



Using Virtual Servers for the CERN Windows infrastructure

**Emmanuel Ormancey, Alberto Pace
CERN, Information Technology Department**



A demo before we start

◆ <http://cern.ch/Winservices/Services/WoD>

WoD - Windows On Demand

- You need a Windows 2003 Server for testing purposes ?
- You need a Windows 2003 Server where you are local administrator ?
- You want to test your new Windows software ?
- You want to play dangerously with Windows XP but avoid reinstallation ?

Windows On Demand is the solution:
Get your Windows environment in less than 10 minutes !
New: WoD serves also Linux operating systems !

WoD Server request

Your login: **pace**
Your Login will be set a local administrator of the test machine.
For security reasons you cannot specify an alternate login.

OS: [Select an OS]

Domain: Join CERN Domain Standalone

Duration: [Select a Server Lifetime]

Budget code:

Cost:

Usage:

Message of the day

Linux SLC4 image currently has problems with AFS login, due to Time Synchronization problems. Investigations in progress.

Guest Servers Status

Hostname	Guestname	Status	CPU
WINNEW05-005	VirtualMachine01	TurnedOff	
WINNEW05-004	VirtualMachine02	Running	
WINNEW05-003	VirtualMachine03	Running	
WINNEW05-002	VirtualMachine04	TurnedOff	
WINNEW05-001	VirtualMachine05	TurnedOff	

Available OS

Name

- Windows 2003 Standard Edition SP1
- Windows 2003 Standard Edition SP1 with Terminal Services
- Windows 2003 Enterprise Edition SP1
- Linux SLC3
- Linux SLC4



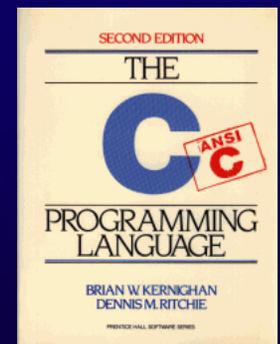
Renaissance of dying technologies





The “virtual” Computer

- ◆ An old dream, a long story
 - ◆ make the software independent from the hardware
 - ◆ Follow hardware evolution without rewriting the software
- ◆ Between 1977 and 1979, portability of Unix was being demonstrated with the C language
 - ◆ “write it in C, it will run on any UNIX”





The “virtual” Computer (cont’ed)

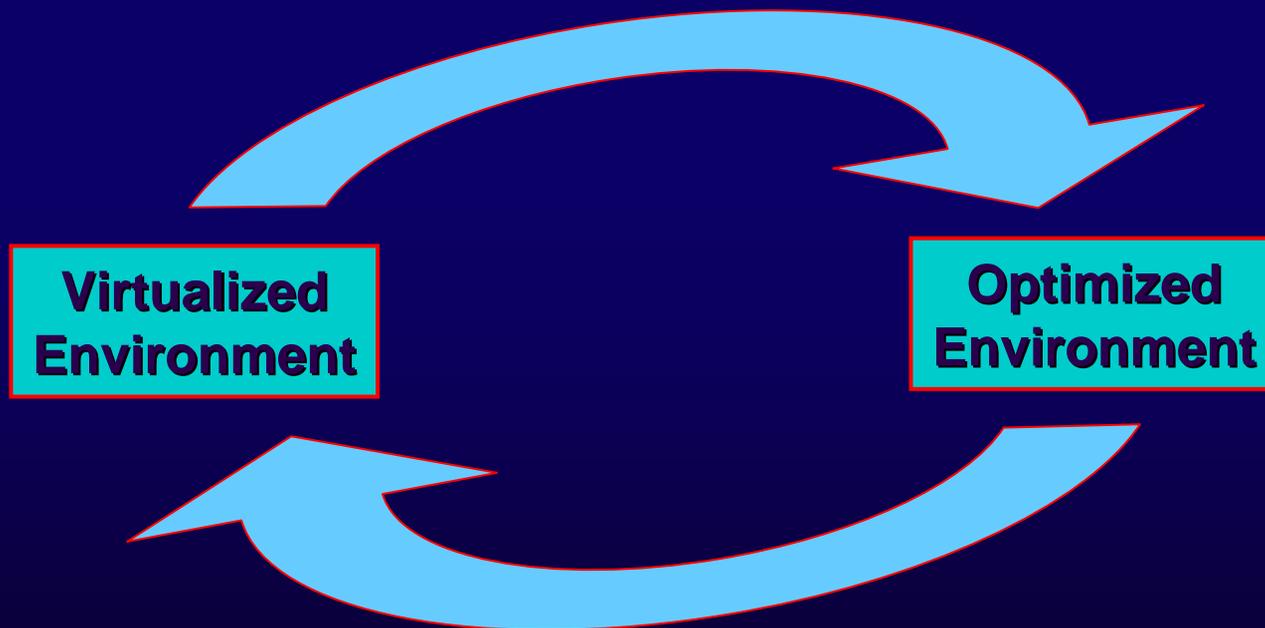
- ◆ In the '80s IBM's System/390 had specialized circuits in the CPU to allow it to virtualize itself
 - ◆ VM – The virtual machine idea
 - ◆ The VM operating system was literally giving each user its own virtual computer
- ◆ In 1995, the web required another virtual machine
 - ◆ The Java Virtual machine
 - ◆ New idea: independent from the underlying OS
 - ◆ “write it in Java, it will run everywhere”





Vicious circle

Need for better performances
(in terms of speed but also development times)



Need for portability, investment preservations



What's new today ?

- ◆ The “Intel” PC is a consolidated standard.
- ◆ The Virtual Intel PC becomes now the new virtual machine
 - ◆ It can run any flavors of Windows and Unix
- ◆ When running on Intel or AMD. On all processor families (Itanium, Pentium 4, Opteron)



Why Virtual servers

- ◆ **More and more requests for dedicated servers in the CERN computer centre**
 - ◆ **Excellent network connectivity, to the internet and to the CERN backbone (10 Gbit/s)**
 - ◆ **Uninterruptible power supply**
 - ◆ **24x365 monitoring with operator presence**
 - ◆ **Daily backup with fast tape drives**
 - ◆ **Hardware maintenance, transparent for the “customer”**
 - ◆ **Operating system maintenance, patches, security scans**
 - ◆ **“customer” focus only on “his application”.**
 - ◆ **Customer not willing to share his server with others, but ready to pay lot of \$\$, €, CHF**
- ◆ **Frame for this server hosting service:**
 - ◆ **<http://cern.ch/Win/Help/?kbid=251010>**



A real success

- ◆ Several request per month received from LHC controls, Technical services, LHC experiments, ...

<http://cern.ch/Win/Help/?kbid=251010>

The screenshot shows a Microsoft Internet Explorer browser window titled "Winservices Help - Microsoft Internet Explorer - [Working Offline]". The address bar contains the URL "https://websvc02.cern.ch/WinServices/Help/?kbid=251010". The page content includes a navigation menu with "CERN Home", "IT Department", and "IT/IS Group". The main heading is "Nice XP - Windows Services". Below this, there is a "Contents" sidebar with a tree view containing "NICE Environment", "NICE Security & Antivirus", "NICE Recommended Applications", "Support for NICE", "VPN Service", "Windows Server Service", and "Server hosting service" (which is selected). The main content area is titled "Custom Windows-based servers hosted in the computer centre" and contains several paragraphs of text, including an "IMPORTANT NOTE" and a list of "Examples of these requests are:". The list includes: "From the IT division: Application servers for Crystal Report, CAD, CAE and database servers (Euclid)", "From the administration: Business object servers, Media servers for technical training", "From the accelerator sector: Servers for controls, file servers, database servers", and "From the research sector: Data and application servers for experiments (LHCB)". Below the list, there is a section titled "Requirements from service providers" and a numbered list starting with "1. Space in the computer centre".



However, after an inside look ...

- ◆ Installing and maintaining custom servers is time consuming ...
 - ◆ Lot of management overhead
- ◆ Space in the computer centre is a scarce resource
- ◆ Several of these servers are underused
 - ◆ Hardly more than 2-3 % CPU usage
- ◆ Excellent candidate for virtualization



Goal of virtualization

- ◆ **Clear separation of hardware management from Server (Software) management**
 - ◆ **Could be even be made by independent teams**
- ◆ **Hardware management**
 - ◆ **Ensure enough server hardware is globally available to satisfy the global CPU + Storage demand**
 - ◆ **Manages a large pool of identical machines**
 - ◆ **Hardware maintenance**
- ◆ **Server (Software) management**
 - ◆ **Manages server configuration**
 - ◆ **Allocates server images to machines in the pool**
 - ◆ **Plenty of optimization possible**
 - ◆ **Automatic reallocation to different HW according to past performances**
- ◆ **Little overhead**
 - ◆ **Emulation of PC on real PC is very efficient**



Foreseen advantages

- ◆ Simplified management
 - ◆ “installing a server” becomes “loading an image”
 - ◆ Unprecedented automation can be achieved
- ◆ “Server on Demand” Service

WoD - Windows On Demand

• You need a Windows 2003 Server for testing purposes ?
• You need a Windows 2003 Server where you are local administrator ?
• You want to test your new Windows software ?
• You want to play dangerously with Windows XP but avoid reinstallation ?

WoD Server request

Your login: **pace**
Your Login will be set a local administrator of the test machine.
For security reasons you cannot specify an alternate login.

OS: [Select an OS]
Domain: [Select an OS]
Duration: [Select an OS]
Budget code: [Select an OS]
Cost: [Select an OS]

Usage: [Select an OS]

Message of the day

Linux SLC4 image currently has problems with AFS login, due to Time Synchronization problems. Investigations in progress.

Guest Servers Status

Hostname	Guestname	Status	CPU
WINNEW05-005	VirtualMachine01	Saved	
WINNEW05-004	VirtualMachine02	TurnedOff	
WINNEW05-003	VirtualMachine03	TurnedOff	
WINNEW05-002	VirtualMachine04	TurnedOff	
WINNEW05-001	VirtualMachine05	TurnedOff	

Available OS

Name

- Windows 2003 Standard Edition SP1
- Windows 2003 Standard Edition SP1 with Terminal Services
- Windows 2003 Enterprise Edition SP1
- Linux SLC3
- Linux SLC4

You are member of IT/IS Group, you have access to:
[Administrative/Win/Server]



Server on Demand

- ◆ Chose from a set of “predefined” images
 - ◆ Windows server 2003
 - ◆ Windows Server 2003 + IIS + Soap + Streaming
 - ◆ Windows Server 2003 + Terminal Server Services
 - ◆ ...
 - ◆ Scientific Linux CERN 3 or 4
 - ◆ ...
- ◆ Takes resources from the pool of available HW
 - ◆ Multiple, different, OS can be hosted in the same box
- ◆ Available within 10 minutes
 - ◆ Before: between one week and one months
- ◆ Cost: much cheaper, especially manpower
- ◆ Performances: unnoticeable difference



"VirtualMachine04" Remote Control

Remote Control ▾

```
Initializing random number generator: [ OK ]
Starting pcmcia: [ OK ]
Starting up APM daemon: [ OK ]
Found system call table at 0xc0363bf0 (scan: close+chdir+write)
Starting AFS services.....
```

"VirtualMachine04" Configuration

<u>General properties</u>	"VirtualMachine04"
When Virtual Server starts:	Never automatically turn on virtual machine
When Virtual Server stops:	Save state
<u>Virtual Machine Additions</u>	Virtual Machine Additions information not available
<u>Memory</u>	256 MB
<u>Hard disks</u>	1 virtual hard disk installed; Undo disks are disabled
Virtual hard disk 1	Attached to primary channel (0) Virtual hard disk file "VirtualHD-Dynamic-04.vhd" Maximum size is 40 GB; Currently expanded to 6.4 GB
<u>CD / DVD</u>	1 virtual CD / DVD drive installed
Virtual CD / DVD drive 1	Attached to secondary channel (0) No media captured
<u>SCSI adapters</u>	No virtual SCSI adapters installed
<u>Network adapters</u>	1 virtual network adapter installed
Virtual network adapter 1	Connected to "External Network (Intel(R) PRO_1000 MT Dual Port Network Connection)" Current Ethernet (MAC) address: 00-03-FF-36-69-06
<u>Scripts</u>	Scripts disabled
<u>Floppy drive</u>	No media captured
<u>COM ports</u>	2 COM ports installed
COM port 1	Attached to none
COM port 2	Attached to none
<u>LPT ports</u>	1 LPT port installed
LPT port 1	Attached to none

Microsoft
Virtual

Navigation

Master Stat

Virtual Serve

Virtual Mach

Create

Add

Configure

Virtual Disks

Create

Inspect

Virtual Netw

Create

Add

Configure

Virtual Serv



What's next ?

- ◆ We can expect request for more “Server types”
 - ◆ Various combinations of OS and applications
- ◆ We can expect request for custom server types
 - ◆ User creates and manages his server images
- ◆ Future server on demand
 - ◆ “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”
 - ◆ I need 10 Macintosh instances ...
 - ◆ ...



Conclusion

- ◆ **Server virtualization a strategic direction for (windows) server management at CERN**
- ◆ **HW and SW management can be independent**
- ◆ **We can expect consequences also for traditional batch systems**
 - ◆ **Instead of allocating CPU time for jobs submitted for a rigid OS configuration one could allocate bare “virtual PC time”**
 - ◆ **User would submit “PC image hosting the job”. Farm independent of OS, less security implication (for the farm management), unprecedented flexibility for users**