Virtual Server

Virtualization is back

Virtualization today is more the renaissance of a dying technology than a brand new concept. Like the brand new Geneva tramway or the videophone, this technology has existed for decades.


Virtualization overview

The Virtual Computer is an old dream, a long story associated with computer evolution. The goal has always been to make the software independent from the hardware, allowing to follow hardware evolution without rewriting the software. We all suffer from service migrations, server reinstallations, hardware renewals, which generate work interruptions, habit changes and some inconvenience.

Between 1977 and 1979 Unix was trying to address the issue with the portability of the C language: “Write it in C, it will run on any Unix”. A few years later it was proven that the C language did not solve all the problems.

In the 1980’s, IBM 390 System had specialized circuits in the CPU allowing to virtualize itself. The Virtual Machine (VM) idea was there: the Virtual Machine Operating System was literally giving each user his own independent Virtual Computer.

In 1995, the Web generated another Virtual Machine: the Java Virtual Machine. The new idea was to have a virtual machine and environment independent from the underlying Operating System. The recurrent issue was addressed in some way: “Write it in Java, it will run everywhere”. Again all problems were not solved by Java, but this time maybe more due to commercial interests than to technical difficulties.

Vicious circle

Virtualization has always been facing performance problem: the need for better performance is always opposed to the need for portability and investment preservation.

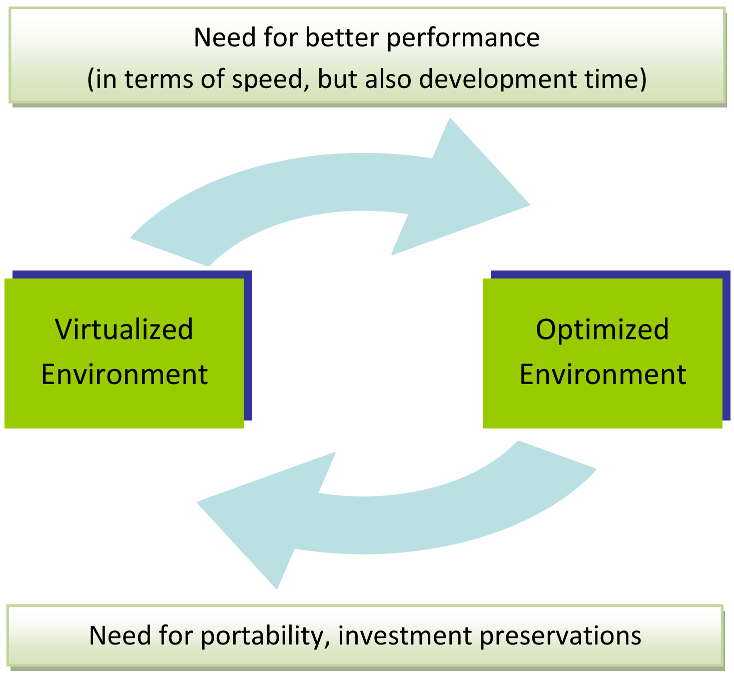

What's new today ?

Today the technology has become standardized: the "Intel PC" is now a consolidated standard. The Intel Virtual PC becomes the new Virtual Machine base layer. It can run any OS supported by the Intel PC, from Windows to Linux, and possibly Mac OS. It can run on any PC processor vendor: Intel or AMD, and any processor model (Itanium, Pentium 4, Opteron, etc…).

Today's emulation overhead is small, emulation of PC on PC is very efficient. Note also that the new generations of processors (Intel or AMD) now integrate a dedicated Virtualization Technology layer, to enhance performance of emulated machines.

Why use Virtual Servers ?

The NICE Custom Servers, a good candidate for virtualization

The CERN IT/IS Group is providing a Custom Windows Server hosting service (NICE environment) for CERN departments and groups requiring an "out of the box" Windows Server solution (see **http://cern.ch/Win/Help/?kbid=251010**).

These custom Windows servers are provided with a high availability service level agreement: servers are hosted in CERN Computer Center, connected to uninterruptible power supplies, monitored 24x365 with operator presence, daily backups with fast tape drives, hardware maintenance, operating system maintenance, security scans and patches are “transparent” to the customer.

With that kind of solution, the customer can focus only on his applications, not on the server hosting them. Note also that most of the time this customer is not willing to share his server with others, and is ready to pay a kind of “rental service” for that.

This service is today very popular, several requests per month are received from various sources, such as LHC Controls and experiments (like ALICE), Technical services, etc… as well as Video Streaming services.

More than 60 servers are running today this service, which revealed some weaknesses that could be solved using virtualization: installing and maintaining physical servers is time consuming, it requires a management overhead for logistics and resource planning. The space in CERN Computer Center is also a scarce resource, and cannot be extended infinitely. IT/IS Group also noticed that several of these servers are underused, hardly more than 2 or 3% of the CPU is effectively used.

Goal of virtualization

Virtualization creates a clear separation between hardware management and server (software) management. It could even be done by different teams.

The hardware management task includes basic hardware maintenance, but is mostly a server pool management: a large pool of identical machines is maintained, and the team has to ensure that enough server hardware is globally available to satisfy the global CPU and storage demand.

The server (software) management allocates server images to machines in the pool, manage server configuration, and can imagine various possible optimizations and automations (for example an automatic reallocation to different hardware according to past performance requirements).

Virtual Servers in real life

Foreseen advantages

The server management is dramatically simplified: ***“installing a server” becomes “loading a Virtual Machine image”***. Unprecedented automation can be achieved.

The CERN IT/IS Group has started automating the Microsoft Virtual Server 2005 product, by providing a Web Interface for its Custom Server Service customers. See http://cern.ch/WinServices/Services/WoD .

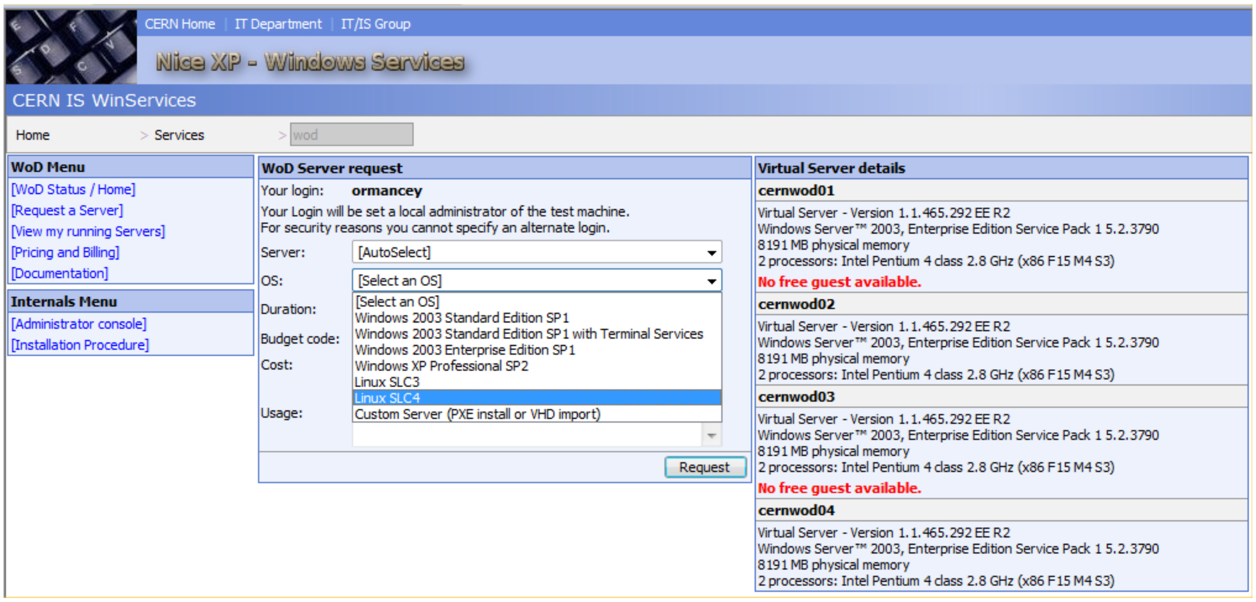

The customer can request a brand new server in 2 clicks: Select OS, lifetime of server, budget code and click Request: **10 minutes later the Virtual Server is available and ready for use**. The requester will automatically be added to the local Administrators group for Windows servers or to the sudo users list for Linux servers.

Note that a budget code is required, to make sure that customers don't abuse this virtual server facility: getting a new server is just one click ahead, but computer resources at CERN are not free !

Once his server installed and ready, the customer can interact on it, without leaving his office: Reset and power buttons are links on a web site:

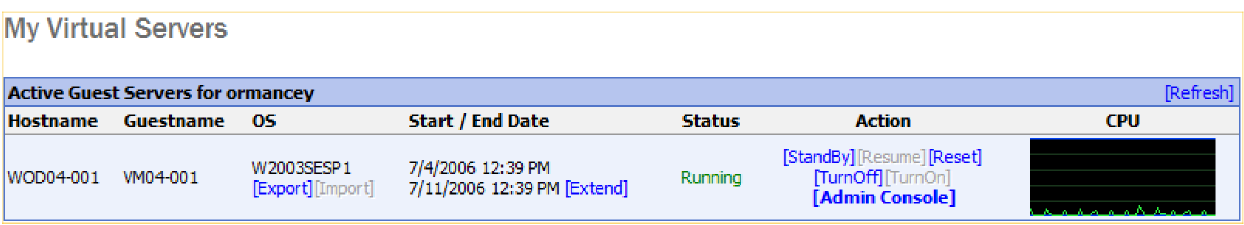

Note that the customer can export and import the disk image through this interface: he can store for later use his server, after having configured it in a specific way.

A Web interface also provides server Console access, exactly like on a real server console, allowing BIOS settings change, and various configuration options that can be done only on the console:

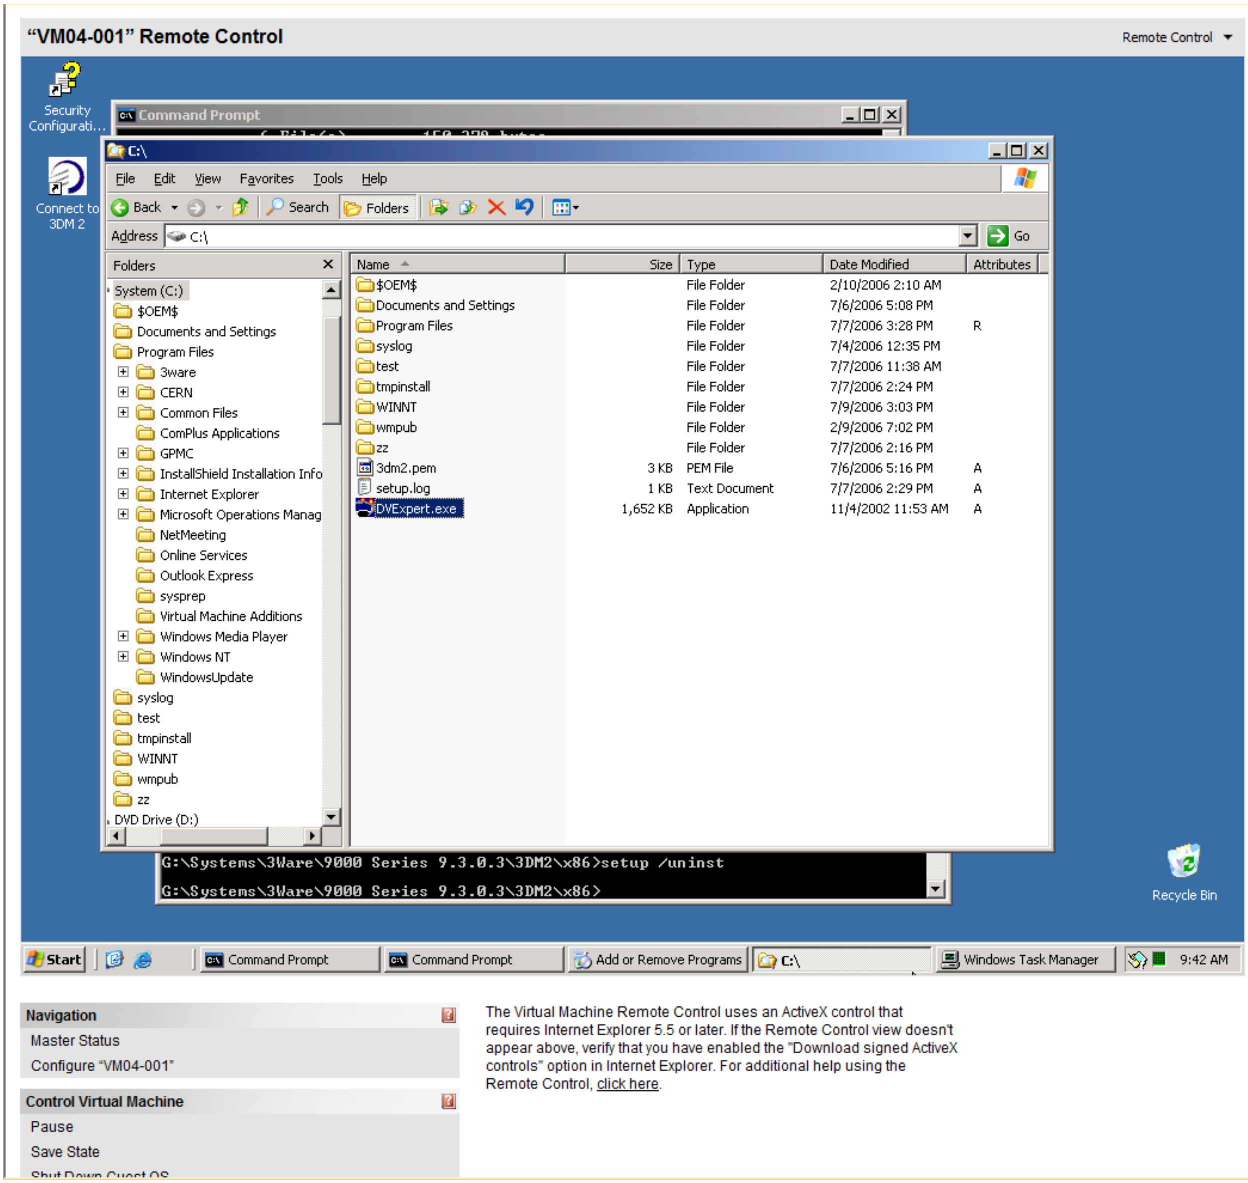

The server physical configuration can also be edited using a web site, giving access to network card, CD / DVD reader, floppies, etc…:

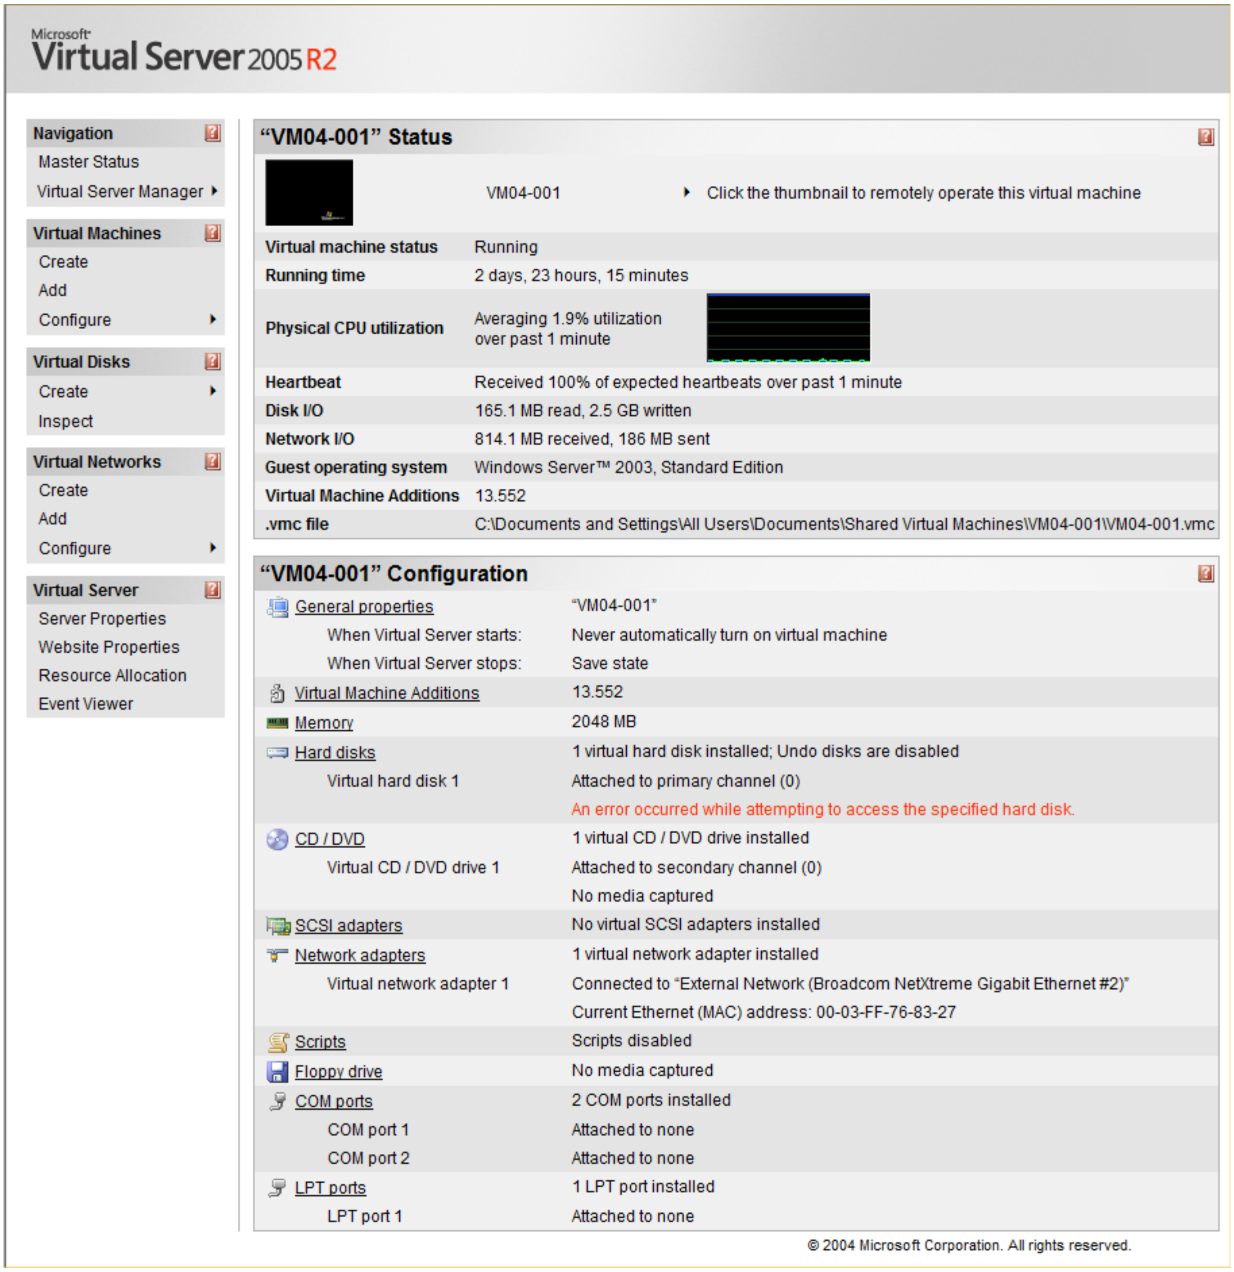

Virtual Server On Demand

This virtualization was automated to fit generic CERN needs: various OS can be requested, like Windows 2003 classic, Windows 2003 with IIS (Web server), Linux SLC3 and SLC4, etc…

An empty image can also be requested: a virtual server with no OS will be provided, the customer will have to install manually his OS on this server. To achieve this there are 2 ways:

> Like on a real physical server, he can mount a CD/DVD image and boot from it to run the installer, or press F12 to boot from the network (PXE) and start various network installations.

Taking advantage of disk images used in virtualization, the user can import a disk image to this virtual server guest. This disk image can be generated from a physical computer (any software compliant with VHD file specification can achieve this, like WinImage shareware), or come from a previous virtual server on which the user exported the disk image using the provided web interface.

When a customer requests a new server, the resource is taken from the pool of available hardware: several multiple and different OS can be hosted in the same box. The server is available in 10 minutes (time to copy the image, then join the CERN domain for Windows OS, or install updates for Linux OS).

Sharing hardware leads of course to an important cost reduction: the current IT/IS platform hosts 3 virtual servers (guests) per physical machine (host).

Performance

Customers of the virtual servers provided by IT/IS did not notice any difference. Of course an IO and CPU intensive application could suffer from virtualization, but classic applications do not.

The current service emulates servers with 1 CPU, in 32 bits mode, running at the same speed as the host (2.8 Ghz, or more). CPU can be pre-allocated, one guest can have 100% of one CPU reserved. By default the 3 guests share the 2 host's CPUs following their needs. The memory allocated is 2GB per guest.

Concerning IOs, the emulated disk drives can be IDE or SCSI, and can be dynamic or static. A dynamic drive will use only what it needs on the real physical disk, compared to the static disk which will use physically the declared emulated size. Note on the following table the read and write speed, on dynamic and static drives compared to their physical host disk:

	RAID1		
	Physical	**IDE emulated**	
		Dynamic	Static
Write (Mbyte/s)	56.11	30.27	55.43
Read (Mbyte/s)	54.16	40.07	42.16

The idea is however to use centralized file space, like DFS or AFS, to avoid storing important data on virtual server guests: this helps to replace a guest server quicker, and easily allow data to be accessed from multiple guests.

Concerning the network, emulated cards have their own generated MAC address, and share the physical host Gigabyte network card.

What's next ?

Server On Demand

IT/IS Group expects to have more server types to provide, with various combinations of OS and Applications. Requests for custom server types are also expected, where the customers will import (and export) their own server images. Note that a physical server can be virtualized by simply building an image of its hard drive(s).

The future of a "Server On Demand" service will have to address these samples:

> "I need 20 servers with this image for one month"
>
> "I need an image for this server replicated 10 times"
>
> "I need more CPU / Memory for my server"
>
> "I do not need my server for 2 months, give me an image I can reuse later"
>
> "I need a test environment, OS version n+1, to which I can migrate my current production services"

Batch systems ideas

We can also imagine the future of batch systems: instead of sending a piece of code to a pool of batch servers, the batch creator could virtualize his own machine (desktop, development environment, etc…), and send this virtualized image to a pool of virtual servers. He would then use his favorite environment, run his batch in his preferred mode, with the custom setup he cannot live without.

Note also that running a virtualized machine could prevent batch scripts or programs to create problems on their batch server: all processes are restricted to their own virtual machine, and in no way a batch execution could interfere with others. If security improvement and user experience are following the same track, we can expect virtualization to move quickly forward.

Conclusion

Server Virtualization is a strategic direction for server (currently Windows) management at CERN, as it reduces server hosting and management cost while providing more flexibility for server customers. Managing a pool of identical servers simplifies operations procedures, virtualized guests increase the range of available OS and applications while they remain independent of hardware.

We can also expect consequences for traditional batch systems, where the full user's environment will be sent to the batch pool in a sealed virtualized box, instead of sending only a piece of code that can interfere with others, giving a more robust batch service.

Emmanuel Ormancey, Alberto Pace - CERN IT/IS