

Techlabs Virtually Four Years On

Dr Eduardo Correia

School of Computing, Christchurch Polytechnic Institute
of Technology

correiae@cpit.ac.nz

Ricky Watson

School of Computing, Christchurch Polytechnic Institute
of Technology

watsonrp@cpit.ac.nz

Abstract

Techlabs is a network the authors implemented some years ago in order to provide a rich learning environment through the use of virtualisation. They outlined the background to and reasons for employing virtualisation in a paper to the conference of the National Advisory Committee on Computing Qualifications (NACCQ) in 2004. In this paper they now describe some of the changes they have introduced during the past four years, in the context of recent developments and the widespread adoption, both in industry and the academic sector, of virtualisation of one form or another.

Keywords: Virtualisation, VMware, Virtual PC, Microsoft, network, laboratory, computing, education.

1 Introduction

In 2002 the authors built a network called “Techlabs” in the School of Computing at the Christchurch Polytechnic Institute of Technology (CPIT) for the express purpose of harnessing the power of virtual technologies to provision a flexible, cost effective, and powerful learning environment. Two years later they outlined their reasons for making use of virtualisation in a teaching environment, as they were early adopters of this technology and its use at that stage was still uncommon, especially in the academic sector (Correia and Watson, 2004). Since then, this technology has been widely adopted in industry, with one recent poll in the United States finding that half of respondents expected that over 80% of their servers will be virtual machines (VMs) in the next two years (CIOL Bureau, 2007). Many in the academic sector have been somewhat less enthusiastic, preferring instead alternative hardware-based solutions that contain a virtual component through the use of Virtual Local Area Networks (VLANs) (Belles, R., Gorka, S. and Miller, J.R., 2002, and Farega, L., Massaguer, J., Jove, T. & Merida, D., 2002). It appears now that increasing numbers of tertiary institutions are making use of virtualisation in at least some courses and have adapted their teaching environment as a result (Bullers, Burd, and Seazzu, 2006; Blezard, 2007; Adams and Laverell, 2005, Collins, 2005, Villanueva and Cook, 2005). (Even Microsoft have published articles outlining suggestions for building a network that utilises virtualisation (Ruest, n.d.).)

This quality assured paper appeared at the *21st Annual Conference of the National Advisory Committee on Computing Qualifications* (NACCQ 2008), Auckland, New Zealand. Samuel Mann and Mike Lopez (Eds). Reproduction for academic, not-for profit purposes permitted provided this text is included. www.naccq.nz

Still, none of these sources suggest implementing virtualisation in the same way that Techlabs does, that is, through a physical infrastructure comprising of two fully-fledged parallel networks, one that connects only VMs to one another, the other, only physical machines. This paper reflects on the past four years of Techlabs, considers some of the changes that have occurred, and offers a glimpse of possible future directions.

2 Background

Techlabs was built in 2002. Since then, it has been used for a variety of courses, including ones on networking and programming. Techlabs comprises of three classrooms, two with 20 student machines and one with 22 student machines (in addition each room also includes an instructor machine). Each student machine is part of two physical networks, one called the “Blue” network, the other, the “Red” network. These two physical networks have their own physical switches, and part of a Layer 3 switch ensures that the networks are separate through the use of VLANs. The Blue network is a conventional network: it is the one students log onto and it connects all the physical machines, namely student machines, tutor machines in classrooms, tutor machines in offices, domain controllers, Domain Naming Services (DNS) servers, a mail server, firewalls, a proxy server, a server for Microsoft updates, a centralised antivirus solution and so on. The Red network is a physical infrastructure that comprises of switches, network interface cards, cabling and connecting firewalls that are available to VMs. These VMs are mostly stored on student machines, but some basic template VMs are also stored on the file server connected to the Blue network.

3 VMware and Virtual PC Side by Side

Microsoft Official Courses (MOC) began utilising virtualisation long after it was introduced at CPIT, so that when Microsoft finally began releasing courseware employing its Virtual PC product four years ago (Correia and Watson, 2004) it meant that some courses at CPIT needed to be set up differently. Microsoft course setups had always been based on the bare metal, so that each student was given a physical machine with the operating system and associated software pre-installed. The authors simply virtualised that single physical machine and provided students with a VMware, virtualised copy. Each of these VMs was unique and networked to other students’ VMs as well as to one, perhaps two VMs controlled by the instructor. The instructor course-specific VMs were left running and available on a VMware GSX server in the Techlabs server room throughout the course on a 24/7 basis. As a result, students could complete labs at any

time, if they had been unable to attend class for some reason. This is possible because, unlike VMware Workstation (the version of VMware on the student machines), on VMware GSX Server various people can log on and off the physical machine and hosted VMs will continue to run.

The Techlabs infrastructure still allows course setups like these, should teaching staff wish to design practical work in this way, but the network infrastructure needed also to incorporate Microsoft Virtual PC (VPC) machines when Microsoft started bundling pre-built VMs with its courses. (This has now become a standard part of Microsoft courseware.) These VPC machines are copied to all physical machines, so that each student has an identical set of VMs. Most if not all the networking between course VMs do not extend beyond students' physical machines, but are instead part of a private, entirely virtualised network. One advantage of this approach is that students can do lab work on their own without involvement from either other students or the instructor, because they have their own copy of each client and server machine in their course setups. This also tends to reduce the time required for course setups, though this has at times been negated by Microsoft's own poorly implemented processes for activating VMs, a procedure quite separate from activating the software itself. As a result each student may be running four or five virtual machines simultaneously, and configuring them in order to implement a particular technology. More recently, teaching staff have started adopting Microsoft Official Academic Courseware (MOAC) for some courses, and as such the setups in these courses have no pre-built VPC VMs and require the use of a centralised instructor class server and each student VM to be networked.

Whether courses employ VMware or VPC, they now take far less time to set up. In reality courses only have to be set up once because a pristine copy of each VM is placed in an area to which students have only read access. At the beginning of the course students copy the appropriate VM(s), and it is that copy that they use to complete all practical exercises. Furthermore, when they copy their VM(s) the permissions set on the folders on student machines ensures that the student becomes the owner of the folder and as has complete control over their copy of the VM(s). This means that one student cannot change or delete another student's VMs. Furthermore, if the student ever wishes to obtain another pristine copy of the VM, this is available, simply by copying the original folder(s) that houses the VM(s).

4 Software Upgrades

Techlabs makes extensive use of Microsoft systems, just as it did four years ago. Like any network it has evolved over time; newer technologies have been implemented, older ones discarded and the infrastructure modified in order to make systems easier to manage and improve the learning environment for students and staff alike. (Obviously, staff members too continue to learn and develop and can make good use of Techlabs.) As has always been the case, VMs are mainly stored on a second (now 250GB) hard drive on local student machines because of current hardware constraints and technical

challenges to implementing this rather than on a centralised server, and the students informed that the onus is on them to make backup copies of their VMs. VMware Workstation 4 has now been upgraded to VMware Workstation 6 on all machines, and Windows Vista has replaced Windows XP Professional in both the classrooms and in the offices of staff members, with most machines utilising the 64-bit rather than the 32-bit version of the operating system. With the exception of a Unix server, other servers on the Blue network run Windows Server 2003, though the authors plan to upgrade these machines to Windows Server 2008 this year, in order to exploit some of the functionality and performance improvements in the new platform. In addition, Techlabs now also has a secure wireless network, using certificates and other technologies.

The virtual environment mirrors the real world in the way certain technologies become obsolete. For example, while it is still possible to make use of virtual floppy discs, these have all but disappeared, much as they have done outside the virtual environment. Still, if teaching staff wish to do so because there is educational value in it, they may still opt to use outdated technology long after the physical media is no longer available because that too is virtualised.

5 New Directions on the Red Network

To store and run increasing numbers of VMs requires more physical machine resources. VPC has assisted in reducing the costs of storage through the use of parent and differencing discs. This means that a base image with the operating system can be created and then each cloned machine requires only differencing disks. A Windows Server 2003 VM may therefore comprise of core files stored in the base image (or as Microsoft call it "parent") that is shared with other Windows Server 2003 VMs as well as a differencing disc. This can save a great deal of storage space, especially when the same parent disc is used multiple times (Connell, 2008). For example, the MOC for Internet Security and Acceleration Server (ISA) 2004 uses 13 VMs (Microsoft, 2005). The differencing discs of these VMs consume 19GB of disc storage, but are combined with parent discs, shared with VMs of other courses that consume 3GB. The total storage requirements for the student lab VMs are therefore 22GB, but would have otherwise have been 39GB, a saving of 17GB for each student.

The authors have only recently started utilising a similar technology (called cloning) in VMware Workstation. The Web Server Administration course uses a basic VM, which is, with the exception of the installation of VMware tools, a vanilla copy of Windows Server 2008. The unique identifiers are stripped out through Microsoft's Sysprep tool and then shutdown, and stored in a location that all users can access. An additional VM is then created by cloning the initial Windows Server 2008 VM, using the linked clone option. The size of the initial VMware guest is less than 5MB. The initial Server 2008 files are then right protected, and the linked clone placed in area that students can access. As the students require two VMs (one for Internet Information Services (IIS) and one for Apache), they can simply copy the linked clone twice in a new folder, and in so doing utilise less than 10MBs.

VMware Workstation 6 also has other features that are useful for learning, including being able to make a “movie” that captures, through the use of an AVI file, someone configuring a VM. A lecturer can therefore demonstrate how to do something in a VM and may also wish, perhaps as part of an assessment, to have students show they can do something.

Now that VMware Workstation 6.0 is used throughout Techlabs, it is possible to use a feature called “teams”. By using teams teaching staff can (a) manage a number of VMs as a coherent group or set of machines, and (b) create customised virtual networks between the team member VMs. The VMs in the team can also be started individually or as a group with a delay between each start up. Furthermore this configuration is centralised in a single place, instead of having to access the VMs individually (figure 1). The virtual networks in teams can have a maximum throughput and introduce packet loss, useful features in a course where for example students are required to complete a project that simulates the connection between two sites over the Internet using a Virtual Private Network (VPN) (Microsoft, 2005). The configuration of networking in the team is simple to perform, both the network definition and how each server connects to different networks. The prime purpose of teams in projects is to keep machines together as one unit, but this feature can also be combined with cloning to reduce the disk space VMs use. The team of VMs for one particular practical project was 5.9GB for the seven machines, all clones of a Windows Server 2003 base machine. This is a saving of more than 6GB per student from the previous time this course was run.

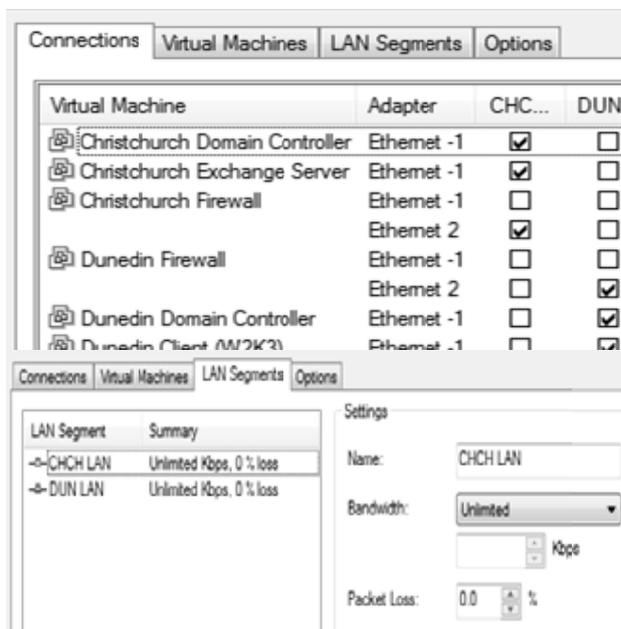


Figure 1: Networking in Teams

While Microsoft are making a concerted effort in the area of virtualisation, particularly with the advent of Hyper-V in Windows Server 2008, the general consensus is that their software is still not as mature as that of VMware. Even commentators who are generally supportive of the Microsoft platform, such as Paul Thurrott, the editor of Windows IT Pro, have said as much (Thurrott, 2007).

Others are somewhat less complimentary, describing Microsoft’s efforts in Windows Server 2008 as “roughly analogous to VMware Server 1.0, although not as polished” (Venezia, 2007). Still, Microsoft may be able to attract some market share because of lower cost, especially if they provide preferential pricing for its own licensed software running on VMs and its virtualisation products have better hardware support than those of VMware. With the acquisition of XenSource last year, Citrix has shown that it too has serious intentions in the area of virtualisation, and has confirmed its commitment to the Xen open source community (Citrix 2007). Developments will continue and competition to VMware will doubtless increase, but at present its ESX Server is the most mature virtualisation product of its kind on the market, and is capable of running enterprise server applications, such as Exchange and Oracle (Wolf and Halter, 2005). It is for this reason that the authors decided to use this platform in Techlabs to extend the benefits of virtualisation to the Blue network.

6 New Directions on the Blue Network

Four years ago VMs were restricted to the Red network and used exclusively for learning. Today, many (not all) of the major servers in Techlabs are now no longer physical machines but instead VMs running on a VMware ESX server. Both VMware Workstation and VMware GSX Server (a different product to VMware ESX Server and now superseded by VMware Server) both run as an application on top of a host operating system, such as Windows Server 2003 or Linux. This means that calls from the VMs to the hardware are first translated by the host operating system. By contrast VMware ESX Server “operates independently from any general-purpose operating system” (VMware, 2007) and is tuned to running multiple VMs efficiently on the given hardware. This has significant performance benefits and makes it a scalable solution to server virtualisation.

One of the VMware ESX Server servers in Techlabs (called both “VM03” and “esx-01” but referred to only as “VM03” for purposes of this paper), hosts a number of the core servers on the Blue network. As the left panel of figure 2 shows, this includes three firewalls (the Blue, Red and Green firewalls), an application server (AP03, which hosts a number of web sites), a database server (DB03), a domain controller for authentication of staff and students (DC03), a mail server (EX03) and a Linux server (UX03). (It was decided that VM servers on the Blue network would be designated with 03 and above.) Previously these VMs each existed in a separate physical box but now exist as VMs, consolidated onto a single hardware platform, with disc redundancy (RAID 1 and RAID 5), and a reasonable amount of memory (16GB). Techlabs does not have its own separate Internet connection but instead connects to the Demilitarised Zone (DMZ) of the main CPIT network. Virtualisation enables much greater flexibility and control over available hardware resources, including memory, CPU and network bandwidth, but there are some risks at having so many servers on a single set of hardware. The authors intend to purchase additional server hardware that is at least as powerful as VM03, in order to procure some measure of redundancy.

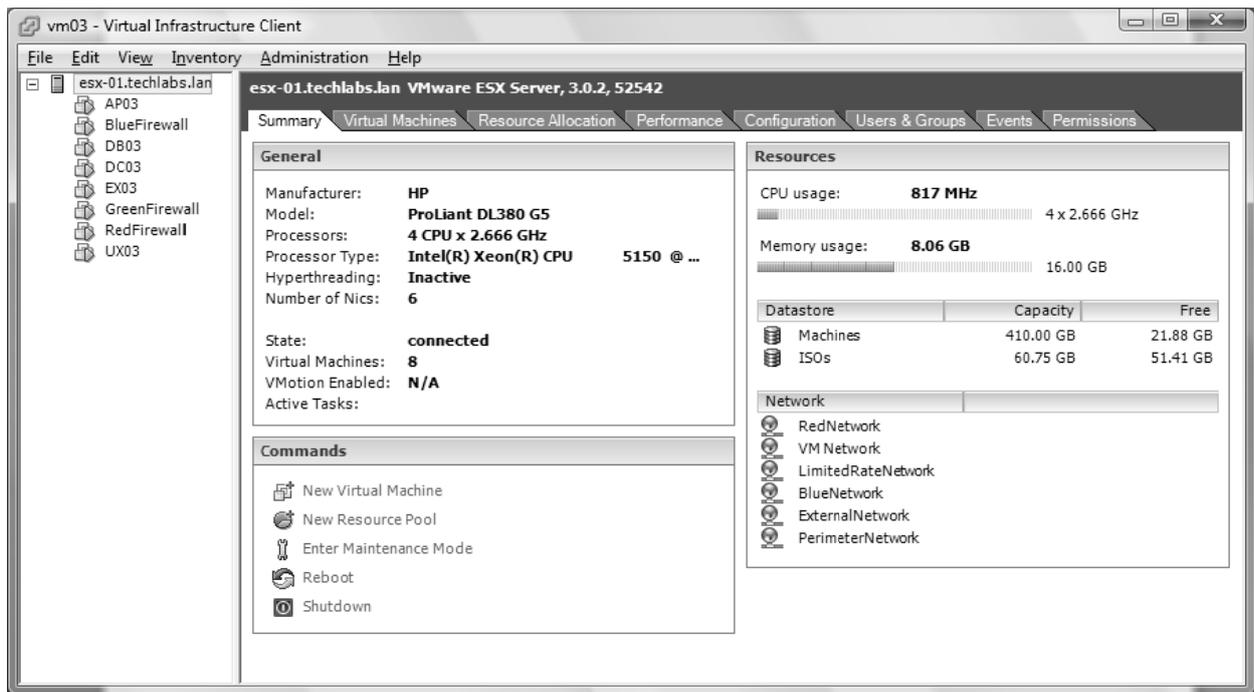


Figure 2: VMware ESX Server

7 Networking VMware ESX Server

One useful feature of VMware ESX server is its ability to utilise multiple virtual switches (vSwitches). These vSwitches map to one or more physical network interfaces, and may even be entirely virtual entities with no mapping to any physical device whatsoever (figure 3). This ensures maximum flexibility and control over network resources. In VM03 the six physical network adapters are allocated to various VMs in different kinds of ways. As figure 3 shows, vSwitch2 enables the Blue Firewall, DC03, EX03, AP03, UX03 and DB03 to connect to the Blue network through two interfaces, vmnic1 and vmnic2. vSwitch7, by contrast, is not mapped to any physical network adapter, but used for the perimeter network, through which traffic to the Internet from the Blue network is directed (as figure 4 also shows). All VMs students build are able to access course-related servers only on the Red network and can only access machines connected to vmnic5, which has considerably less bandwidth available compared to the adapters on the Blue network. Students can access the Internet through their VM with this interface, but cannot connect to the servers on the Blue network - with one exception, they can use the DNS servers on the Blue network for name resolution on their VMs.

As figure 4 shows, VM03 contains three firewalls, all running Microsoft ISA 2004. The Blue network and the Red network each have a firewall, with these two firewalls connected to an upstream Green firewall, which utilises a public IP address. All network traffic to and from the Internet passes through this interface, whether it has its origins in the Blue or Red network. The Blue firewall has three interfaces (figure 4): vSwitch 1, vSwitch2, vSwitch7 (figure 3). The Red and Green firewalls are connected to vSwitch6, and its Virtual Machine Port Group of LimitedRateNetwork (figure 3) (see also figure 4). While vSwitch6 is not mapped to any physical network adapters, it

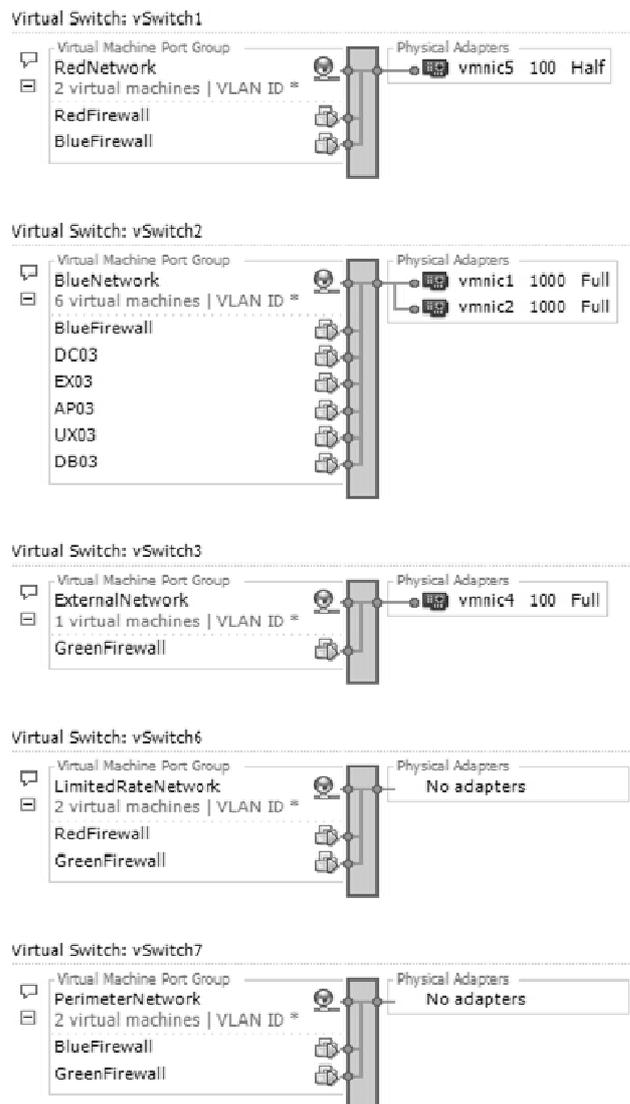


Figure 3: VMware ESX Server

is configured in such a way that ensures that student VMs on the red network have full Internet access but at a greatly reduced bandwidth compared to the physical machines on the Blue network. This means that the Internet connection from the Red network is sufficient for the purpose of study but discourages abuse of this resource on a network that is not closely monitored because it is meant to provide students with a great deal of freedom to experiment and learn.

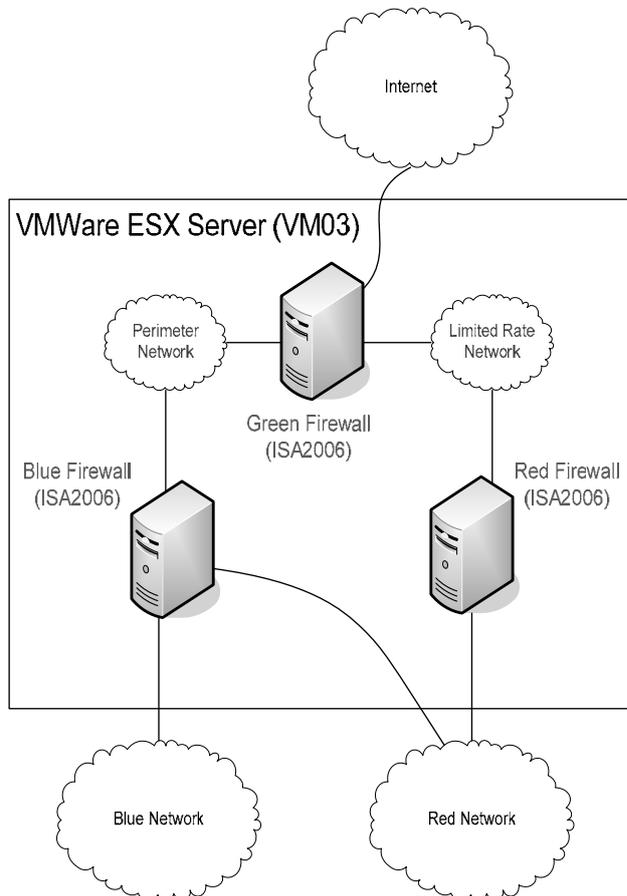


Figure 4: Network Design on VMware ESX Server

8 Conclusion

Over the past four years, Techlabs has become an established part of the School of Computing at CPIT. It has proven itself to be an invaluable resource for the purpose of learning, not just for students but also for staff. Anyone with access to this network can copy a VM pre-installed with for example Windows Server 2008, Windows Vista, or Suse, Debian, Ubuntu or another distribution of Linux. These VMs can easily be configured to make use of the infrastructure to access the Internet or connect to other VMs, both those dispersed among various student machines and those located on a central server. The “physical” resources on these VMs may be extended to include more discs or network interfaces, just like any machine. What is important is that people with access to Techlabs can have full control over their own (virtual) systems, while ensuring that the shared physical infrastructure is robust and secure.

It appears that Techlabs is still the only network of its kind anywhere in the world. This seems to be an extraordinary

claim, so why do we make it? Certainly it is true that virtualisation is increasingly being implemented in tertiary institutions (Bullers, Burd, and Seazzu, 2006; Blezard, 2007; Adams and Laverell, 2005, Collins, 2005, Villanueva and Cook, 2005), but no one else has to date suggested implementing virtualisation in the form of a dual parallel network structure, one dedicated to the physical (largely student) host machines, the other, through virtualisation, dedicated exclusively to learning with its own network cards, switches and making use of the guest TCP/IP stack rather than that of the host.

This initial design of Techlabs has now been adapted, so that many key servers that previously formed part of the physical infrastructure have now also been virtualised, and migrated to the VMware ESX Server platform. The reason for this is that VMs make much more efficient use of hardware resources and are more flexible and easier to manage than conventional systems. The dual nature of the network infrastructure, its use of parallel Red and Blue networks, though, remains because it supports the imperatives of computing education so well. One of the major differences between Techlabs now and Techlabs four years ago therefore is that the Blue network contains mostly virtual machines whereas previously it comprised entirely of physical machines.

Nothing, it appears, is certain, but virtualisation is here to stay. Well, for the moment at least.

9 References

- Adams, J.C. and Laverell, W.D. (2005): Configuring a Multi-Course Lab for System-Level Projects. *Proc. of the 36th ACM SIGCSE* (pp.525-529). St. Louis, Missouri, USA.
- Belles, R., Gorka, S. and Miller, J.R. (2002): Flexible Network Topologies for your Computing Lab: VLAN and Router Technology at Work in the Classroom. *Journal of Computing Sciences in Colleges*, **17** (3), 53-59.
- Blezard, D.J. (2007): Multi-Platform Computer Labs and Classrooms: a Magic Bullet? *Proc. of the 35th annual ACM SIGUCCS Conference on User services Conference on User services* (pp.16-20). Orlando, Florida, USA.
- Bullers, W.I., Burd, S.D. and A. Seazzu. (2006): Virtual Machines - An Idea Whose Time Has Returned: Application to Network, and Security, and Database Courses. *ACM SIGCSE Bulletin (Inroads)*, 102-106.
- CIOL Bureau (2007). IT managers planning for mass virtualization rollouts. Accessed 28 February 2008. <http://www.ciol.com/content/201107101581.aspx>
- Citrix (2007): Citrix Completes Acquisition of XenSource. Accessed 19 February 2008. <http://www.citrix.com/English/NE/news/news.asp?newsID=683171>
- Collins, D. (2006): Using VMWare and live CD's to Configure a Secure, Flexible, Easy to Manage Computer Lab Environment. *Journal of Computing Sciences in Colleges*. Vol.21, 4 (April 2006), 273-277.
- Connell, A. (2008): HOWTO: Use Virtual PC's Differencing Disks to your Advantage.

- <http://www.andrewconnell.com/blog/articles/UseVirtualIPsDifferencingDisksToYourAdvantage.aspx>. Accessed 19 February 2008.
- Correia, E. and Watson, R. (2004): A Virtual Solution to a Real Problem: VMware in the Classroom. *Proc. of the 17th NACCQ Annual Conference*. Christchurch, New Zealand. 8 July 2004. 250-253.
- Farega, L., Massaguer, J., Jove, T. & Merida, D. (2002): A Virtual Network Laboratory for Learning IP Networking. *Proc. of the ITICSE 02 Conference*, Denmark, 161-164.
- Microsoft. (2005): *2824B Implementing Microsoft Internet Security and Acceleration Server 2004*. Redmond, WA: Microsoft Corporation.
- Ruest, N. and Ruest, D. (n.d.): Build a Virtual Laboratory with Virtual Server 2005. Part 1 of two parts. <https://mcp.microsoft.com/mcp/resources/resources.mspx>. (Available on the Microsoft "MCP only" site.) Accessed 11 February 2008.
- Ruest, N. and Ruest, D. (n.d.): Working with a Virtual Laboratory. Part 2 of the article "Build a Virtual Laboratory with Virtual Server 2005." <https://mcp.microsoft.com/mcp/resources/resources.mspx>. (Available on the Microsoft "MCP only" site.) Accessed 11 February 2008.
- Thurrott, P. (2007): Microsoft Offers Hyper-V Holiday Surprise. In *Windows IT Pro Perspectives* (InstantDoc #97838). <http://windowsitsecurity.com/Windows/Article/ArticleID/97838/97838.html>. Accessed 18 February 2008.
- VMware (2007): The Architecture of VMware ESX Server 3i. http://www.vmware.com/files/pdf/ESXServer3i_architecture.pdf. Accessed 25 January 2008.
- Villanueva, B. and Cook, B. (2005): Providing 24/7 Virtual Access and Hands-On Training using VMware GSX Server. *Proc. of the 33rd annual ACM SIGUCCS Conference on User services*. Monterey, CA, USA. 421-425.
- Venezia, P. (2007): Next-gen virtualization from Microsoft. http://www.infoworld.com/article/07/12/19/51TC-microsoft-hyperv-preview_2.html. Accessed 19 February 2008.
- Wolf, C. and Halter, E.M. (2005): *Virtualization: From the Desktop to the Enterprise*. Berkeley, CA: Apress.