

Business continuity planning refers to any methodology used by an organization to create a plan for how the organization will recover from an interruption or complete disruption of normal operations. International Organization for Standardization (ISO) / International Electrotechnical Commission (IEC) 17799:2000 Information Technology is the Code of Practice for Information Security Management, an international version of the British Standards 7799-1:1999, published in December 2000 [14]. It contains major sections, one of which is Business Continuity Management (Section 11) [15]. ISO/IEC Technical Report (TR) 13335 – Guidelines for Management of IT Security, 13335-2: Managing and Planning IT Security contains requirements for procedural security, including business continuity [15, 13].

Business continuity and disaster recovery are closely related concepts that often exist as a point of contention between information technology (IT) and business management. Unlike disaster recovery, which focuses on IT infrastructure, business continuity represents the processes and procedures that an organization puts in place to ensure that essential functions can continue during and after a degradation or complete loss of critical people, processes or technology.

One big similarity between business continuity and disaster recovery planning is how quickly the plans become obsolete. Changes in core technology, such as server platform, are major indicators that a disaster recovery plan needs to be revisited. Minor changes can quickly insert as well that will put the disaster recovery plan out of alignment with organizational needs.

The ***advantages*** of using virtualization as opposed to going with a more traditional DR solution are [12]:

- One advantage to using virtual servers instead of physical servers is in the testing. When we build the test environment, we basically have two virtual servers running on the same hardware.

- Using virtual servers, it is easy to test because we can shut down and turn on these virtual servers easily. We can simulate server failure. We don't need remote hardware access to power down.

Disaster recovery (DR) is one of the most significant and multifaceted aspects of IT management. The responsibilities drop out IT into other areas of the business. Midmarket CIO's think how to do it in terms of manpower and money. DR planning is often based on well-known technologies such as rudimentary backup software and services. The migration through virtualization technology takes time. For the reason of reduced hardware and software costs, SMBs (not only enterprises) should request service-based disaster recovery.

Disaster recovery plan is a plan for business continuity in the event of a disaster that destroys part or all of a business's resources, including IT equipment, data records and the physical space of an organization. The goal of a DRP is to resume normal computing capabilities in as little time as possible.

The design of any disaster recovery system should be driven by the ability to make available to the business the critical systems and information systems required to conduct normal production activities, without making those systems and information available to the wrong people [13].

Virtualization is changing DR; it changed the way that data center administrators look at fundamental hardware and software considerations. Central issues are disruption, hardware dependencies and cost implications of both. First, migrating copies of workloads between physical hardware is disruptive, limits application and database availability and can potentially result in lost productivity or sales.

Companies typically absorbed the cost of duplicate hardware and upgrade expenses. Some organizations have tried to moderate this cost by omitting non-critical workloads from the DR plan, relying instead on the availability of common backups for later restoration.

Server virtualization also enables hardware consolidation, allowing multiple virtual instances to operate on a properly configured server.

A company can use different servers in the DR site and can manage with fewer physical servers if each one is running multiple virtual instances. [19]

Conclusions

Virtualization is transforming every part of data center operations management.

Server virtualization can be viewed as part of an overall virtualization trend in enterprise IT that includes storage virtualization, network virtualization, and workload management. This trend is one component in the development of autonomic computing, in which the server environment will be able to manage itself.

But equally important, these technologies can also lower costs both directly, through an immediate reduction in power and cooling costs, and indirectly (but not with a lesser impact), through a reduction in IT administrative costs associated with server hardware and the layers of infrastructure software management.

Users could implement virtualization with software applications or by using hardware and software hybrid appliances. The technology can be placed on different levels of a storage area network. Data can easily be transferred and migrated between storage resources. Virtualization permits storage resources to be altered and updated on the fly without disrupting application performance. When properly implemented, storage virtualization eliminates forgotten or partially used disks, allowing superior storage utilization.

Vendors have been selling storage virtualization tools for a while now, but server virtualization techniques have largely centered around installing thin client/fat server products in which all applications run on the server. Who wants to work on a thin client these days? Most SMBs have their employees doing a range of tasks that don’t fall into any one neat category. You have to be careful about which applications you run, check their stabilities and run tests.

No matter what we implement, security is at the forefront of each project. Virtualizing servers is no guarantee of reliability. You still need to make sure that your data is adequately protected.

“Everything starts with the business!” said W. Preston [12]. A plan that protects company data and applications can stand for the difference between staying in business and going bust.

As for plans when the economy picks up, several network administrators (39%) surveyed indicated that endpoint virtualization could increase productivity. Another 27% said they believed that endpoint virtualization could decrease complexity, while nearly 40% thought the technology might decrease costs. According to the Symantec survey, nearly three-quarters of the network administrators polled are "at least considering plans to implement endpoint virtualization." [18]

Storage virtualization can be included in a virtual DR environment, but it's not a requirement for DR.

References

1. NOVELL, Virtualization in Data Center, 2006 www.novell.com/collaboration/
2. CHAD MARSHALL, Creating Business Continuity through Enterprise Storage Solutions, Realtime publishers, 2007, chap.1
3. Intel & VMWare, Virtualization Basics for SMBs, *TechTarget Data Center Media*, 2007
4. ANIL DESAI, A look at different approaches to virtualization, *Storage Magazine*, July, 2006
5. HERMAN MEHLING, Virtual desktops: Cheap and effective, *CIO Magazine*, June, 2007
6. DON JONES, Selecting the Right Virtualization Solution, Realtime publishers sponsored by SWSoft, 2007, Chapter 1
7. JOAN GOODCHILD, IT Consultant discusses virtualization for disaster recovery, *Storage Magazine*, August, 2006
8. MARK SCOTT, PC Restoration and Disaster Recovery, *Realtime publishers*, sponsored by Attachmate, 2007, chapter 4: Managing the Disaster Recovery Process
9. National Institute of Standards and Technology, Special Publication 800-34: Contingency Planning Guide for Information Technology Systems, June 2002, Section 2.2
10. NIST, Special Publication 800-12: An Introduction to Computer Security: The NIST handbook, October 1995, Chapter 11
11. Texas Department of Information Resources, Business Continuity Planning Guidelines
12. W. CURTIS PRESTON, Backup and Recovery, O'Reilly Media, 2007, chapter 24: It's all about Data Protection
13. JENNIFER MEARS, *The rush is on to virtualized servers*, *Network World*, January, 2005
14. MIKE ROTHMAN, *Preparing for security unknowns* - article in Virtualization security, sponsored by SOURCE *fire* and Tripware, 2009, pag.24
15. JED SCARAMELLA, *Next generation Technology Virtual I/O and Blade Servers*, Sponsored by Hewlett Packard, November 2008, 5272_215119 New VC Paper 111908, pag.1-2
16. DENNIS FISHER, *Virtual threats* – article in Virtualization security, sponsored by SOURCE *fire* and Tripware, 2009, pag.4-5
17. POSEY M. BRIEN, *Virtualization Offers No Escape from Security Concerns*, CIO Decisions EZINE, February/March 2009, volume 5, pag.13
18. http://www.networkworld.com/newsletters/nsm/2009/032309nsm1.html?nlhtnsm=ts_032309&nladname=032309networksystemsmanagemental
19. STEPHEN J. BIGELOW, *Virtualization concepts in disaster recovery*, Chapter 1, pag.3-4, Introduction to Virtualization E-book, TechTarget, 2009