# VIRTUALIZATION TECHNOLOGY FOR SMB

**Adriana BARNOSCHI**[*]

**Abstract**

*Data is growing at an alarming rate. Storage administrators are struggling to handle a spiraling volume of documents, images, audio and video files and large emails. 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. I give some ideas that guide you to protect yourself against disaster, what to back up and how backups should work.*

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

## Introduction

The earthquakes, 9/11 event, summer storms and California fires have drawn attention to the importance of protecting data and applications. Most companies rely on computing system as their business infrastructure. 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.

The first question was: “Where shall I start from?” I found the answer: disaster recovery (DR) and business continuity plan (BCP). When I paid attention to the disaster planning, I attempt to find out what choices of storage system do I have (both conventional and new solutions) and if the researches are viable options for SMBs with limited resources and budgets.

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, a few suggestions for SMEs intended for successful virtualization implementations and the benefits of desktop virtualization for SMEs, since this technology has become so popular in storage management.

[*] Associate Professor, Ph.D., Social and Administrative Sciences Faculty, “Nicolae Titulescu” University.

Section 5 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 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.

According with the chosen virtualization solution from section 5 and going from the idea that "everything starts with the business" [16], I proposed a disaster recovery plan that protects business data and applications of company. There are simple steps SMBs can take to prepare for the worst [15]. The business sets the IT budget and therefore the RTO and RPO metrics need to fit with the available budget.

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

## Literature review of virtualization

Virtualization is almost as old as enterprise computing itself. First introduced in the 1960s to allow partitioning of mainframe hardware, it has been a foundation of high-end proprietary server environments ever since. Today, virtualization is once again a hot topic of conversation in the data center because emerging technologies have the potential to remedy issues relating to resource utilization, efficiency, scalability and manageability [1].

The idea of virtualization in computing systems is to add a layer of abstraction between two layers in that computer system. 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:

▪ *By virtualizing a software application*, we eliminate its direct hooks into its host operating system (OS) allowing to more easily install, remove and modify that software installation without affecting the host OS.

▪ *By virtualizing an entire computer system*, we encapsulate its configuration into a data structure that is more portable, easier to manage and has more capability for being backed up and restored [2].

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 [3].

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 [4].

Server virtualization provides a path toward server consolidation that results in significant power and space savings, while also offering high availability and system portability. Today, vendors are building hardware and software platforms that can deliver virtualization solutions at near-native performance.

Products from Microsoft and VMware lead in domain. VMware Infrastructure 3 running on Intel-based platform creates high performance virtualized environments that meet today's requirements, giving customers confidence, reliability and security [5].

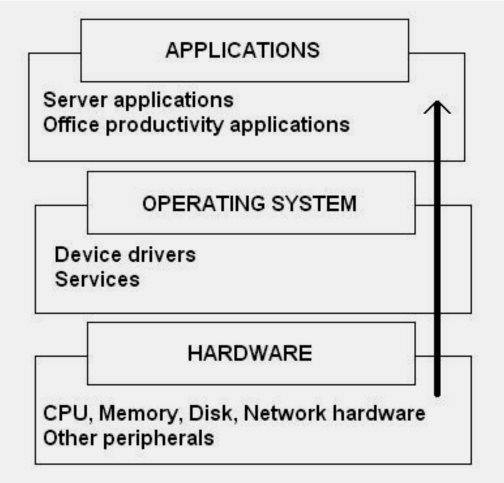

**Figure 1:** Standard server stack

## 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. Figure 1 provides a high level overview of the areas of a standard server stack that can be virtualized.

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].

## Successful virtualization implementation

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 underlying hardware system as possible,

3. To do all of this while minimizing performance overhead. That's no small set of goals, but it can be done (and in more ways than one).

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

Natalie Lambert, a senior analyst at Forrester Research Inc. in Cambridge, Mass, makes a few ***recommendations*** for SMBs for a successful virtualization implementation:

- Break down users into groups based on mobility, resource requirements and sensitive data requirements.
- Conduct a pilot program with specific user groups, such as contractors using unmanaged PCs.
- Consolidate and standardize machines to support the desktop virtualization effort.
- Conduct pilot projects running problematic applications.

To get the most out of 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 offer mobile and contact workers with secure, corporate-approved desktops, 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 ***benefit*** 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."

There are other performance issues with server virtualization that will be addressed using optimized hardware chipsets, such as Intel Corp's vPro Processor Technology and Q35 Express Chipset [4].

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 may do very little to stem the tide of complexity and inefficiency that has overwhelmed even administrators' best plans.

## Virtualization in DR

***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 plan

*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].

"Everything starts with the business!" said W. Preston [15].

***a. Define the Core Competency of the Organization***

It means understanding an organization's activities and how all of its resources are interconnected. Answer at the first question: "What are the core products and/or services that the organization offers?"

The second question is: "What is the required information that provides that product or service, and what applications are required to effectively use this information?" The answer to the second question defines your organization's intellectual property (IP).

***b. Prioritize the Business Functions Necessary to Continue the Core Competency***

It is important to review an organization's vulnerability in all areas, including operating procedures, physical space and equipment, data integrity and emergency planning.

***c. Correlate Each System to a Business Function, and Prioritize***

A great example of this type of prioritization can be found in a publication of the U.S. Federal Communications Commission. It shows the FCC's different types of data and its criticality, and it is published at *http://www.fcc.gov/webinventory/*.

***d. Define RPO and RTO for Each Critical System***

The Recovery Time Objective, or RTO, means the time you want the system to be recovered. RTOs can range from zero seconds to many days, or even weeks. Each application serves a business function, so the question is how long you can live without that function. The Recovery Point Objective, or RPO, defines the point in time that is reflected once you have recovered a system, also referred to as how much data you can afford to lose.

You must create an RTO and RPO for each *protected* system. You should also know what your *budget* is, how much data needs to be backed up, how much data changes and what your RTO and RPO requirements are.

***e. Create Consistency Groups***

It is often necessary to recover several systems to the same point in time. This is first and foremost caused by applications that pass data to one another. If your company has several systems that perform related business processes, those systems need to be in the same consistency group. In addition to determining an RTO and RPO, you must identify those systems that are related to each other because they need to be recovered to the same point in time [13].

***f. Determine for Each Critical System What to Protect from***

After you have made a list with prioritized business functions and you have assigned each system to a business function, you should identify the things that can happen, that trigger a recovery scenario. You'll make a list of levels and types of disasters on each level that are expected for your area and type of business. The *Disaster Recovery Institute* states that each company should define its own levels of disasters.

***g. Determine the Costs of an Outage***

Once you have created the list from above, establishing all types of disasters and their associated probability, you must assign a cost to each type of disaster, for each type of system.

***h. Plan for all types of disasters***

Do not allow to go any particular type of disaster! Thinking "that will never happen". Murphy's Law will find you. The disaster you do not prepare for, is the one that will strike you!

An ***example*** of planning a disaster recovery solution is given by Michael Osterman [19] beginning from the answers to several important questions: How much e-mail data loss is acceptable following a disaster? How much time between the beginning of an e-mail service loss and e-mail recovery is acceptable? What quantifiable or other benefits would there be in speeding up the recovery process?

The ***best plan*** in the world is not worth much, if people aren't available to implement it!

## 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.

In general, as you move up, from hardware- to server- to application-level virtualization, you gain scalability at the cost of overall independence. Many companies are seeing the benefits of virtualization software because they can reduce their capital expenditures.

Server virtualization can be used to eliminate server sprawl, to make more efficient use of server resources, to improve server availability, to assist in disaster recovery, testing and development, and to centralize server administration.

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.

Interest in virtual machine technology has been growing. IDC says the market reached more than $300 million in 2004 and is on pace to grow at a rate of about 18% over the next few years. "It's been one of the faster growing technologies that we've encountered," says Galen Schreck, a senior analyst at Forrester Research [16]

According to the Disaster Recovery site: "Despite the number of very public disasters since 9/11, only about 50 percent of companies reports having a disaster recovery plan. Of those that do, nearly half have never tested their plan, which is equivalent to not having one at all."

## References

1. NOVELL, Virtualization in Data Center, 2006 www.novell.com/collaboration/
2. http://searchservervirtualization.techtarget.com/ Definitions, *Powered by Whatis.com*
3. CHAD MARSHALL, Creating Business Continuity through Enterprise Storage Solutions*,* Realtime publisher*s,* 2007, chap.1
4. Intel & VMWare, Virtualization Basics for SMBs, *TechTarget Data Center Media*, 2007
5. ANIL DESAI, A look at different approaches to virtualization, *Storage Magazine*, July, 2006
6. www.virtualization.info/glossary
7. www.searchwmware.techtarget.com
8. HERMAN MEHLING, Virtual desktops: Cheap and effective, *CIO Magazine*, June, 2007
9. DON JONES, Selecting the Right Virtualization Solution, Realtime publishers sponsored by SWSoft, 2007, Chapter 1
10. JOAN GOODCHILD, IT Consultant discusses virtualization for disaster recovery, *Storage Magazine*, August, 2006
11. MARK SCOTT, PC Restoration and Disaster Recovery, *Realtime publishers,* sponsored by Attachmate, 2007, chapter 4: Managing the Disaster Recovery Process
12. National Institute of Standards and Technology, Special Publication 800-34: Contingency Planning Guide for Information Technology Systems, June 2002, Section 2.2
13. NIST, Special Publication 800-12: An Introduction to Computer Security: The NIST handbook, October 1995, Chapter 11
14. Texas Department of Information Resources, Business Continuity Planning Guidelines
15. W. CURTIS PRESTON, Backup and Recovery, O'Reilly Media, 2007, chapter 24: It's all about Data Protection
16. MICHAEL OSTERMAN, Recovering from disasters, *Network World's Unified Communications Newsletter*, November, 2007
17. JENNIFER MEARS, *The rush is on to virtualized servers*, *Network World*, January, 2005