

DATA CENTER VIRTUALIZATION AND ITS ECONOMIC IMPLICATIONS FOR THE COMPANIES

Senior lect. Ph.D. Logica BĂNICĂ, Professor Ph.D. Mariana JURIAN,
Assistant Cristian ȘTEFAN

University of Pitesti
olga_banica@yahoo.com

Abstract: *In the current situation of the economic crisis, when companies target budget cuttings in a context of an explosive data growth, the IT community must evaluate potential technology developments not only on their technical advantages, but on their economic effects as well. More than ever, the old cliché “doing more things with fewer resources” is true today. Many IT companies started building very large facilities, called data centers (DCs) or Internet DC (IDCs), which provide businesses a wide range of solutions for systems deployment and operation. In recent years, the IT departments around the world have moved from data center and infrastructure consolidation to virtualization.*

Data center virtualization is the process of aligning available resources with the actual needs of the offered services, moving from physical servers to virtual servers, sharing and provisioning servers, networks, storage, and applications. By taking advantage of three basic innovations — virtualization, tiered storage architectures and dynamic provisioning software — an organization can achieve greater efficiencies in their current computing environment.

Such a unified computing architecture offers end-to-end virtualization; all structures are optimized for virtualized environments, from the CPU to the aggregation layer. In combination with embedded management, this new approach increases responsiveness and reduces the opportunities for human error, improving consistency and reducing server and network deployment times.

Key words: data centre; virtualization; tiered storage; provisioning software; unified computing.

JEL codes: C91; C88

1. INTRODUCTION

In the current situation of the economic crisis, when companies target budget cuttings in a context of an explosive data growth, the IT community must evaluate potential technology developments not only on their technical advantages, but on their economic effects as well. More than ever, the old cliché “doing more things with fewer resources” is true today. So, the increased global competition, shrinking budgets and rising energy costs are driving up virtually every aspect of doing business. Dramatically increasing amounts of data are produced by the business, the rising management costs of complex storage networks, and the pressure to hold or reduce IT costs leaves no place for waste or error.

IT operations are an important aspect of most organizational operations. One of the main goals is **business continuity**: companies rely on their information systems to run their operations continuous. If a system becomes unavailable, the company’s services may be disturbed or stopped completely. It is necessary to provide a reliable infrastructure for IT operations, in order to minimize any chance of discontinuity. Today, the majority of IT budgets are spent just for maintenance and management of the existing infrastructure. As a result, IT departments must

continually invest in more people, time, and other resources to maintain a growing, inflexible infrastructure instead of rapidly and effectively responding to business needs.

Now, during the global economic crisis, IT departments struggle to be more responsive to the business at a lower cost. Power, cooling, real estate, and the ever-increasing demand for computing, storage, and application resources all continue to push costs higher.

Companies needed fast Internet connectivity and nonstop availability to establish a serious presence on the market. Installing IT&C equipment for this purpose was not financially acceptable and viable for many smaller organizations. So, many important service providers started building very large facilities, called *data centers (DCs)* or *Internet DC (IDCs)*, which provide businesses with a range of solutions for systems deployment and operation. Information security is also a concern, and for this reason a data center has to offer a secure environment which minimizes the chances of a security break. A data center must keep high standards for assuring the integrity and functionality of its hosted computer environment.

In recent years, the IT departments around the world have moved from *data center* and infrastructure consolidation to *virtualization*. In 2008, 38 percent of all IT workloads were virtualized. This trend is continuing as the majority of IT executives are planning to virtualize more than 50 percent of their computing infrastructure within the next two years. In 2008, the average enterprise deploying virtualization could expect a return on investment (ROI) of more than 400 percent and payback within 12 months, largely as the result of server consolidation, increased utilization, and higher availability.

However, customers are facing challenges in achieving the full benefits of virtualization. Despite the widespread adoption of virtualization, data center operating costs are still at a high level. Server capacity and performance have increased exponentially faster than data center cooling capacity. This disparity accelerates the obsolescence of data center infrastructure. According to Forrester (December 2008), more than 70 percent of today's average IT budget is spent maintaining today's inflexible data center equipment. Data center owners are challenged with unifying and scaling separate virtualized environments.

The history of enterprise computing has been marked by compromises between scale and simplicity. As the number of data center systems has scaled, they have also increased in complexity, and as complexity has increased, so has the expense of deployment and ongoing management [5]. They are struggling to:

- Maintain consistent and coherent infrastructure policies in this new virtualized environment;
- Adapt current management tools and methods to the changing and increasingly interdependent roles of data center administrators;
- Simplify operations when their infrastructure is inflexible, underutilized, and lacking cohesive integration.

Today's concept is data center virtualization and will be our target likewise the distinct approaches of that by two important companies: Cisco, Hitachi and VMware.

2. MEETING VIRTUALIZATION WITH DATA CENTER

It may be helpful to give an appropriate definition for this concept: a **data center** is a facility used to house computer systems and associated components, such as telecommunications and storage systems. It generally includes redundant or backup power supplies, redundant data communications connections, environmental controls and security devices [4].

A data center must therefore keep high standards for assuring the integrity and functionality of its hosted computer environment.

The TIA-942: Data Center Standards Overview [6] describes the requirements for the data center infrastructure:

- *Physical layout* – the equipment such as mainframe computers and storage devices;

- **Network infrastructure** - communications in data centers are most often based on networks running the IP protocol suite; data centers contain a set of routers and switches that transport traffic between the servers and to the outside world; some of the servers at the data center are used for running the basic Internet and Intranet services needed by internal users in the organization: e-mail servers, proxy servers, and DNS servers.
Network security elements are also usually deployed: firewalls, VPN gateways, intrusion detection systems, etc. Also common are monitoring systems for the network and some of the applications. Additional off site monitoring systems are also typical, in case of a failure of communications inside the data center.
- **Applications** - the main purpose of a data center is running the applications that handle the core business and operational data of the organization. Such systems may be proprietary and developed internally by the organization, or bought from enterprise software vendors. Such common applications are Enterprise Resource Planning (ERP) and Customer Relationship Management (CRM) systems.

Often these applications will be composed of multiple hosts, each running a single component. Common components of such applications are databases, file servers, application servers, middleware, and various others. Data centers are also used for off site backups. Companies may subscribe to backup services provided by a data center. Larger companies may also send their backups off site for added security. Encrypted backups can be sent over the Internet to another data center where they can be stored securely. For disaster recovery, several large hardware vendors have developed mobile solutions that can be installed and made operational in very short time. Vendors such as Cisco Systems, Sun Microsystems, and IBM have developed systems that could be used for this purpose [2].




Figure 1 Data center architecture

Virtualization has created a market transition in which IT departments are trying to reduce costs and increase flexibility. Despite this, IT organizations are constantly working against existing rigid, inflexible hardware platforms. As a consequence, data center administrators have had to spend significant time on manual procedures for basic tasks instead of focusing on more strategic, proactive initiatives.

Virtualization allows **multiple virtual machines**, with **heterogeneous operating systems** to run independently, on the same physical machine. Each virtual machine has its own set of virtual hardware (RAM, CPU, NIC, Disk) upon which an operating system and applications are loaded. The operating system sees a consistent, normalized set of hardware regardless of the actual physical hardware components [1].

An important advantage of virtual machines is that they are encapsulated into files, making it possible to rapidly save, copy and provision a virtual machine. Full systems (fully configured applications, operating systems, BIOS and virtual hardware) can be moved from one physical server to another for zero-downtime maintenance and continuous workload consolidation. They can even be sent over the network in real time, without stopping the running system.

Some of the advantages of Virtualization are [7]:

1. Partitioning hardware and software – meaning that multiple applications and operating systems supported on a single physical system, computing resources are treated as a unified pool from which they are allocated to virtual machines in a controlled manner
2. The isolation between host machine and virtual machines, the functional independence of virtual machines and their applications, that can only communicate over network connections
3. The possibility to save as a single file of the complete virtual machine environment (named also Encapsulation).

This blend of virtualization technologies - or **virtual infrastructure** provides a layer of abstraction between computing, storage and networking hardware, and the applications running on it. Virtual infrastructure gives administrators the advantage of managing pooled resources across the enterprise, allowing IT managers to be more responsive to dynamic organizational needs and to better leverage infrastructure.

The PC's functional structure after the virtualization gives us important advantages:

- Virtualization offers a hardware independence of Operating Systems and applications;
- Virtualization offers high availability and error recovery solutions by encapsulating entire systems into single files that can be replicated and restored on any destination server.
- Virtualization offers the ability to run multiple operating systems on a single physical system and share the underlying hardware resources;
- The virtualization solves the problems of insecurity by layering a security policy in software around desktop virtual machines.

The virtualization has two aspects typically used with software-based partitioning: hosted and hypervisor architectures [3].

A **hosted approach** provides partitioning services on top of the operating system and supports the broadest range of hardware configurations. The main characteristic of this virtualization solution is the capability to run as an application, based on host Operating System for device support and physical resource management. There are two kinds of applications on the same machine: depends on host OS and based on Virtualization Layer.

A **hypervisor architecture** (Figure 2) considers the virtualization as the first layer of software installed on a clean x86-based system (often referred to as a *bare metal* approach). Since it has direct access to the hardware resources, a hypervisor is more efficient than hosted architectures, enabling greater scalability, robustness and performance. This approach has a virtualization central kernel (Virtualization Layer) and a Service Console for applications.




Figure 2. Hypervisor virtualization architecture [7]

Virtualization technology enables you to create a set of logical resources that share underlying physical resources. A group of resources, such as servers, applications, databases, and networks, behave as a single resource from which all services can draw. It requires application components, access software, virtualized operating systems, storage, and an overlying infrastructure layer to make it appear that all components are equally available. The goals are high availability, load balancing, increased utilization, improved scalability, flexible capacity on demand, and simplified systems management.

Data center virtualization is the process of aligning resources with the needs of services. In the earliest stages of virtualization, data centers typically move from *physical servers* to *virtual servers*. In support of maturing IT processes, data center virtualization extends to *sharing* and *provisioning servers, networks, storage, and applications* to maintain application service levels through policy-based management.

Data center virtualization enables more efficient sharing of resources. It has been a proven solution for IT management in mainframe environments for more than 20 years, where the need for reliability and efficient resource utilization has been crucial. Now the distributed systems industry is searching for the same level of efficiency that data center virtualization affords to mainframe IT managers.

By taking advantage of two basic innovations — **virtualization** and **tiered storage** architectures — an organization can realize *greater efficiencies* in their *current storage environment* and discover a way for easier data migrations in the future. In other words, they can reduce cost and complexity while increasing quality of service. Beyond these organizational benefits, they are economically superior architectures that can reduce both the initial cost of acquisition as well as ongoing costs, leading to a lower total cost of ownership and a better return on investment.

Flexible data center automation is achieved by unifying computing, network, storage access, and virtualization within a single, highly available, cohesive system, decreasing costs and increasing flexibility.

As a result of the virtualization revolution of the past few years, new standards, technologies, and integration conventions are emerging. These innovations enable the design of a pre-integrated data center solution from industry-standard components, rather than requiring IT staff or consultants to integrate their own network, computing, and virtualization platforms to create a data center solution.

3. EXAMPLES OF DATA CENTER VIRTUALIZATION

There are many hardware and software companies that offer to the clients the possibilities of virtualization their data center and two of them are: Cisco and Hitachi.

Data center 3.0, the latest iteration of **Cisco** is based on a shared, virtualized infrastructure model: services abstracted from the underlying physical assets that can be dynamically adapted to changing application and service level requirements. This architecture enables the emerging enterprise and service provider "cloud" models. This transformation is more heavily reliant on the network than ever before. In many ways *the network* now becomes ***the underlying foundation*** upon which the virtualized data center infrastructure runs.

Data Center 3.0 is Cisco's strategy and vision for facilitating data center transformation to this new virtualized model by using the network as the underlying platform. It combines an innovative architecture with an integrated technology roadmap, design best practices, and services, all designed to help customers transform their data center infrastructures, processes, and organizations with maximum effectiveness and minimum risk, cost, or disruption.

The Cisco Data Center 3.0 technology portfolio includes the most comprehensive range of networking products and solutions designed specifically to help IT organizations transform their data center infrastructures:

- **Data center network infrastructure:** Network infrastructure platforms that provide the physical and logical connectivity and access to all physical and virtualized data center resources (compute and storage). This infrastructure layer supports all heterogeneous servers and storage, including standalone servers, blade servers and virtual machines.
- **Data center infrastructure services:** These networked services reside on the infrastructure platforms to help ensure scalable protection, manageability, and reliability for the applications and information hosted in the data center. These services include virtualization, storage, security and application network services.
- **Application network services:** Cisco Application Networking Services (ANS) increase application performance and availability, enable the consolidation of expensive branch office and data center infrastructure, and accelerate new application deployments. These services include data center application services and wide area application services.
- **Data center automation and provisioning tools:** As data center technologies become increasingly sophisticated, Cisco is developing innovative tools and working with partners to help reduce the complexity of deploying and managing a virtualized environment.

This market transition, which Cisco is embracing as unified computing, marks the natural evolution of the data center to link data center resources within a cohesive management framework, bringing infrastructure policy consistency and transparency, traditionally delivered by the network, to the virtual machine level. In essence, the network is becoming the foundational infrastructure of the data center.

This architectural shift will have a profound, positive effect on traditional data center administrator roles. Unifying and embedding management at the system level creates the opportunity to increase flexibility and incorporate multiple administrator roles within the same system. With embedded role-based management, organizations can make more efficient use of their administrator resources.

Another interesting approach is given by **Hitachi Company**, which demonstrates that virtualization technology, tiered storage and provisioning software can all deliver significant economic benefits to the organization by helping to increase storage utilization and overall efficiency while reducing costs and complexity. By decreasing capital expenses as well as operational expenses, these elements represent an economically superior solution, a sustainable approach to run multiple applications and operating systems on a single physical system, to

improve the performance and efficiency of IT infrastructure, high availability and multiple storage solutions for today and for the future [2].

Because of the way storage is packaged, businesses must buy up to 75 percent more capacity than they actually need. It's cheaper to simply buy more storage than to hire somebody to manage it, so an organization's first response to a storage crisis will be to throw more storage at the problem. Of course, the many different business units in a given organization may not want to share storage, so they each purchase it piecemeal through a variety of different vendors with an eye to the lowest cost option. The result is a confusing tangle of heterogeneous storage systems and management software, which requires more IT staff expertise, and massively underutilized storage assets and "stranded storage," which are a waste of already shrinking budgets. Ultimately, coping with exploding data requirements by simply purchasing more storage can lead to very costly consequences.

By taking advantage of three basic innovations — *virtualization, tiered storage architectures and Hitachi Dynamic Provisioning software* — an organization can realize greater efficiencies in their current storage environment. In other words, they can reduce cost and complexity while increasing quality of service. Beyond these organizational benefits, tiered storage, virtualization and dynamic provisioning are economically superior architectures that can reduce both the initial cost of acquisition as well as ongoing costs, leading to a lower total cost of ownership (TCO) and a better return on investment (RoI).




Figure 3 Data Center Virtualization architecture

Another competitor on the datacenter virtualization market is *VMware, Inc.*, a company with tradition and long-time experience in the field of virtualized environments. Their flagship product, *Virtual Datacenter OS*, brings the paradigm of cloud computing one step closer, by using transparently all available physical resources, independent of their location, and offering integrated management through *vCenter*, their solution to reduce costs associated with managing distributed environments. Applications run wherever there are resources available, and services are offered on a full time basis wherever the clients are demanding them, without the hassle of redirection and downtimes. Near 100% availability can be achieved at substantially lower operating costs, and all industry-used technologies can be deployed instantly, without the OS-related incompatibility issues, as the VDOS can support anything from Windows to Linux.

4. CONCLUSIONS

This new approach of data center *allows organizations to increase effectiveness* by maintaining IT disciplines and accountability while increasing teamwork and collaboration. Server, network, and storage administrators can preserve accountability for their domain policies while interoperating within a single integrated management environment. As data center roles and policies evolve, individual responsibilities and system privileges can be easily modified and new roles quickly created.

Unified computing will accelerate adoption of virtualization and magnify the returns. Integration of virtualization, network, and compute resources creates the opportunity to provide the same level of network visibility for virtualized environments that is expected and required for physical servers. Data center virtualization integrates policy-based virtual machine connectivity, mobility of network properties with the virtual machine, and a consistent operational model for both physical and virtual environments. This approach simplifies network and security policy enforcement when virtual machines are migrated from one host to another.

A unified computing architecture offer *end-to-end virtualization*; all structures are optimized for virtualized environments, from the processor to the aggregation layer. The latest industry-standard processor technology enables better virtualization performance, scalability, and enhanced flexibility. In the server, by balancing CPU and I/O capabilities while increasing memory capacity, more virtual machines can be hosted per server than ever before.

By providing visibility and portability to network polices and security all the way to the virtual machine, a consistent operational model can be implemented between the physical and virtual environments. In combination with embedded management and the service profile construct, this automation *increases responsiveness* and *reduces the opportunities for human error*, improving consistency and reducing server and network deployment times.

A unified computing architecture based on virtualization helps companies:

- a) Reduce total cost of ownership (TCO) at the platform, site, and organizational levels;
- b) Increase data center staff productivity by accelerating infrastructure provisioning for both virtualized and non-virtualized environments;
- c) Enable scalability without the increase in complexity typical of older environments, with hundreds of servers and thousands of virtual machines now operating in a single highly available management domain;
- d) Increase data center energy efficiency and resource utilization;
- e) Reduce the interventions of management; with disparate technologies integrated into one cohesive system, management is consolidated and embedded, enabling infrastructure policies to be consistently applied across the system.

REFERENCES

1. L. Bănică, D. Roșca, C. Ștefan, *Virtualization: an Old Concept in a New Approach*, Informatica Economica Journal (EI no.7), no. 1/2007, Academy of Economic Studies, Bucharest
2. D. Merrill, V. Ramaswamy, *The economic effects of virtualization, tiered storage and Hitachi dynamic provisioning software*
3. <http://www.vmware.com>
4. Managing Data Center Virtualization, <http://www.bmc.com/bmceducation>
5. Unified Computing Overview, 2009 Cisco Systems, <http://www.cisco.com/go/unifiedcomputing/>
6. <http://www.adc.com/Library>
7. L.Bănică, F. Smaranda, C. Ștefan, *Virtualization platforms: a comparative approach*, Scientific Bulletin, Computer Sciences series, no.13 (2007), Pitesti