Virtualisation for local cloud computing: DWU data centre

Peter K Anderson

Abstract

Virtualisation, meaning in effect but not in reality, is a concept of increasing importance in computing. DWU ICT Services Department has developed a Data Centre where both Server functionality and data storage are based on this concept with resulting economy of purchasable hardware as well as power and maintenance costs. This study was based on personal observations and interviews with senior ICT Services staff. It points to the way of the future where various forms of virtualisation of computing services, known as "Cloud Computing", could become as accessible a public utility as power on the electricity grid.

Key words: Data centre, virtualisation, cloud computing, storage area network, virtual machines, logical and physical servers.

Introduction

Whilst we all like to think we are living in a real world, there is a paradigm shift occurring in the world of computing today whereby the devices we deal with or experience are more often actually virtual or logical devices which simulate the function of remotely located real or physical devices. In some cases the real device can spawn multiple virtual devices, each delivering a nearly similar level of performance as the single physical device and thereby significantly conserving computer processing resources. It should be noted, however, when considering the advantages of virtualization that some performance penalty may be incurred with the extra demands made on hardware.

This process, known as virtualization, essentially involves a decoupling process (Abts, D, & Felderman, B., 2012) which separates, in various ways, the physical implementation from the virtual effect. Decoupling is here meant to convey a process whereby the interdependence of previously strongly linked systems is loosened, separated or detached so that they can operate with various degrees of independence. Forms of virtualization becoming more significant today involve decoupling by locating hardware separately from its effects leading ultimately to what is now known as cloud computing.

This paper seeks to document the forms of virtualization occurring in the Divine Word University (DWU) Data Centre hosted by the DWU ICT

Services Department. Both network servers and data storage units are installed using forms of virtualization. This will be further discussed as a local version of cloud computing, where the actual physical devices are installed locally but remotely from where their effects are experienced. It will then be shown that the physical devices can be further decoupled from their effects by being placed at a central point on a wide area network (WAN) and accessed by many consumers of processing power on an as-needed basis much as electrical power consumers access electricity on demand via electrical grids which are regional transmission and distribution networks.

Virtualisation

Virtualisation is a term already widely used in computer architecture where, for example, the amount of available processing memory is expanded by setting aside a portion of the hard drive (pagefile.sys) to simulate physical memory as virtual memory (Figure 1, Adapted from Andrews, 2006, P 275, Figure 6-25). Decoupling is produced by the virtual memory manager (VMM), that part of the operating system (OS) which translates virtual memory addresses to corresponding real physical memory addresses.

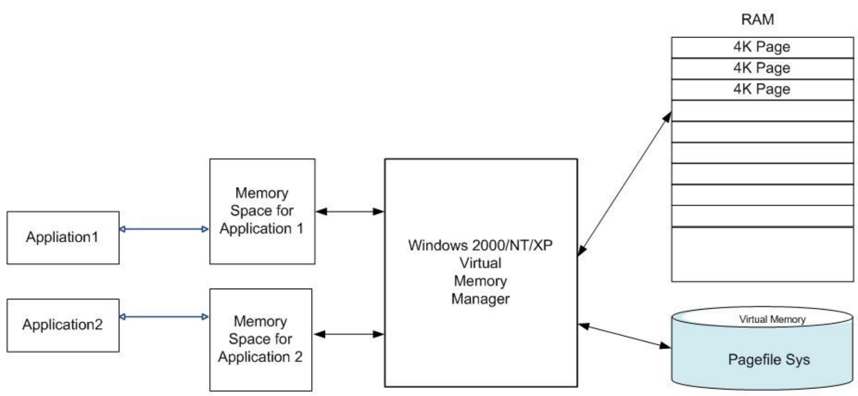

Figure 1 Decoupling by a VMM

A VMM decouples the memory space allocated to an application from the actual physical memory locations. There is also a second level of virtualisation as space on the hard drive (PageFile.sys) known as virtual memory functions to simulate real memory to produce the effect of extra Random Access Memory (RAM), albeit with significant decrease in access time. The VMM enables multiple applications to run apparently simultaneously in protected memory areas.

Virtualisation of a whole computer system has also long been used to allow multitasking and the implementation of several virtual machines (VM) on a single physical system (Figure 2, adapted from Andrews, 2006, P 608, Figure 13-2). Multitasking allows the processing of several applications seemingly simultaneously, but actually each application receiving regular periodic time slices of the processing power of the CPU. Here several Microsoft operating systems function as VMs with each VM appearing to have exclusive access to the processing hardware.

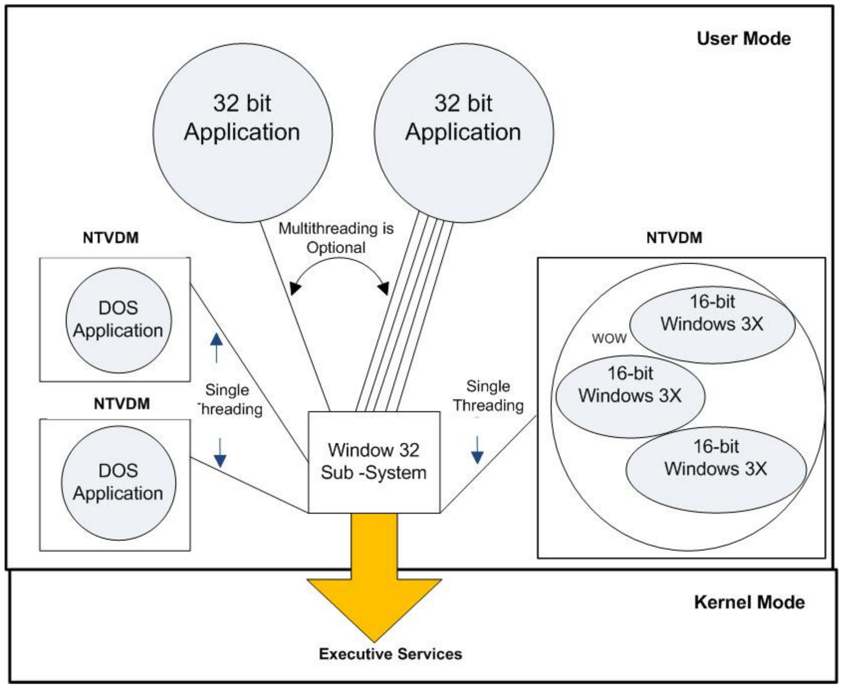

Figure 2 Traditional applications of virtual machines

This access to the processing hardware is via the software operating in kernel mode which manages hardware. Examples of operating systems are multiple copies of DOS (NTVDM, NT Virtual DOS machine). Another NTVDM runs Windows 3.x applications (Windows on Windows, WOW) and each 32 bit application running in its own VM. Also shown are threads which are single lines of processing instructions each of which functions with apparent exclusive access to the system. Each thread, in turn, is given a time slice of access to the processing power of the CPU**.**

Other well-known examples are Virtual Private Networks (VPNs) and Virtual Local Area Networks (VLANs). Access to the Internet at DWU is via a Virtual Private Network (VPN) which is a path through a public network (typically the Internet) traversed over a virtual path by encrypted (thus private) data. A logically isolated private network is provided by cryptographic methods to secure data traversing the public network (Waldspurger & Rosenblum, 2012).

At DWU, student and staff networks exist on separate Virtual Local Area Networks (VLANs), each VLAN being a part of a switched network separated from other parts of the physical network based on switch port configuration rather than physical separation.

In all these examples, the virtual simulates a physical effect. Having multiple operating systems accessing one hardware system (Figure 2) is the forerunner of new forms of virtualisation which are today rapidly transforming enterprise networked data centres which use server virtualisation.

Server virtualisation

As noted, virtualisation is based on the use of software to simulate the existence of hardware. A single physical instance of a device presents itself as multiple logical devices from which it is deliberately separated. It should be further noted that server virtualisation is only one of a number of forms of computer hardware virtualization. Others include Storage virtualization (to be discussed later), I/O virtualization (multiple VMs accessing I/O devices through a single physical channel), and Desktop virtualization where a desktop package resides separately on a server and provided on demand. All involve using software to simulate multiple instances of hardware for high performance, availability and efficiency.

Server virtualisation enables the running of multiple server guest OSs (VMs) on a single physical machine. VMs are formed with guest OSs, where the term “guest” is used in contrast to “host”, the OS which provide the hardware abstraction or emulation. Each VM (Figure 3) can, in turn, run multiple applications.

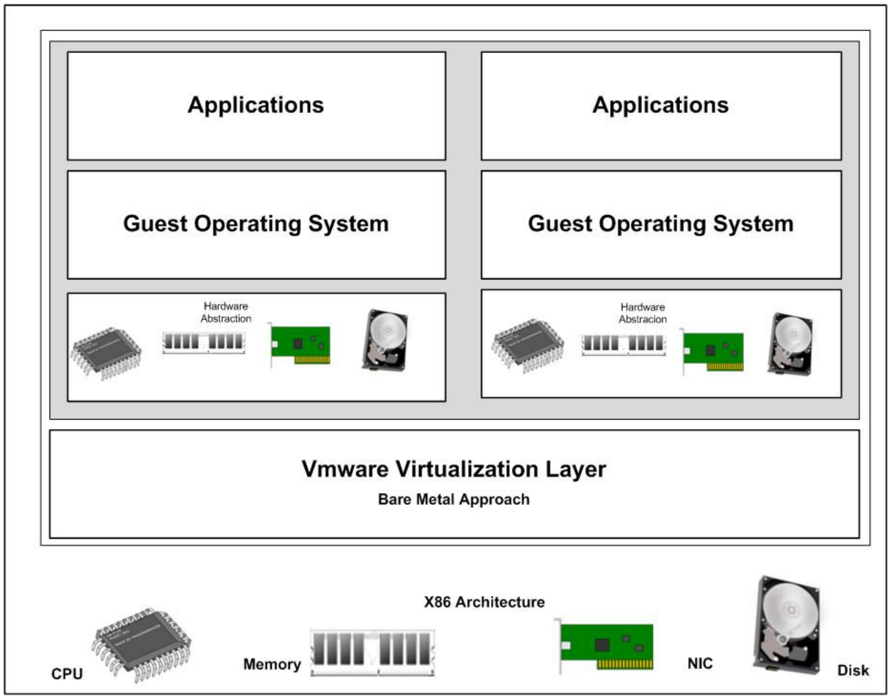

Figure 3 Bare metal approach to hardware virtualisation. (Adapted from Golden, 2008, P 24, Figure 1-2).

Virtualisation is effected by hardware emulation software in the virtualisation layer. Shown here are two (of many possible) installed guest OSs running applications and forming VMs.

Server virtualisation can be achieved by the emulation of hardware using software in the virtualisation layer which runs directly on the hardware and known as the "bare metal approach". The emulation software contains the core only of an OS known as the kernel, that part of the OS which interacts with hardware and manages the running of applications. It virtualises or provides an abstraction of the hardware to each guest operating system thereby forming a decoupled VM.

This hardware emulation is only one of a number of methods of developing VMs, another being OS virtualization whereby each application uses an abstraction or logical view of the single physical OS rather than of the hardware.

The DWU data centre uses four rack mounted physical server machines (Figure 4) to provide a much larger number (35 at present) of virtual servers, each provided with its own abstraction of the hardware, both for processing and storage. All devices are rack mounted in a secure air conditioned data centre.

Figure 4 Four clustered rack mounted physical servers

The four Dell R710 PowerEdge physical servers are configured to function as a single system or cluster to ensure service continuity by means of load balancing, high availability, and fault tolerance against the failure of any one system. All servers have access to all storage devices on an as-needed basis. The cluster allocates sufficient resources to service many requests coming from the busiest periods of network operation. Each server has 8 Gb Ethernet interfaces supporting the required cable connections (Figures 4 and 6). The cluster provides processing power for each of the multiple VMs as required. A second level of service availability will be the availability of multiple virtual servers after hypervisor virtualization.

Each physical server has installed VMware's virtualisation hypervisor known as ESXi (Figure 5, ESXi 1 to ESXi 4) and available free on the Internet. The term "hypervisor" derives from operating systems being sometimes referred to as supervisors of hardware infrastructure and applications requiring processing. The term hypervisor is then used to indicate a supervisor of supervisors or a higher level supervision. The hypervisor uses the bare metal approach installing directly onto hardware without an intervening OS. It partitions each physical server into multiple VMs, each VM appearing to run simultaneously and having access to the resources of the underlying hardware. Each VM in turn can host multiple applications using pre-emptive multitasking whereby each application is given a time slice of the CPU's processing power, which in turn is a time slice of the time slice allocated to the respective VM.

The R710 servers feature embedded hypervisors, large memory capacity with 18 DIMM slots, and 4 integrated network connections for machine-per-server capacity. They use Intel® Xeon® processor technology adapting to software in real time, processing many tasks simultaneously. The servers use the Intel Xeon 5500 quad core processor series, with 2.26 GHz speed on each core giving a total processing power of 72 MHz, of which 64 MHz is available for server software (Figure 14). Of a possible 144GB of RAM, 64 GB are presently installed. Each motherboard has 2 processor sockets with 4 cores per socket. Each physical server has 8 Ethernet ports enabling multiple cable connections (Figures 4 and 5).

Every physical server is connected to every switch to enable clustering, sharing of resources by VMs, and centralised management for the each ESXi and of the multiple VMs hosted. The administration is effected by VMware vCentre™ Client management software on a DELL R610 server machine (Figure 5) and is accessed via a management console (Figure 6, lower RHS).

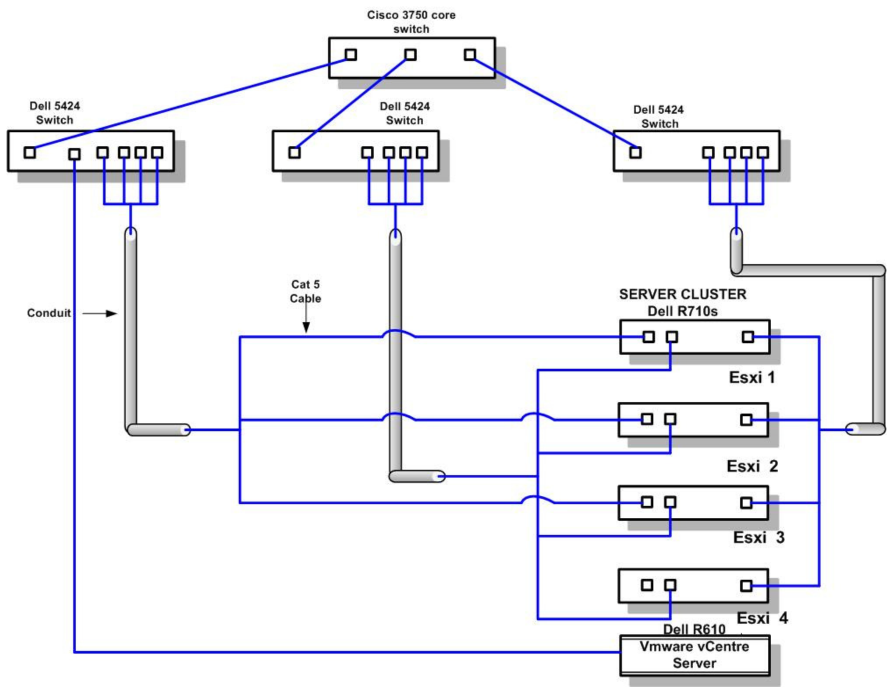

Figure 5 DWU Data Centre server Ethernet connections

Server Ethernet interfaces

Server Ethernet interfaces are distributed across multiple switches which in turn provide client machines with access to the servers and also server access to devices providing storage space (Figure 6).

Cabling distributed to multiple switches also connects to PS6000 storage devices, giving all VMs access to storage in a Storage Area Network (SAN) network. This mesh cable distribution between client machines, switches, server machines and storage devices is designed to provide high performance and availability.

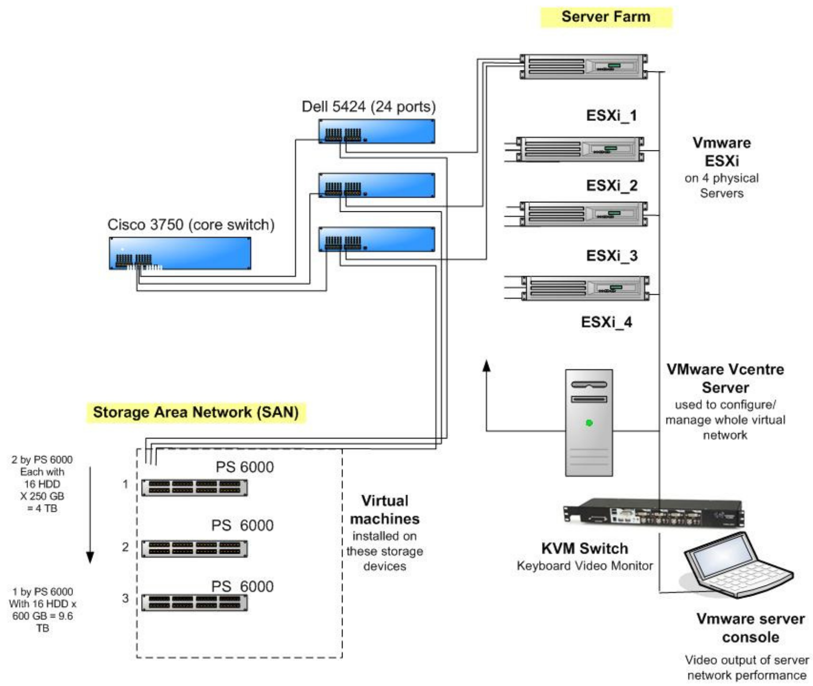

Figure 6 Data Centre cabling

All VMs have access to virtual storage in a SAN network**.** The Key, Video, Monitor (KVM) switch enables the console to access to any selected server for virtual server management and maintenance.

Virtual Machines (VMs)

Operating systems installed on virtual servers together with software abstracted emulation of the hardware form VMs (Figure 7). The Operating systems installed to form the VMs include Windows Server 2008 R2, CentOS (an enterprise class Linux distribution), UNIX (Free BSD), and Linux. Each VM can, in turn, run multiple applications. Here each VM runs one installation of specific server software. Although software is installed on VMs on specific physical servers, one server can call on other servers in the cluster to provide processing power as required in the interests of high availability.

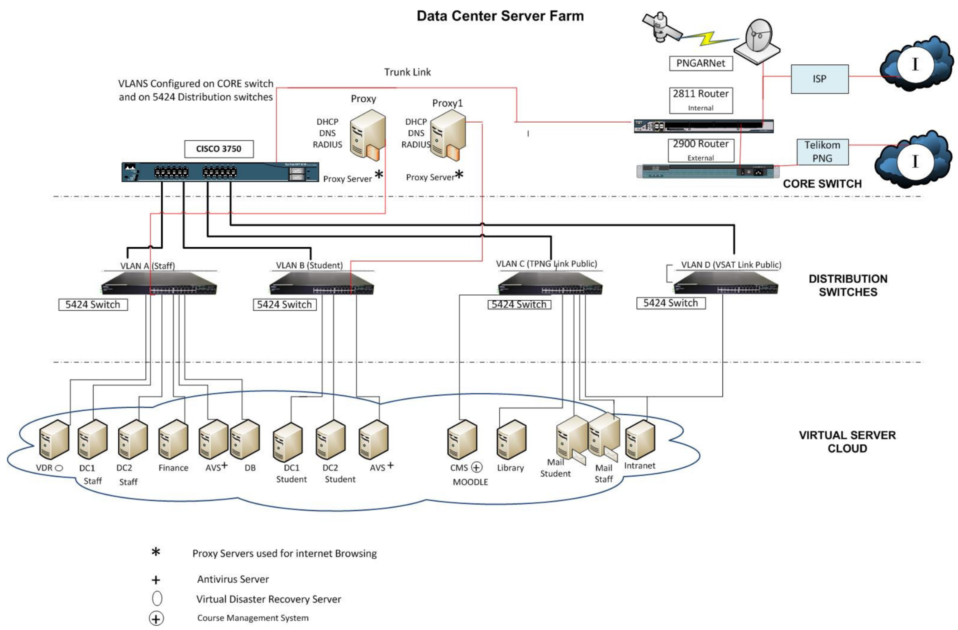

Figure 7 Mapping of virtual servers

Server virtualisation at the Data Centre presently provides 35 virtual servers (some shown within the cloud graphic) each on its own VM, with the VMs located on 4 physical servers thereby constituting a server cloud. The cloud symbol is used as a metaphor for the Internet as well as an abstraction of the complex infrastructure it represents. It used to indicate a service when details of its functioning are not relevant to the context under present focus. The virtual servers are located on four separate VLANs configured on the ports of the core and distribution switches.

The previously mentioned server management software (Figure 8) allows administration of the VM Network (LHS panel) of the 4 physical servers (Cluster 01) and 35 production and test VMs, where names of these production VMs, used with permission, are effaced for security purposes.

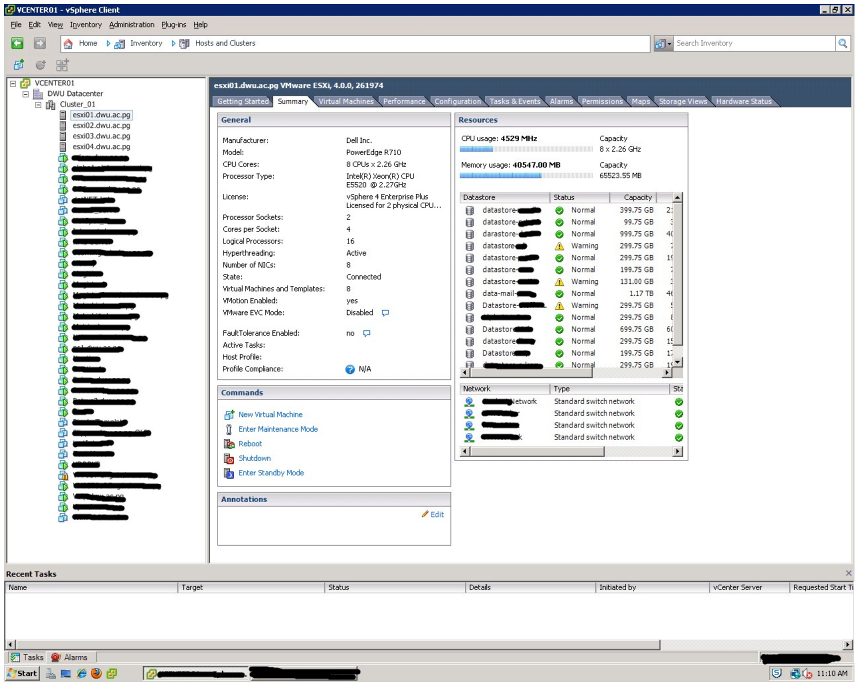

Figure 8 Server management software

The virtual servers include domain controllers (2 for the student VLAN and 2 for the staff VLAN providing network access authentication), and servers for the Course Management System (CMS), library reference software, Antivirus software, Email and Intranet software as various other database services. There are two Internet gateways each with a proxy server to supervise access to the Internet, one through the PNGARNet ISP and the other through PNG Telikom's fiber optic access to the PPC-1 undersea cable (Sydney-Guam PIPE Pacific Cable, Anderson & Kim, (2011)).

As well requiring access to processing power, the VMs require access to storage space which is provided by a Storage Area Network (SAN).

Storage Area Network (SAN)

Data Centre Storage Area Network (SAN) has 3 PS 6000 storage devices (Figure 9, with only 2 shown here for simplicity). Through the switching hierarchy, all switches and storage devices are connected through a mesh network, reducing bottle necks, latency, and a single point of failure.

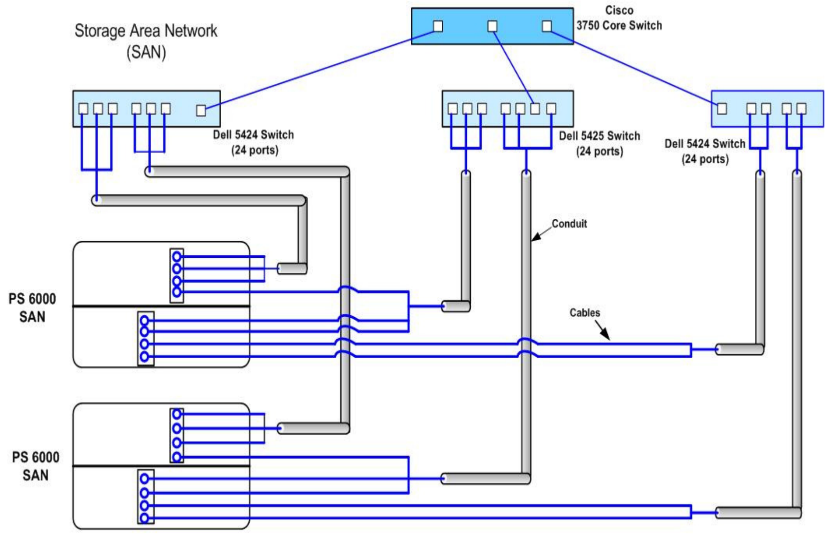

Figure 9 Data Centre Storage Area Network (SAN)

SANs are dedicated specialized high-speed networks designed especially to give efficient access to an array of storage devices in a data centre (Velte & Velte, 2007). Unlike packets in IP networks, the data units enabling faster data transfer as raw blocks of data are here transferred on Cat 6 cables, which are IEEE physical layer standards for Ethernet networks capable of carrying data in the gigabit (Gb) range.

The SAN uses a form of storage virtualization, with each server and VM being provided with a logical or abstracted view of the physical storage hardware by storage management software.

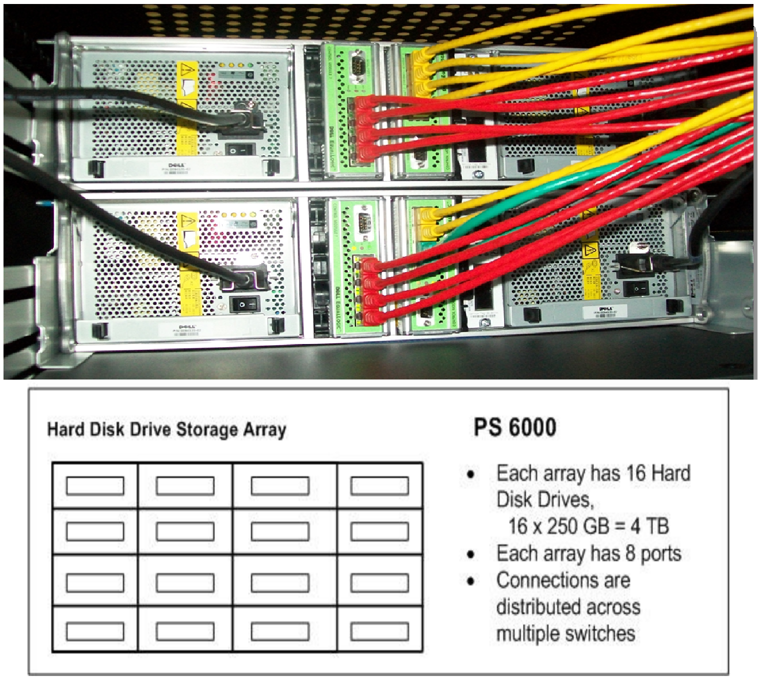

Figure 10 SAN physical connections

Thus the SAN is designed to make any storage device available to any user or device on the main network on an as-needed and shared basis. The SAN is preferred ahead of aalternative storage device configurations such as Directly Attached Storage (DAS) with each server having its own (possibly under or over utilized) storage device, and Network Attached storage (NAS) with a specialized file server providing access to the storage devices.

Any virtual server has access to any disk making all disks are available for storage, thereby maximizing disk utilization with none being underutilized. The mesh network ensures no single point of failure, and so high availability (HA).

Each of the first two PS6000 arrays (Figure 10) consists of 16 hard disk drives each drive having a 250 GB capacity giving a total capacity of 4 TB. A third array, not shown, has a total capacity of 9.6 TB.

Cat 6 **c**ables (yellow and red) connect the 8 ports on the storage devices to switches. The black cables provide power. Cable connections are distributed across multiple switches to form a mesh network.

Cable connections from the 8 Ethernet interfaces available on each array are distributed across multiple switches in a mesh network topology similar to that shown for the physical server's cable connections to the switches (Figure 5).

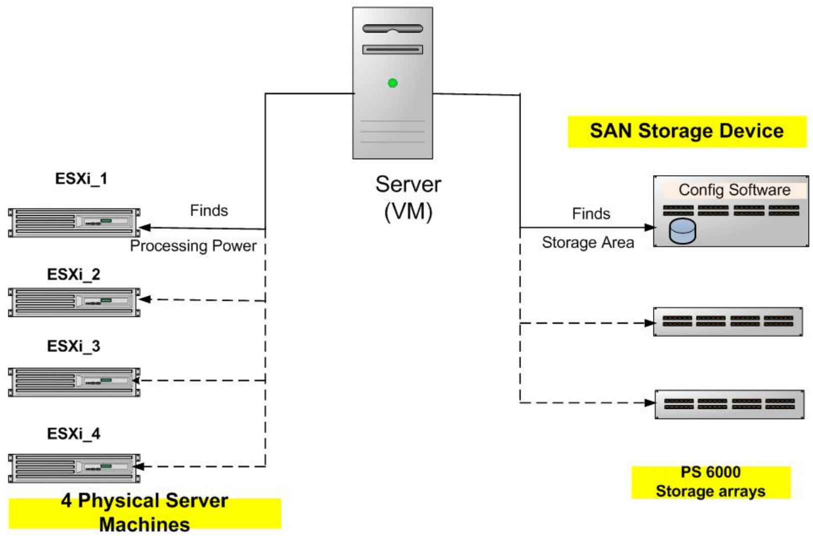

Figure 11 Processing power and storage for each VM

Storage management is provided by the server console (Figure 6) which configures software on the SAN storage devices. This software provides the abstraction of the storage device as seen by each VM with each VM then appearing on a physical area of the storage device.

Each VM receives processing power from the cluster of physical servers seen as a single system (Figure 11) and will access a data storage area provided on the cluster of storage devices. Configuration software is provided on the storage devices (Figure 12).

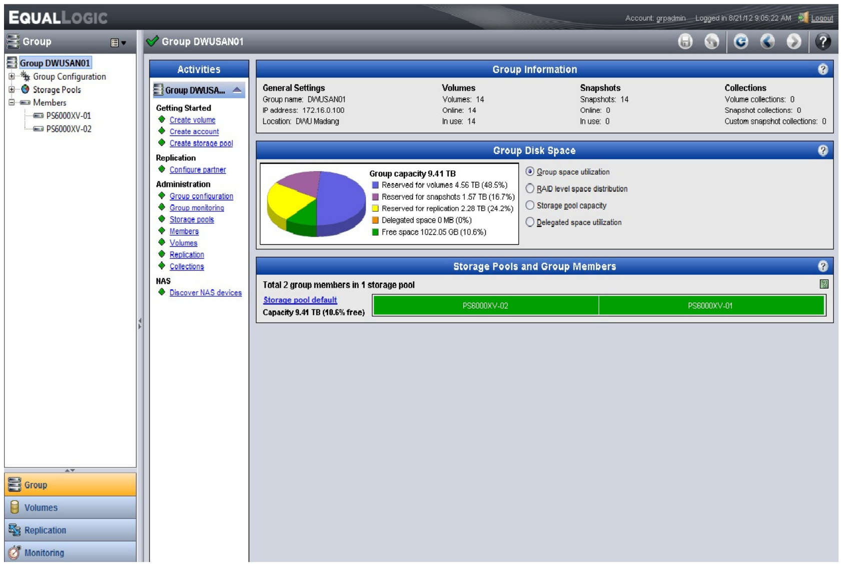

Figure 12 Equal Logic storage management software

Equal Logic software to provide storage management (Figure 12, is also located on the storage devices. It allows administration of virtual partitions (named as Volumes) on PS 6000 storage devices (named as Members). The software allows administration of virtual partitions (Volumes LHS panel) on PS 6000 storage devices (Members LHS panel).

DWU Data Centre Network

The three main components of the Data Center, the core and distribution switches, the physical and virtual servers and the storage array can now be considered as an integrated mesh network (Figure 13) providing services to client devices and also Internet connection for the whole network of clients and servers.

Services to client machines are provided from the core switch via two distribution switches which serve the underground campus fiber-optic cables connecting to staff and students throughout the main campus and beyond to other local campuses. Figure 13 summarises the topology of the main Data Centre components: Internet gateway, core and distribution switches, and physical servers as part of a local version of a server cloud, and the SAN.

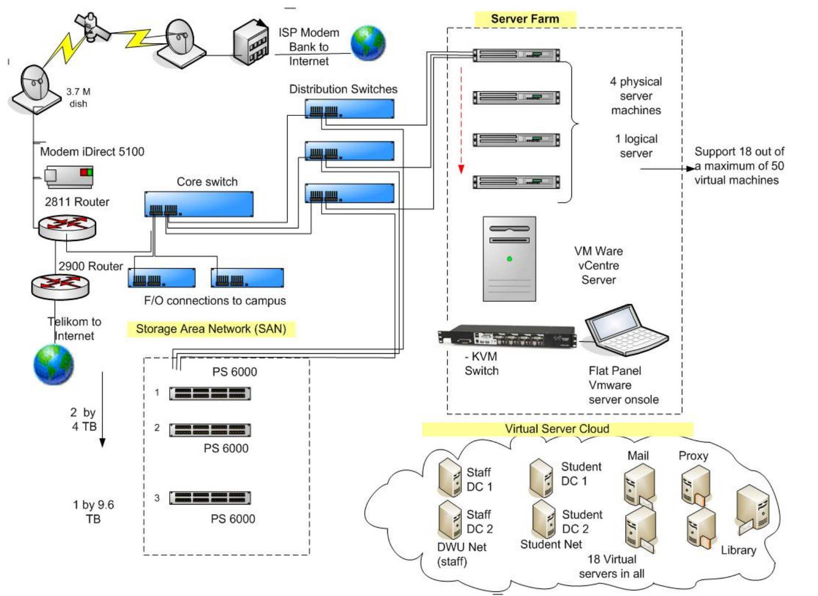

Figure 13 Summary logical diagram of the DWU Data Centre

Of interest for motivating cloud computing (both local and remote), it should be noted that with even just four physical servers used in the Data Centre provide processing power and available memory far in excess of present needs (Figure 14). Total available processing power (64 MHz) and available memory (220GB) from 4 physical servers in the Data Centre far exceeds the present demand even with 35 VMs. Although an improvement on the use of separate physical machines for each required server this is clearly shows we still have underutilisation and uneconomical provision.

With remote cloud computing there is no possibility overprovision of computing resources from the perspective of an end-user organisation, because billing will always be on the basis of user demand.

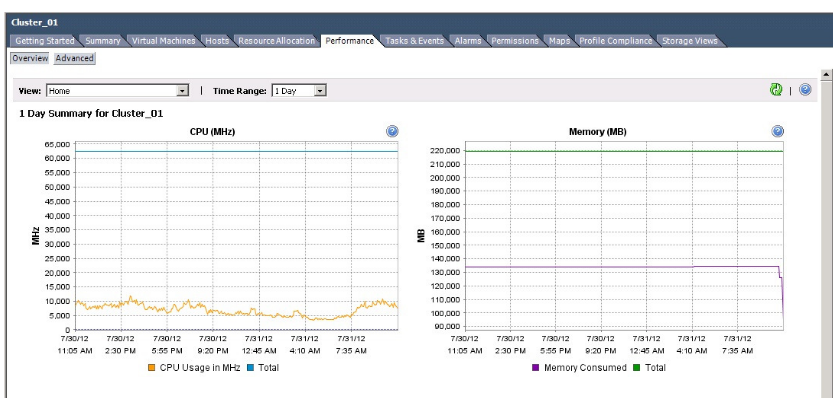

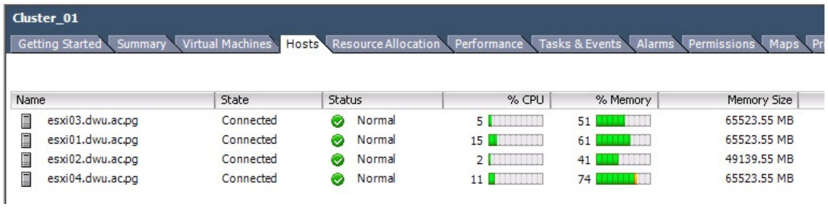
Cluster_01

Getting Started | Summary | Virtual Machines | Hosts | Resource Allocation | Performance | Tasks & Events | Alarms | Permissions | Maps | Pr

Name	State	Status	% CPU	% Memory	Memory Size
esxi03.dwu.ac.pg	Connected	Normal	5	51	65523.55 MB
esxi01.dwu.ac.pg	Connected	Normal	15	61	65523.55 MB
esxi02.dwu.ac.pg	Connected	Normal	2	41	49139.55 MB
esxi04.dwu.ac.pg	Connected	Normal	11	74	65523.55 MB

Figure 14 Demand for available processing power

Cloud computing

The essential feature of the DWU Data Centre is the decoupling of the server services from the physical machines by hypervisor software running directly on hardware. This is a reduced form of an operating system known as a microkernel, where the kernel is the part of the operating system which manages hardware to provide access to applications.

It has been noted that hardware emulation using hypervisor software without an intervening operating system (OS) is only one of a number of approaches to server virtualization. Another possibility is OS virtualization where each VM is provided with an environment in which it appears to have access to a separate instance of the operating system. Multiple guest OSs appear to be running on a single host OS.

VMs also view abstractions of the physical clustered storage device, another example of decoupling. Once decoupling has occurred, the physical devices providing processing power and storage can be located remotely, and so not necessarily even on the same campus receiving the virtual services.

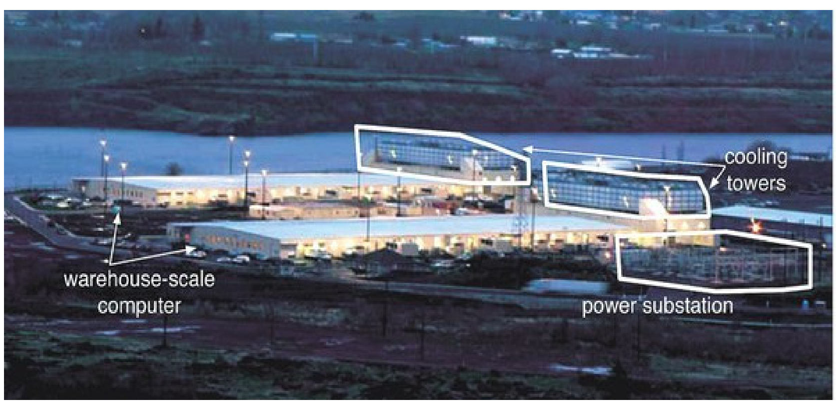

Figure 15 Data center and warehouse scale computer

A remotely located computer processing and data storage centre providing regional services on an on-demand basis would be cloud computing properly so called (Figure 15). A depiction of a remote data center and warehouse scale computer (WSC, Abts & Felderman) which could provide regional computer processing power over a WAN. The computer could consist of tens of thousands of hosts each with one or more processors packaged into racks and allocated as clusters.

The computing platform of interest in a data centre providing cloud computing is a warehouse sized building full of computers. Large portions of the hardware and software resources in such a facility must work in concert to efficiently deliver required levels of Internet service performance. The datacenter can be viewed as one massive warehouse-scale computer (Barroso & Hölzle).

The main components of a data centre (Figure 16) include a UPS, power generation systems, cooling systems, air conditioning units (CRAC), processing power and data storage systems.

End-user remote access would then be by means of a Web browser running on a device with minimum processing power (thin client) including mobile devices. A grid connecting clients to such a remote warehouse sized data centre with its own cooling towers and power substation (Abts & Felderman, 2012) would require a high bandwidth wide area data network WAN). Computer processing power would thus become another public utility along with electricity, water, gas and sewerage which also provide services on demand and are paid for on an as-used basis.

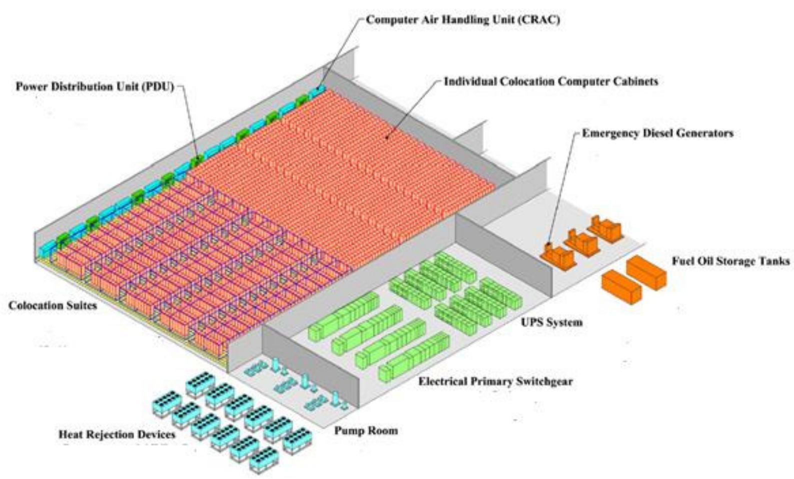

Figure 16 Main components of a cloud computing data center

Whilst the present PNG Telikom microwave network (Figure 17 Kim & Anderson, 2012) would not provide sufficient data bandwidth, it is not inconceivable that the fiber-optic network presently being installed on PNG Power poles (Optical Ground Wire, OPGW,) along the National Highway could, sometime in the future, provide the necessary bandwidth to allow thin client network and Web browser access to remote cloud processing power and data storage, thereby necessitating a radical restructure of the whole IT industry in PNG as is happening in other parts of the world.

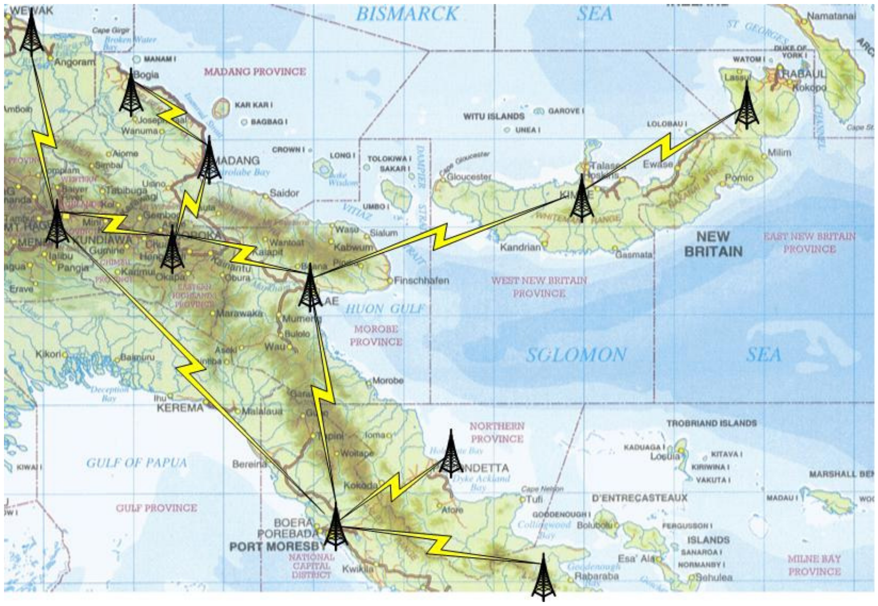

Figure17 PNG national telecommunication microwave network

Conclusion

This paper has described a recent innovation of a virtualised Data Centre at the DWU campuses where a central hub for PNGARNet is also provided. The innovation provides multiple virtual servers on a cluster of a much smaller number of physical servers.

The advantages of this new configuration include reduction in capital expenditure on hardware (fewer physical servers), reduction of server footprint allowing more scarce building space for offices, lower power consumption given reduction in required hardware, lower maintenance costs, and finally greater utilisation of available computing resources.

This level of decoupling is sometimes referred to as local cloud computing. The logical further challenge would be to pool all processing power in a national or regional warehouse scale computer providing all necessary processing power on an as-needed basis.

This regional large scale processing source could generate its own electrical power to replace, or at least supplement, a sometimes unstable national power grid. Processing power would then be available on a 24/7 basis as is needed, for example, for Web servers providing e-commerce in a globalised world of business activity. Server downtime should be substantially reduced

Replacing today's PCs and laptops in such a scenario, much simpler devices providing only keyboard, mouse, monitor and network connection together with minimal processing power (thin clients) would be sufficient.

Acknowledgements

The DWU Data Centre in its present form was constructed in 2010 by Mr Chandana Silva, Head, Dept. ICT Services. Mr Silva also provided reviews of the paper and permission to publish. Advice on network design for the diagramming was provided by Mr Charith Silva and Mr Kevin Kosala, both of DWU ICT Services. However, any errors of fact or omission rest entirely with the author. Ms Melanie Livingston and Ms Gloria Arabagali, both of DWU, provided assistance with drafting of the diagrams.

References

Abts, D, & Felderman, B. (2012). *A Guided Tour of Data-Center Networking,* Communications of the *ACM*, Vol 55, No. 6. pp 44- 51

Anderson, P.K. & Kim, J. (2011). Sydney-Guam PIPE Pacific Cable: New international gateway for PNG via Madang. In *Contemporary PNG Studies, DWU Research Journal* Vol. 15.

Andrews, J. (2006). *Managing and Maintaining Your PC*, (5th ed.). US: Thomson, Course Technology.

Barroso, L. & Hölzle, U. (2009). The Datacenter as a Computer: An introduction to the design of warehouse-scale machines, in *Synthesis Lectures on Computer Architecture*, Mark D. Hill, (Ed) University of Wisconsin, Madison. *http://www.morganclaypool.com/doi/pdf/10.2200/S00193ED1V01Y200905CAC006 Available: 22-Jul-12*

Golden, B. (2008). *Virtualisation for Dummies*, US: Wiley Publishing Inc.

Kim J. & Anderson, PK. (2012). A Model to close the digital divide in Papua New Guinea in *Contemporary PNG Studies: DWU Research Journal*, Vol. 15, in press, PNG: DWU Press

Velte, T.J. & Velte, A.T. (2007). *Cisco: A Beginner's Guide*, 4th ed. NY: McGraw-Hill, pp 411-420.

Waldspurger, C. & Rosenblum M. (2012). I/O Virtualisation, *Communications of the ACM*, Vol. 55, No. 1, pp 66-72.

Author

Associate Professor Peter K. Anderson is a senior lecturer at DWU and is foundation head of the Department of Information Systems where he specialises in data communications. He holds a PhD in thermodynamic modelling from the University of Queensland. His research interests include documenting major technology developments in PNG. Email: panderson@dwu.ac.pg

Glossary

CMS	Course Management System
GB	Gigabyte (10^9)
HA	High Availability
ICT	Information Communications Technology
I/O	In Out
ISP	Internet Service Provider
KVM	Keyboard Video Monitor
OPGW	Optical Ground Wire
OS	Operating system
PNGARNet	PNG Academic & Research Network
PPC	PIPE Pacific Cable
RAM	Random Access Memory
SAN	Storage Area Network
TB	Terabyte (10^{12})
VM	Virtual Machine
VMM	Virtual Memory Monitor
VLAN	Virtual Local Area Network
VPN	Virtual Private Network
WAN	Wide Area Network
WOW	Windows on Windows
WSC	Warehouse-sized Computer