

Data centre virtualization at UTAD – A hybrid consolidation approach

Rio-Costa A. ¹, Reis A. ¹, Borges J.^{1.}, Vasconcelos A. ¹, Santos J ¹, Barroso J.², Bulas-Cruz J³

¹ Centro de Informática - Departamento de Engenharias, Universidade de Trás-os-Montes e Alto Douro, Quinta de Prados, 5001-801 Vila Real, Portugal, acosta@utad.pt, ars@utad.pt, jborges@utad.pt, albertov@utad.pt, jgsantos@utad.pt. ² Grupo de Investigação em Engenharia do Conhecimento e Apoio à Decisão, Instituto Superior de Engenharia do Porto, R. Dr. António Bernardino de Almeida, 431, 4200-072 Porto, Portugal, jbarroso@utad.pt. ³ Centro de Investigação e de Tecnologias Agro-Ambientais e Biológicas, Universidade de Trás-os-Montes e Alto Douro, Quinta de Prados, 5001-801 Vila Real, Portugal, jcruz@utad.pt.

Keywords

Virtualization, IT strategies, IT infrastructure, Vision IT, data centre consolidation.

1. EXECUTIVE SUMMARY

Data centres have been evolving and growing throughout the years. Examples of this growth and evolution include mainframe to distributed, single to clusters and now, organizations begin the task of server consolidation, physical to virtual.

Virtualization technology serves as a great enabler of data centre optimization, consolidation and migration. Virtualization, by offering the ability to host multiple operating systems on a single physical server, offers the ability for organizations to liberate data centre space. Therefore, saving, space, energy, hardware and IT resources.

During the past year the computer centre of the University of Trás-os-Montes e Alto Douro (UTAD) started their data centre consolidation. This consolidation was achieved by the development of a hybrid architecture that mixed both virtualized and physical infrastructures.

With this case study scenario this paper, intends to report on the set up architecture and his implementation. To give some considerations on the gains achieved by the reduction of space, energy and other factors and also, to pinpoint some strong and weak facts regarding to and when should we step to virtualization.

2. INTRODUCTION AND PROJECT SCOPE

In traditional data centres there is a tight relationship among the particular computers, particular disk drives, particular network ports, and the applications they support. Virtualization allows the breaking of those bonds. With virtualization we can dynamically move resources where they are needed and consequently we can move processing where it makes more sense.

Virtualization technology serves as a great enabler of data centre optimization, consolidation and migration. Virtualization, by offering the ability to host multiple operating systems on a single physical server, offers the ability for organizations to liberate data centre space. Therefore, saving, space, energy, hardware and IT resources.

Achieving the full value of virtualization requires work through integration points between people, process and technology. Sustainable, virtualization adoption involves a comprehensive approach that should reach out to all the areas of an organization.

We can look at virtualization as the use of software to enable multiple pieces of software to interact with hardware in a way that each piece of software “thinks” that has the complete control of the hardware. Virtualization refers to the abstraction of computing, storage and network devices. It allows detaching physical hardware from the operating system (OS), to deliver greater resource utilization and flexibility.

There are mainly four reasons that brought virtualization to the front centre of the computer environment:

- ✓ Underutilized hardware - The increase of computing power brought the ability to have unused capacity in smaller servers.
- ✓ Data centres space - The IT infrastructure growth and consequently, server’s growth caused that data centres started to run out of space. Building new data centres costs money.
- ✓ Energy costs - The cost of running full data centres at low utilization levels are really high.
- ✓ System administration - IT system administrators, need to do hardware monitoring, Operative OS patching, upgrading and backup each system case by case.

With virtualization it’s possible to reduce all of these factors in a way that data centres can be greatly optimized. Also, IT management becomes easier and less time consuming. But, traditional systems management no longer apply. With virtualization, we can get the potential to increase great economic advantages and efficiency, nevertheless IT management needs to face new challenges of managing application performance across multiple IT domains.

During the past year the computer centre of the University of Trás-os-Montes e Alto Douro (UTAD) started their data centre consolidation. This consolidation was achieved by the development of a hybrid architecture that mixed both virtualized and physical infrastructures at the data centre.

3. UTAD DATA CENTRE CONSOLIDATION

During the past year, the computer centre at UTAD noticed that there was a need for IT infrastructure consolidation at his data centre. The available space, the energy and power costs needed to cool and maintain the entire data centre, and also, some environmental issues started to come up. In the top of those problems, it came along the need for efficiency concerning to system administrators to maintain and to manage the server infrastructure. But this wasn’t yet the real problem. The real problem was to solve the question in how to increase the overall infrastructure, by providing more services with more quality and at the same time by reducing the infrastructure at data centre and management costs.

To fulfil these questions the computer centre, started to deploy a virtualized hybrid architecture in manner that it could get the most advantages possible of the existing hardware available. Making in this manner all the data centre consolidation efforts at the less cost possible.

The set up architecture was based by deploying servers with virtualized machines in a separated data storage group (OS Storage group) and, also with directly mapped physical data storages where the server's data lies (Data Storage Group), on the top of a physical hardware infrastructure. The next figure shows the developed architecture.




Figure 1 - Set up architecture

In the virtualized layer, we putted the virtual machines. Those can be Microsoft OS servers, Linux Servers or even virtual switches. This layer, is were the system administrator installs, and maintains all the aspects related to his server. The system administrator gets a set of standard drivers for the servers provided by the physical layer. The virtual machines are treated and maintained as physical.

In the physical layer, the physical servers were stacked and grouped with an abstraction software module that provides hardware emulation to the upper layer. This minimal virtualized software module is usually called hypervisor, which creates a virtualized machine by emulating the entire hardware environment. The operating system that is loaded into each virtual machine in the top layer is a standard unmodified product. As it makes system calls the hypervisor catches that system calls and redirects them to be manipulated by data structures provided by the hypervisor. Then the hypervisor itself makes the calls to the physical hardware underlying the entire software agglomeration.

In the storage layer, were created two storage groups, the “Os storage group” and the “data storage group”. The “Os storage group” is the place where the virtualized server image files are stored. Each data store in this group contains the virtual disk image of one virtualized server. The “data storage group” contains in each group the data store for each virtualized machine. This data is not accessed in a virtualized environment, but is directly mapped by each server independently. This implementation permits that if one virtualized server fails, is possible to access that particular data store from a non virtual environment.

The next image shows the first virtual machines connected to the physical resource in UTAD’s data centre.




Figure 2 - Virtual Servers connected to the resource pool

As shown in the next figure, virtual servers can share virtual switches inside the virtual infrastructure even if is only an single physical server

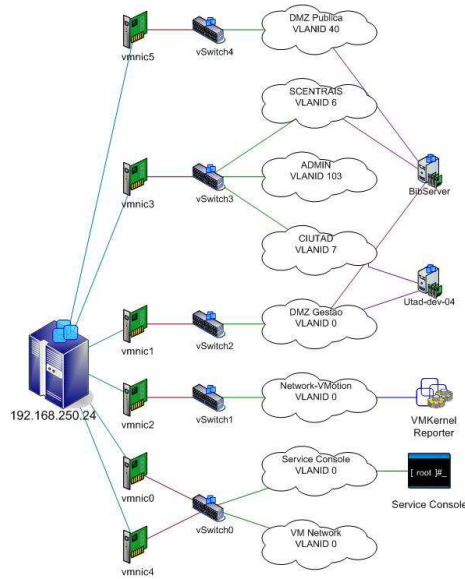


Figure 3 - Sharing virtual switches

The actual state of the data centre at UTAD is shown in the following figures.




Figure 4 - One virtual machine network and data storage relationship map

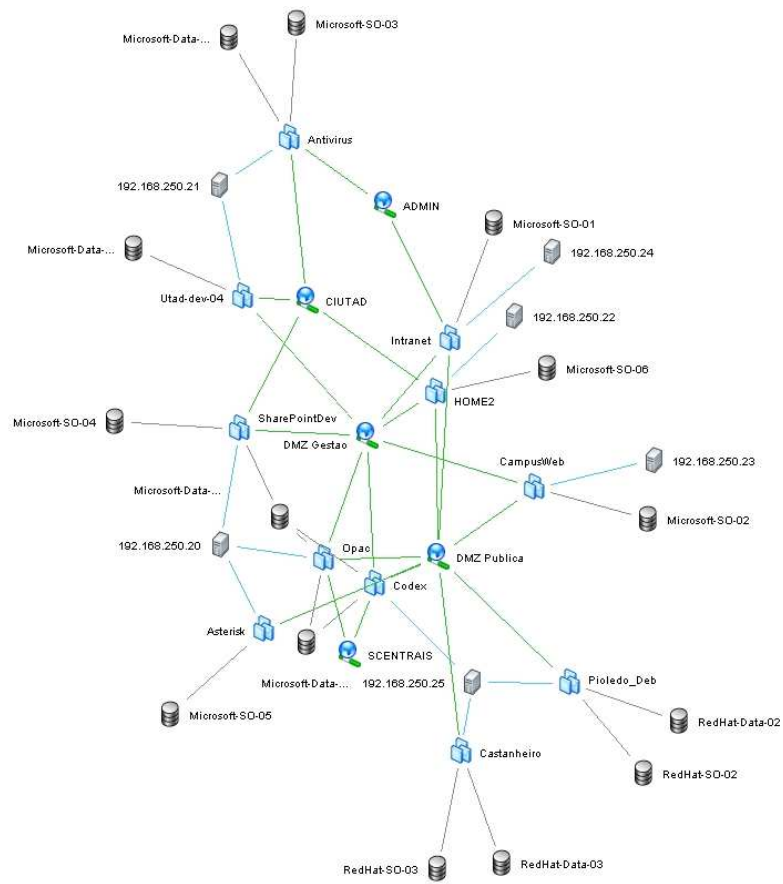


Figure 5 - Virtual data centre relationship map

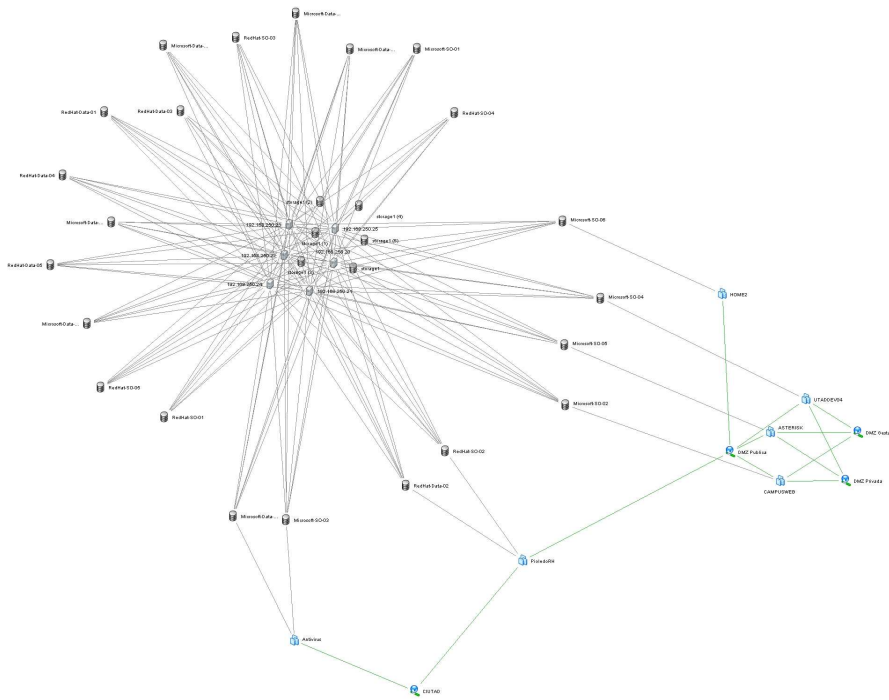


Figure 6 - Virtual data centre physical data store relationship map

This architecture was deployed using the following components of hardware and virtualization software:

- 6 HP Blade Servers in 1 single enclosure
- 1 EMC San Storage with 24TB available
- 36 Ethernet 1 Gbps ports
- Backup solution - 30 TB
- Virtualization software - VmWare Esx server
- Infrastructure management - VmWare Infrastructure

4. RESULTS AND CONSIDERATIONS

At present UTAD data centre has his virtual environment 18 virtualized production servers. In these servers we included the university web servers, the intranet servers, the e-learning server, the student's information support system servers, and also other servers that provide a various number of applications to the university. All of these virtualized servers are supported by 6 physical machines. Therefore all of energy costs by supporting cooling and maintaining 18 physical servers online where reduced to 1/3. The achievement of hardware maintenance, hardware monitoring, in preventive patching and upgrading concerns were strongly reduced. Also, the implemented model permitted to get better performances that if we had deployed the entire infrastructure in a virtual environment.

But to step to virtualization there are some concerns that must be taken in consideration when to go to virtualization.

The first is that virtualization isn't the cure for all IT illness. Virtualization infrastructure provides a stable environment, but in some cases there are hard or even impossible implementation scenarios. Therefore starting by small and not so demanding solutions can be a good experience. Be aware that choosing servers with great I/O (like email servers) or very particular hardware dependent servers (like IPpbxs) often don't work. To choose what to make virtual, should be made by looking carefully into the services that virtual servers need to provide and, also by the ability in improving the quality of that particular services.

Second is that virtualization infrastructure is dynamic, so you can move around virtual servers and users and even system administrator won't notice. Also this gives you the chance to constantly evaluate the infrastructure and improve by adding physical servers or only by making simple software adjustments. All of this can be done without any service disruption.

Third when you go to virtualization, the organization should be involved. Most of the technical projects fail not because of the technical issues.

5. NEXT STEPS

Our aim is to keep moving to virtual. The achievements made during this last months are really encouraging. This hybrid approach permits that we keep growing and in an transparent way by adding more physical servers to the infrastructure and also to go straight to physical any time if something goes wrong.

6. REFERENCES

Golden, B. (2007). Virtualization for dummies, HP Special edition. Wiley Publishing, Inc.

White Paper (2007). Real Help for Migrating to a Virtual Infrastructure. Portsmouth, Ecora Software. Located [07.02.2008] on World Wide Web: <http://www.ecora.com/ecora/news/whitepaper.asp>

White Paper (2007). Aligning Virtualization with Business Applications. Littleton MA, Akorri. Located [07.04.2008] on World Wide Web: <http://www.akorri.com>.

White Paper (2008). Managing Virtualized Systems – Pinpointing performance problems in the virtual infrastructure. APM experts. Located [10.05.2008] on World Wide Web: <http://www.apmexperts.com>.