

Virtualized Network Services

by Paul Lawrence

FILE	3346_PAUL_LAWRENCE_VIRTUALIZED_NETWORK_SERVICES_100686_1 409966520.DOCX (8.64M)		
TIME SUBMITTED	14-FEB-2018 11:22PM (UTC+0000)	WORD COUNT	12231
SUBMISSION ID	83196716	CHARACTER COUNT	67522

Virtualised Network Services

Paul M Lawrence (P152216) Undergraduate Student, Mcomp Computer Networks
University of Wales Trinity St David, Swansea, Mount Pleasant Campus
plawrence.152216@student.uwtsd.ac.uk

Abstract – This report aims to provide a practical and theoretical introductory guide to virtualised network services and network architectures, the role they play within modern networking and the technologies used to provide high availability load balanced networks.

Keywords-component; ARP, DHCP, DNS, Virtualisation, Servers, Clustering, RAID, Routing, RIP, Hypervisor, LAN, WAN

I. INTRODUCTION TO NETWORK SERVICES AND VIRTUALISATION

Network services are at the heart of any businesses network infrastructure, their key roles are to deliver network services to clients within a Local Area Network (LAN) [1] or across a Wide Area Network (WAN) [2], the services they provide range from application services to fundamental administration operations i.e. IP address allocation, end user authentication, domain creation and authorisation, network segmentation, redundancy and much more. [3]

Typical legacy servers are very resource intensive, expensive to build, maintain and administer. Within modern day computing, services are in high demand with more and more end users connecting to the internet every day to access; video streaming, social media, banking etc these services can be resource intensive.

Taking this into consideration the typical legacy way of providing services are inefficient and have become economically unviable within large production networks. To combat this problem Hypervisors [4] were developed to virtualise environments in order to efficiently provide a range of services utilising all available hardware and converging administration tasks.

Hypervisors typically contain multiple Operating System (OS) images called 'virtual machines' they typically consolidate network services and makes them easier to manage utilising all hardware resources effectively and efficiently in the process. This solution is often referred to as virtualisation, businesses have adopted this model in abundance to benefit from all the advantages hypervisors have to offer.[4]

To demonstrate the technology a 'Tier-1' hypervisor, 'Windows Hyper-V' will be utilised within a lab environment using two 'Windows server 2012 RS' OS images named 'Swansea' and 'Cardiff'. A client OS image (windows 7) named 'Lab' will also be installed to emulate the services being provided within a production network. The 'Cardiff' server will be a Wide Area Network (WAN) Customer Edge (CE) facing server which will be configured for routing capabilities between the VM's as well as providing Dynamic Host Configuration Protocol (DHCP) and Domain Name System (DNS) services.

The environment the lab will be installed in consists of two servers that are configured as a failover cluster; they provide redundancy and load balancing which will be discussed later in the report. The servers are aptly named 'Babbage' and 'Von Neumann', they are accessed through a virtual domain called 'tesla' which is the virtual default gateway for the cluster.

The 'Hyper-V' is accessed through 'Windows' System Centre Virtual Machine Manager (SCVMM) [5] client application. 100GB of virtual hard disk partition space has been allocated for the lab. Topology shown in *figures A*.

(Full size image can be found in Appendix-A)




Figure A: Depiction of Hyper-V Lab Topology with Legend [6]

II. VIRTUAL INTERFACE CONFIGURATION AND CONNECTIVITY

The two server images and the client OS images previously discussed in the report have been successfully installed within the Virtual Machine (VM) Manager shown in *figure 1*.




Figure 1: OS Images Installed on VM Manager

The virtual Network Interface Cards (NIC) are installed for each of the VM's and an additional card installed for the virtual router VM (Cardiff) as these interfaces will be the network gateways. The virtual **caballing** is also configured for connectivity between the VM's emulating physical caballing and switching within a physical environment shown in **figures 2-5**.




Figure 2: Cardiff Server 1st Virtual NIC Card installed and caballed to LAB NIC (Network 1).




Figure 3: Cardiff Server 2nd Virtual NIC Card installed and virtually caballed to Swansea Server NIC (Network 2).




Figure 4: Swansea Server Virtual NIC Card installed and virtually caballed to Cardiff Server NIC (Network 2).




Figure 5: Lab Client Virtual NIC Card installed and virtually caballed to Cardiff NIC (Network 1).

The interfaces are renamed relating to their respective connections for easier management of the network shown in **figure 6-8**.



Figure 6: Cardiff Server Interfaces Named



Figure 7: Swansea Server Interface Named



Figure 8: Lab Client Interface Named

Static IPv4 addresses and Subnet Masks are applied to the interfaces shown in **figures 9-12** to identify each VM's interface(s) corresponding sub-network (subnet) within the Local Area Network (LAN).




Figure 9: Lab Gateway Interface IPv4 Address and Subnet Mask Configuration at Cardiff Server




Figure 10: Swansea Gateway Interface IPv4 Address and Subnet Mask Configuration at Cardiff Server




Figure 11: Cardiff Connection Interface IPv4 Address, Subnet Mask and Default Gateway Configuration at Swansea Server



Figure 12: Cardiff Connection Interface IPv4 Address, Subnet Mask and Default Gateway Configuration at Lab Client

Using the 'Windows' command Prompt a verification command 'ipconfig /all' is entered within each of the VM's to view information about the interfaces e.g IP Address, Subnet Mask, Default Gateway to ensure that the correct addresses have been applied to each of the interfaces shown in figures 13-15.




Figure 13: Interfaces IP Address and Subnet Mask Verification using command prompt command 'ipconfig /all' at Cardiff Server




Figure 14: Interface IP Address, Subnet Mask and Default Gateway Verification using command prompt command 'ipconfig /all' at Swansea Server




Figure 15: Interface IP Address, Subnet Mask and Default Gateway Verification using command prompt command 'ipconfig /all' at Lab Client

Within a Windows server environment additional services have to be installed as needed, this minimises the resources being used unnecessary i.e. CPU, RAM, HDD. All new services are installed at the services manager dashboard using an install wizard shown in figure 16.




Figure 16: Server Manager Dashboard for adding roles and services

A virtual routing role is installed and configured at Cardiff server called; Routing and Remote Access service (RRAS). This service provides a virtual routing capability to the server that can dynamically create routes to each of the LAN's i.e Swansea and Lab for connectivity and reachability between each of the VM's using the dynamic routing protocol; Routing Information Protocol (RIP) shown in figures 17-20.




Figure 17: Remote Access Role Selected to be Installed at Cardiff Server




Figure 18: Routing and Remote Access Setup Wizard, LAN Routing Selected




Figure 19: Routing and Remote Access Control Panel, RIP Selected on the left-hand window pane. 'Swansea' and 'Lab' interfaces selected to be used within RIP dynamic routing protocol as these are the two gateways for the two subnets in which need to communicate through 'Cardiff'.




Figure 20: Within the RIP protocol panel; Swansea Properties were selected and configured for RIPv2 outgoing broadcasts and to accept both RIPv1 and v2 incoming Packets. Periodic update Mode has also been used. This step was also repeated for the Lab Gateway interface within the RIP properties.

After configuring the RIP protocol, verification of connectivity between all VM's will be performed by using the 'ping' command within the windows command prompt shown in figures 21-23.




Figure 21: Ping from 'Cardiff' server to 'Lab' and 'Swansea' to verify connectivity via RIP are both successful




Figure 22: Ping from 'Lab' client to 'Swansea' to verify full connectivity via RIP is successful




Figure 23: Ping from 'Swansea' server to 'Lab' client to verify connectivity via RIP is successful

Finally, a 'route print' command is issued from the command prompt at 'Cardiff' to view the virtual routing table created at Cardiff in order to verify the routes (Subnets) that RIP has learned dynamically shown in figure 24.




Figure 24: CMD prompt 'Route Print' output verifying all routes are present within the 'Cardiff' server routing table that were dynamically learned via RIP

Identify the benefits of Type 1 Hypervisor software services and briefly evaluate some suitable products.

As discussed earlier in the report a Tier-1 Hypervisor also referred to as a 'bare metal' installation is a low-level small OS that runs directly on top of a systems hardware that fully

utilises all the systems resources i.e. Solid-State Drives (SSD), RAM, CPU etc. A 'Tier-1' configuration also minimises latency between the VM image and the hardware response time. Once configured Operating System (OS) images can be installed within the hypervisor manager and hardware can be allocated to each OS as needed by the systems administrator. [4] [7]

This method is very efficient and economical as all the hardware is fully utilised. VM's are easier to manage/troubleshoot within a production network using a hypervisor as opposed to legacy standalone servers which take up a lot of memory, waste hardware and could be situated in different locations. [4]

There are many 'Tier-1' hypervisor products available on the market some offering additional functionality for a higher price. The most popular are: *Research conducted Dec' 2017*

- **VMware ESX and ESXi;** This product is at the mid-range of the monetary scale. It offers advanced enterprise features and is also scalable which is attractive to larger companies. VMware currently dominates virtualisation within today's market due to its price compared to Microsoft Hyper-V for similar functionality. [8] [9]
- **Microsoft Hyper-V;** This product is at the high end of the monetary scale with similar functionality compared with VMware but offers higher integration with 'Microsoft' products i.e. 'Windows Server', 'Windows Clients' etc which is attractive to company's running a 'Windows' environment. [10] [9]
- **Citrix XenServer;** This product was one of the first on the market and is currently at the lower end of the monetary scale as it offers less functionality. This product is suitable for Small-to medium sized companies that only provide internal services for a small amount of end users. [11] [9]

Unlike a 'Tier-1' hypervisor a 'Tier-2' hypervisor is software installed within an existing OS that must access the hardware through the native OS; this architecture creates an extra layer between the VM and the hardware maximising latency. The native OS also uses a lot of system resources needlessly. This method of virtualisation does not effectively utilise the available hardware making it inefficient within a production network. Similar to the 'Tier-1' hypervisor once configured OS images can be installed within the VM manager. A 'Tier 2' hypervisor is relatively inexpensive and is most commonly used by end users wanting to run multiple OS's on their devices for various purposes. [7]

Define the role of Routers in the Network Architecture, what types can be used?

A Router is responsible for forwarding packets at Layer-3 of the Open Systems Interconnections (OSI) [12] model between networks within a WAN or sub-network within a LAN.

A routers primary role is defined by how it is placed within a LAN or a WAN's architecture i.e core or distribution layer. The most commonly used routers within a LAN are:

Core routers or backbone routers as they are also referred to: These routers are high end routers that's primary goal is to forward packets from one subnet to another as fast as possible without providing any additional functionality such as DHCP etc so that all resources are focused on routing packets across the backbone of a network. These routers are usually aggregated within a full mesh design to eliminate a single point of failure and also utilise etherchannel (port aggregation) for added bandwidth. [13]

MPLS Routing; Modern day core routers have adopted a Multiple Protocol Label Switching (MPLS) solution for delivering packets, this type of routing is a lot more efficient than conventional packet routing. The mechanics of MPLS work by instead of inspecting the header of each packet to learn the route, a label is applied to each packet on a router and then the packets are routed to the neighbour router that the label corresponds with saving bandwidth and hardware resources at the router which in turn substantially increase throughput. MPLS is described as Layer-3 routing at Layer-2 (Switching) speeds. [14]

ISP router or Customer Edge (CE) routers as there also referred to: Are routers that sit on the edge of a LAN and are the gateway to the WAN which is owned by an Internet Service Provider ISP. These routers often have security or firewall capabilities as they are the first line of defence against malicious traffic coming from the WAN. [15]

Virtual routers within a virtualised Hyper-V environment are used by small businesses with a small amount of end users and minimal services/traffic. [16]

Identify and discuss the usage of common Routing Protocols, what are the associated benefits of using Routing Protocols?

A dynamic routing protocol is defined as an algorithm that calculates a best path for a packet to traverse within a network autonomously. It communicates with neighbor routers configured with the same routing protocol using multicast messages. [17]

There are three common Interior Gateway Protocols (IGP) [18] that are used within networks around the world, they each have their advantages and disadvantages:

1. **Routing Information Protocol (RIP):** This protocol has an administrative distance (AD) of 120 and is categorized as a distance vector protocol. It calculates routes based on hop counts using the 'Bellman-Ford' algorithm and is the most basic of the IGRP's. Its advantages are that it is simple to configure/troubleshoot, and convergence time is relatively fast. However, this protocol has its limitations because it has a maximum hop count of 15 routers and any router after that is considered unreachable. RIP does not consider Metric within its calculations so the path it selects may not

always be the fastest path. It only uses a routing table so does not know the entire network topology only neighbor routers. Periodic updates of the entire routing table regardless of any topology changes are also a disadvantage as this uses unnecessary router resources and bandwidth within the network. This protocol is generally used within small businesses with a small amount of end users and an IT team with limited networking knowledge. [19]

2. Enhance Interior Gateway Routing Protocol (EIGRP):

EIGRP is categorized as an advanced distance vector protocol with an AD of 90. It uses the Diffusing Update ALgorithm (DUAL) and is considered the fastest IGRP due to the convergence time and the ability to adapt quickly to any changes within the network topology. It calculates best path using metric based on 'K' values; bandwidth, delay, load, reliability and Maximum Transmission Units (MTU). EIGRP is able to pre-compute expected loops within the network using Reported Distance (RD) from neighbor routers. Using this information and the local metric Bandwidth speed from the cable throughput, EIGRP can calculate if a destination is reachable this is called the Feasible Distance (FD) which helps to determine the best path. EIGRP is comprised of two tables; a routing table and a topology table, because of this the protocol is aware of the entire network topology and is able to adjust its calculations based on these tables storing backup routes in the event of router becoming unavailable. EIGRP only updates when there is a change within the topology and only updates this information not the entire routing table. EIGRP is not limited by a maximum number of routers and has advanced features which can fine tune a network depending on its architecture. EIGRP also has the ability to summarize routes and load balance packets based on the available metric. This protocol is used by mid-to-large sized companies. The disadvantage of EIGRP is that up until 2016 it was a Cisco proprietary protocol so only companies who owned only Cisco equipment would use this protocol and most vendors are yet to adopt this protocol. It also can be complex to troubleshoot and would need to be managed by a knowledgeable network team. [20]

3. Open Shortest Path First (OSPF):

OSPF has an AD of 110 and is categorized as a link state protocol. It calculates routes using the 'Dijkstra' algorithm based on cost which is similar to EIGRP's metric 'K' values. OSPF is a very robust and complex protocol, it segments the network into areas for management; Area 0 being the backbone of the network. It uses two master control routers within a network which are elected based on highest router ID or priority, to store information about the topology and update other routers within the network. These routers are called the Designated Router (DR) and the Backup Designated Router (BRD), the BDR is used as the master router if the DR goes offline. Unlike RIP and EIGRP, OSPF stores all information within its DR and BDR topology databases about every router within the network to make more complex decisions about the network, this is the primary difference between Distance

Vector and Link State protocols as only the DR makes decisions based on the topology table. OSPF is mostly used in large enterprise networks and is administered by a network team with extensive networking knowledge. The disadvantages of OSPF is that it's a very complex protocol and takes a long time to converge using a lot of resources during this time which can temporarily debilitate the network during convergence. High performance routers are recommended when running this protocol on a large network for this reason. [21]

4. Static Routes; IGRP's are preferred within a network

over static routes because each route must be manually configured and tailored to forward traffic to the best path. In modern networks this approach is unfeasible and a waste of resources as advanced IGRP are autonomous and self-governing as discussed earlier in the report. Static routes are generally used as backup redundancy routes within the backbone of a network in the event of the IGRP failing so that a network still has minimal functionality whilst the network engineers troubleshoot the issue. [22]

5. Border Gateway Protocol (BGP):

This protocol is categorized as an Exterior Gateway Routing Protocol (EGP) [18] it used mainly by Internet Service Providers (ISP) to route traffic over the WAN (Internet) between autonomous systems (AS). [23]

Define the role of Routing Tables, what are they. Breakdown and analyze the information a Routing Table provides us with.

A routing table is a database housed within a router that is used to store all paths to known networks within a LAN shown in figure 25.

Network	Network Mask	Next Hop	Interface	Metric
0.0.0.0	0.0.0.0	0.0.0.0	0.0.0.0	0
127.0.0.1	255.0.0.0	127.0.0.1	127.0.0.1	1
192.168.1.0	255.255.255.0	192.168.1.1	192.168.1.1	1
192.168.1.1	255.255.255.255	192.168.1.1	192.168.1.1	1
192.168.1.2	255.255.255.255	192.168.1.2	192.168.1.2	1
192.168.1.3	255.255.255.255	192.168.1.3	192.168.1.3	1
192.168.1.4	255.255.255.255	192.168.1.4	192.168.1.4	1
192.168.1.5	255.255.255.255	192.168.1.5	192.168.1.5	1
192.168.1.6	255.255.255.255	192.168.1.6	192.168.1.6	1
192.168.1.7	255.255.255.255	192.168.1.7	192.168.1.7	1
192.168.1.8	255.255.255.255	192.168.1.8	192.168.1.8	1
192.168.1.9	255.255.255.255	192.168.1.9	192.168.1.9	1
192.168.1.10	255.255.255.255	192.168.1.10	192.168.1.10	1
192.168.1.11	255.255.255.255	192.168.1.11	192.168.1.11	1
192.168.1.12	255.255.255.255	192.168.1.12	192.168.1.12	1
192.168.1.13	255.255.255.255	192.168.1.13	192.168.1.13	1
192.168.1.14	255.255.255.255	192.168.1.14	192.168.1.14	1
192.168.1.15	255.255.255.255	192.168.1.15	192.168.1.15	1
192.168.1.16	255.255.255.255	192.168.1.16	192.168.1.16	1
192.168.1.17	255.255.255.255	192.168.1.17	192.168.1.17	1
192.168.1.18	255.255.255.255	192.168.1.18	192.168.1.18	1
192.168.1.19	255.255.255.255	192.168.1.19	192.168.1.19	1
192.168.1.20	255.255.255.255	192.168.1.20	192.168.1.20	1
192.168.1.21	255.255.255.255	192.168.1.21	192.168.1.21	1
192.168.1.22	255.255.255.255	192.168.1.22	192.168.1.22	1
192.168.1.23	255.255.255.255	192.168.1.23	192.168.1.23	1
192.168.1.24	255.255.255.255	192.168.1.24	192.168.1.24	1
192.168.1.25	255.255.255.255	192.168.1.25	192.168.1.25	1
192.168.1.26	255.255.255.255	192.168.1.26	192.168.1.26	1
192.168.1.27	255.255.255.255	192.168.1.27	192.168.1.27	1
192.168.1.28	255.255.255.255	192.168.1.28	192.168.1.28	1
192.168.1.29	255.255.255.255	192.168.1.29	192.168.1.29	1
192.168.1.30	255.255.255.255	192.168.1.30	192.168.1.30	1
192.168.1.31	255.255.255.255	192.168.1.31	192.168.1.31	1

Figure 25: Windows Server 2012 RS Routing Table

The information that is stored within a 'Windows' routing table is as follows:

- **Network Destination;** The IP address of a known network.
- **Netmask;** The subnet mask of the network in the 'Network Destination' field.
- **Gateway;** The default gateway used to reach the network within the 'Network Destination' field.

- **Interface;** The IP address associated with the exit interface that is used to reach the network in the 'Network Destination' field of the routing table.
- **Metric;** The metric calculated to reach the network within the 'Network Destination' field. Metrics are only used if there's multiple paths to a destination in which case the same destination would be present twice with a different metric and exit interface, in such a case windows would choose the lowest metric number within the metric field as this is determined as the best path and the alternative route would be a redundant route.

What is the Subnet Mask, what is the significance of this in the Routing Table?

Every Internet Protocol version 4 (IPv4) Address has a subnet mask associated with it. They are layer-3 addresses that are used only by routers to identify smaller networks called subnets within an IP address range of a LAN which is comprised of subnets.

A subnet mask is used for segmenting and identifying a networks IP address range into smaller networks called subnets to preserve IP addresses and for organisation of a network for LAN management, security and administration purposes.

A subnet mask determines how many available hosts there are within a subnet and also how many subnets are available within an IP address range.

The subnet mask within a routing table is used to identify a particular subnet within a network so that packets can be forwarded to the correct subnet in which the destination host is a member.

III. ARP, PACKET FILTERING AND REMOTE ACCESS

Understanding Address Resolution Protocol (ARP) [24] is a fundamental skill that network administrators need to learn to be able to analyse/troubleshoot the network at Layer-2 of the OSI model within IPv4.

The ARP protocol is used to map Layer-2 Media Access Control (MAC) addresses to layer-3 IPv4 addresses. For this reason, ARP in IPv4 is considered the bridging protocol between layer-2 and Layer-3.

Within this lab, a File Transfer Protocol (FTP) service will be configured to demonstrate how users can access files housed on a server within a network locally. IP packet filtering will also be configured to control what protocols and subnets can access the FTP server.

Packet filtering is a key component to securing a network both locally and remotely in order to securely provide accessibility to a network and control what traffic enters the network at Layer-3.

Using the 'Windows' command prompt the command 'arp -a' is entered to inspect the ARP table present at the Cardiff server to learn what IP addresses have been mapped to Media Access Control (MAC) addresses at Layer-2 within the subnets associated with each of the interfaces shown in figure 26.




Figure 26: Cardiff Server ARP table

To enable remote file access at the Swansea server the web server (IIS) role is selected and then the FTP service is installed using the server manager dashboard shown in figure 27-28.




Figure 27: Web Server (IIS) role is selected at Swansea Server




Figure 28: FTP Server role selected to be installed at Swansea Server

Once the FTP service is installed the next stage is to configure the service for remote access within the network.

The first step is to create a local username and password for FTP service called 'ftpuser' within the computer management screen. Change user password on next login is deselected because it is going to be used as an FTP authentication account shown in figure 29.




Figure 29: FTP user account created within computer management screen

The next step is to activate the FTP server and create an FTP site called 'test' and local file path within the Internet Information Services (IIS) manager shown in figure 30.




Figure 30: IIS FTP configuration site name and physical path created

After creating the FTP site basic configuration is applied using the install wizard. Unassigned IP address bindings is selected and default FTP port 21. Secure Socket Layer (SSL) network certificate security is deselected for this basic configuration shown in figure 31.




Figure 31: FTP wizard basic configuration no SSL selected

Basic Authentication and Authorisation is selected for the purpose of this lab demonstration allowing only the FTP user account located at the Swansea server created in the previous step figure 29 (ftpuser) this is shown in figure 32.




Figure 32: Authentication and Authorisation basic configuration

A welcome banner is configured for remote user's login on to the FTP server using the FTP messages screen shown in figure 33.




Figure 33: FTP service welcome banner configured

A test document (testftpdnload.txt) is created within the FTP root folder at Swansea server shown in figure 34.




Figure 34: Test file created at FTP root folder

At the Cardiff server a directory is created at the 'C:' called 'ftpdnloads' this is where all downloaded FTP files will be stored shown in figure 35.




Figure 35: FTP download folder created at Cardiff server

To verify the FTP server is working correctly a connection will be made from Cardiff server using the 'Windows' command prompt command 'ftp [Swansea IP address]' to connect to the Swansea server. Once connected the 'lcd C:\ftpdnloads' command is issued to change the local FTP directory to the directory previously created. After the correct local directory has been selected the test document created at Swansea FTP root folder will be downloaded onto the Cardiff server to the specified directory using the 'get testftpdnloads.txt' command shown in figure 36.




Figure 36: FTP connection from Cardiff to Swansea server, test file downloaded and welcome banner displayed correctly

At the 'ftpdnloads' directory at Cardiff server the test file created at Swansea has successfully been copied to Cardiff server shown in figure 37.



Figure 37: Test file successfully copied to Cardiff servers local FTP directory

Similar to the FTP download step, a directory called 'ftpupload' is also created at Cardiff with a test file named 'testftpupload.txt' to test that files can successfully be uploaded to the Swansea server remotely shown in figure 37A.




Figure 37A: FTP upload directory created at Cardiff server

Similarly to the FTP download step the 'lcd' command is issued followed by the path of the upload directory. Once the ftp prompt is in the 'ftpupload' directory the command 'put' is issued followed by the name of the test file. This command will upload the file to the FTP root directory at the Swansea server shown in figure 38.




Figure 38: FTP upload to Swansea server from Cardiff server

To verify the file successfully uploaded to Swansea server the FTP root directory of the Swansea server is accessed locally shown in figure 39.




Figure 39: FTP root directory at Swansea, test file uploaded successfully

As previously discussed IP packet filtering is essential for administration control of inbound and outbound connections within the network. To demonstrate how they work IP packet filters will be applied to the Cardiff servers interfaces as this is the virtual router for the network and the interfaces are the gateways for the subnets within the network. Filters are configured from within the routing and remote access control panel by selecting the interface in which the filters are to be applied.

The first filter will be an inbound filter applied at the 'Lab' client interface at Cardiff. The filter will not allow Internet Control Message Protocol (ICMP) traffic [most commonly referred to as 'ping'] from the 'Lab' client to the Swansea server routed through Cardiff. Pings to Swansea inbound from the 'Lab' interface will be rejected after the filter is applied. A subnet mask of 255.255.255.255 is configured which translates as; match IP address exactly, this is applied so that only one source and destination host is included within the filter not the entire subnet shown in figures 40-42.




Figure 40: Lab interface selected at Routing and Remote Access screen (IPv4 > General), and inbound filters button selected within the 'Lab' interface screen




Figure 41: At the inbound filter screen within lab interface Source IP, destination IP and protocol configured for filter




Figure 42: After filter is configured it is verified at next screen and configured to allow all packets except those that meet the criteria of the IP filter created

To verify the filter is working correctly a ping command is issued from the 'Lab' client to the Swansea server which is unsuccessful because this packet meets the criteria of the inbound IP filter created. A second ping to Cardiff's Swansea interface is then issued which is successful because this request does not meet the criteria of the IP filter shown in figure 43.




Figure 43: Lab ping to Swansea unsuccessful. Lab to Cardiff Swansea gateway interface is successful. IP Packet filter is working correctly

An FTP filter 'Outbound' from Swansea interface on the Cardiff server will be configured to block all FTP traffic to Swansea's FTP service. The previous filter will be removed to avoid any possible contradictory filters.

The outbound filter will be configured by selecting outbound tab at the Swansea Gateway interface. The source and destination fields will be left blank to apply this filter to all subnet traffic. Transmission Control Protocol (TCP) Port 21 and 20 will be selected as these are the incoming ports used by the FTP service. Transmit all packets apart from those that meet IP filter criteria will also be selected shown in figures 44-45




Figure 44: Outbound filter configured at Swansea Gateway interface on the Cardiff server




Figure 45: Outbound filter verified and transmit all packets except packet filter criteria selected

An FTP connection from Cardiff to Swansea and Lab to Swansea will be attempted to verify the IP filter is working correctly shown in figures 46-47.




Figure 46: FTP connection unsuccessful from Cardiff to Swansea. IP filter working correctly




Figure 47: FTP from Lab to Swansea unsuccessful. IP filter working correctly

The next filter will be configured to only allow Cardiff to access Swansea's FTP service and all other subnets FTP requests should be blocked. The previous filter will be removed to avoid any possible contradictory filters.

The filter will be an inbound filter applied at 'Lab' interface on the Cardiff server. The same steps will be taken to apply the filter as previously outlined. The Source and Destination will be blanked out so that all inbound traffic meets the IP filter criteria, TCP port 21 will be used (FTP) and Transmit all packets apart from those that meet IP filter criteria will also be selected shown in figure 48.




Figure 48: Inbound Filter Created at Swansea interface on Cardiff server to refuse FTP creates inbound from this interface

To verify the filter is working correctly an FTP request from 'Lab' to Swansea will be issued which should be

unsuccessful, then an FTP request from Cardiff to Swansea will be issued which should be successful due to the logic applied to the IP filter shown in figures 49-50.




Figure 49: FTP request to Swansea server from 'Lab' is unsuccessful as expected




Figure 50: FTP request to Swansea server from Cardiff is successful as expected

The next filter will be an outbound filter and will allow FTP access to all subnets except Cardiff. The previous filter will be removed to avoid any possible contradictory filters.

The source address for the filter will be the Swansea gateway interface at the Cardiff server and the destination address will be the Swansea server. The subnet mask will be 255.255.255.255 as used previous to include only the hosts of the interfaces being used not an entire subnet, TCP port 21 (FTP) will also be selected. shown in figure 51:




Figure 51: Block Cardiff to Swansea FTP request IP filter configuration

Transmit all packets apart from those that meet IP filter criteria will also be selected on the next screen the same as previous examples.

Verification that the IP filter works correctly will be confirmed with FTP requests to the Swansea server from 'Lab' and Cardiff shown in figures 51-52



Figure 52: Lab to Swansea FTP service request successful as expected.




Figure 53: Cardiff to Swansea FTP service request unsuccessful as expected

The IP packet filter worked correctly all FTP connections were allowed apart from Cardiff.

Finally, three IP filters are going to be created in order to block all traffic within the network i.e. ICMP, FTP etc apart from FTP traffic from Cardiff to Swansea

The first filter will be an inbound filter at 'Lab' interface on the Cardiff server which will be configured to reject all incoming traffic. This is achieved by setting all the filter fields to any and applying the 'Receive all packets except those that meet the criteria' option shown in figure 54.




Figure 54: Inbound filter applied at Lab interface to block all traffic entering

The second IP filter will be an outbound filter that is going to be created at the Swansea interface also on the Cardiff server that will be configured to block all traffic apart from FTP traffic from the Swansea interface to the Swansea server. 'Drop all packets apart from those that meet the criteria' has also been selected in this filter shown in figure 55.




Figure 55: Outbound filter created to block all outgoing traffic except for FTP traffic to Swansea server

The third filter will be an inbound filter also applied at the Swansea interface with the destination IP address as the Swansea interface and the source IP address as the Swansea server IP address. It will be configured to block all traffic apart from FTP requests from Cardiff server, this is configured so that Cardiff FTP requests can establish a two-way connection. 'Drop all packets apart from those that meet the criteria' has also been selected in this filter shown in figure 56.




Figure 56: Inbound filter configured at Swansea interface to only allow FTP requests from Cardiff server. All other packets will be dropped

Ping commands are executed to verify the IP filter logic works correctly so that traffic has been blocked between all hosts. An FTP connection between Cardiff and Swansea servers will also be established to verify the IP filter only allows this connection shown in figures 57-59




Figure 57: FTP connection successful from Cardiff to Swansea and ping from Cardiff to Swansea fails as expected




Figure 58: Swansea to Cardiff ping fails as expected




Figure 59: Cardiff to Lab ping fail as expected

Q1). Describe (using appropriate illustrations) how the ARP process operates in Networking Contexts?

As previously discussed in the report, Address Resolution Protocol ARP is used to translate IPv4 addresses to MAC addresses; it works by checking its ARP table for corresponding entries, if there's no entry to map a client's IP-to-MAC address then it will send a broadcast message to advertise the IP address of a packet that is trying to reach a client to every MAC address within the subnet. The client who owns the IP address then responds with its 48-bit MAC address and the packet is sent to that client. That address is then stored within the ARP table for future communications without having to re-broadcast to that client in the near future. The default time for ARP entries within the ARP table is 4 hours (cisco - times may vary between vendors), after that the table is flushed (deleted).

Q2). What is Packet Filtering? In our context, why is Packet Filtering termed 'IP Packet' Filtering? Why would you use Packet Filtering in your Network? What are the perceived benefits?

IPv4 works at layer-3 of the OSI model (routing) as discussed earlier in the report, it is called IP packet filtering for this reason. Layer-2 filtering is called MAC filtering because MAC addresses operate at Layer-2 of the OSI.

IP Packet filtering is a very important and useful tool for administrators as it allows granular control of traffic being sent or received from the WAN and/or between subnets i.e. HTTP, FTP etc to or from certain hosts and/or entire subnets within the network. This control provides extra security and efficiency within the LAN as demonstrated earlier in the report.

Q3). Why is it possible for contradictory filters to occur when you use packet filtering? Give an example of a Contradictory Filter:

A contradictory filter occurs when two filters with **apposing** logic are applied to the same solution that contradict each other; cancelling each other out. An example of this would be to create an inbound filter to allow traffic from HTTP and then create an outbound filter that blocks HTTP. This would result in a contradictory filter because two-way communication needs to be allowed for a client to communicate with a server/host.

Q4). FTP Server is just one Role/Service we can configure our Server to provide? What other roles can the Server provide that would benefit our Infrastructure in both a local and remote context?

Remote Desktop Management Services (RDMS) is a tool used within windows server 2012 and later, that allows an administrator to configure/troubleshoot a server remotely. This would be very beneficial to a network in the event an issue occurred during the night or when the administrator was out of the office, this service would allow the administrator to log into the server remotely to troubleshoot the issue.

This tool is also useful locally for an administrator to troubleshoot individual or multiple clients simultaneously within a network without having to visit each client one at a time, the clients could be 30 floors apart which could be expensive to the company in employee downtime.

IV. DYNAMIC HOST CONFIGURATION PROTOCOL

Dynamic Host Configuration Protocol (DHCP) is one of the most utilised and important protocols used within modern production network.

DHCP allows end devices (clients) within a network to autonomously receive an IP address from a DHCP server using broadcast messages at Layer-2.

Before initial configuration of a DHCP service can take place, careful planning of IP addressing schemes and network segmentation using subnets needs to be considered.

Internet Protocol Address Management (IPAM) (*not shown in this Lab*) is often used in conjunction with DHCP to manage address leases and IP address contradictions in order to manage IP addresses within a enterprise LAN which houses a vast amount of end devices.

For the purpose of this 'Lab' the IP addressing scheme is going to remain the same because there are only two subnets, but they will be managed and configured autonomously using the DHCP service.

After the IP addressing scheme has been planned, the next step is to install the DHCP role within the server manager dashboard shown in *figures 60-62*.




Figure 60 : Server Manager Dashboard




Figure 56: Add DHCP server from the roles and features wizard




Figure 57: Install the role on the server

Once the service has been installed, enter the administration tools screen and select DHCP shown in *figure 63*.




Figure 58: DHCP selection at the Administrative tools screen

Before A new DHCP scope can be created, the server needs to be a member of the Active Directory (AD) local group to prevent against rogue DHCP servers providing incorrect IP addresses within the network by malicious actors. To achieve this the server needs to be activated by selecting the server from the console tree; right clicking and then selecting activate from the menu shown in *figure 64-65*.




Figure 59: DHCP server Authorisation




Figure 60: Cardiff server IP address added for authorisation

Next scope creation and configurations of the DHCP service need to be applied; IP Address Range, Lease duration, Default gateway for subnet that the scope will be serving and Excluded addresses for internetworking devices i.e. servers, routers etc that need static IP addresses (Additional addresses will be excluded for future scalability i.e. additional routers, printers etc. shown in *figures 66-72*.




Figure 61: New scope creation by right clicking IPv4 in the console tree menu




Figure 62: New scope wizard, name added




Figure 67: Scope configuration confirmation




Figure 63: Subnet IP address range and subnet mask added



Figure 64: IP address range exclusions added for internetworking devices and future scalability




Figure 65: Default gateway added for subnet DHCP address allocation




Figure 66: 8-day IP address lease configured

Because there are two subnets within the network an additional scope will also be created and configured for the Lab (network 1) subnet by repeating these steps.

From the DHCP client Swansea, within its interface IPv4 properties; 'Obtain IP address automatically' will be selected instead of 'Use the following IP address' (static address that was configured earlier in the report) shown in figure 73.




Figure 68: Obtain IP address automatically (DHCP) selected at client's local interface

The next step is to issue the commands; `ipconfig /release` (to release the static IP address from interface), `ipconfig /renew` (to request a new address from DHCP server) within the 'Windows' command prompt shown in figure 74.




Figure 69: `ipconfig /release` and `ipconfig /renew` commands issued at 'Windows' cmd prompt

To confirm that the DHCP server has issued an IPv4 address to the client the command `ipconfig /all` is issued at 'Windows' cmd prompt to view the interfaces addresses shown in figure 75.




Figure 70: Correct IP address range and lease issued at Swansea DHCP client

All above steps are completed at the Lab client. DHCP is also verified at Lab client shown in figure 76.




Figure 71: Correct subnet IP address range and lease issued at Lab client

Windows DHCP service allows an administrator to reserve preferred addresses for clients. This is accomplished by first making a note of the Network Interface Cards (NIC) 48bit MAC address that can be viewed for 'Lab' in figure 76: *Physical Address* field. This is used to bind the IP address to the MAC address of the client when setting up the reservation.

To create the reservation, open the DHCP console via administrative tools; expand IPv4 and then scope. Right click on reservations and select 'New Reservation' and enter the required information within the fields at the next screen shown in figures 77-78




Figure 72: Create IP Address New Reservation




Figure 73: IP Reservation Configuration, Lab client address changed and reserved to 192.168.1.3 on next DHCP request

To obtain the new IP-address reservation at the Lab client the command `ipconfig /release` and `ipconfig /renew` will be issued within the 'Windows' command prompt at the Lab client shown in figure 79-80.




Figure 74: 'Release' command issued at Lab client




Figure 75: 'Renew' Command issued. The correct IP address that was reserved has been obtained and allocated through the DHCP service.

Q1). Define the role of the DHCP Server in your Network?

A DHCP server is used to allocate IP addresses autonomously within a network from a pre-configured pool of addresses. The IP address ranges and exclusions are first configured and then monitored by the networks systems administrator. This saves time and resources within a network negating the need for an administrator to configure every host manually (static IP addresses) within the network.

Q2). What happens when we have dynamically configured TCP/IP for our NIC, but then consequently cannot contact a DHCP server?

When a DHCP server can't be reached the client is issued with an Automatic Private IP Addressing (APIPA) in the 169.254.x.x range by the OS so that the client can communicate at layer-2 within its own subnet using ARP but not at layer-3.

Q3). Illustrate (with illustrations!!) the 4-stage process involved when DHCP clients look to obtain an IP address via DHCP?

The 4 stage DHCP process can be broken down into an acronym called DORA shown in figure Q3;




Figure Q3: 4 stage DHCP process

1. **Discover;** Client broadcasts a discover message to server.
2. **Offer;** DHCP server makes an offer via a unicast message to client.

3. **Request;** Client accepts offer and requests the IP address via broadcast message.
4. **Acknowledge;** Server acknowledges the request and issues IP address to client via unicast message.

Q4). Describe the advantages of using a relay agent. Explain (with descriptions and illustrations) how and why we would setup the Relay Agent on our Router?

A DHCP client and server communicate using broadcast messages which cannot be transported at layer-3 (though a router). If a DHCP client is in a sub-network that's not directly connected to the DHCP server a DHCP relay agent is used to relay the DHCP request through the router as a unicast message to retrieve an IP address from the DHCP server.

V. ACTIVE DIRECTORY AND DOMAIN NAME SYSTEM

Active Directory (AD) is a widely used service within networks around the world. AD converges network resources within a LAN by segmenting the network further into domains that provides accessibility and security control within a LAN.

AD's architecture consists of a hierarchical domain structure for files and directories that can be managed efficiently and securely by a network administrator for controlling permissions and authentication within a network.

AD is a Microsoft proprietary service that was first used in Windows 2000 and has since been updated through all versions of Microsoft server OS's which is the backbone of all windows production networks.

Domain Name System (DNS) is the most widely used service within networking and the Internet as a whole. It works by mapping memorable domain names to forgettable IP addresses e.g. www.google.com = 8.8.8.8 this is done so that an end devices Network Interface Card (NIC) knows where to send a request as it only understands IP addresses, this mapping is referred to as an 'A' record which is stored within the DNS database.

DNS also has a parent-child architectural structure, the most common industry configuration uses a primary DNS server and a secondary DNS server. The secondary DNS server is periodically updated from the primary DNS service, this model provides load balancing of domain name resolutions and redundancy.

AD utilises DNS to map domain names to IP addresses of server/clients within a LAN in order to locate and forward traffic between one or more domains

For the purpose of this lab an AD domain is going to be configured at Cardiff server in which 'Swansea' server and 'Lab' client will become members.

Also, a DNS primary and secondary zone will be configured for AD to utilise.

The first step is to install the AD Domain Services role within the roles and features dashboard shown in figures 81-83.




Figure 76: Roles and Features Dashboard




Figure 77: Active Directory Domain Services role selected




Figure 78: Role confirmation and installation

Once the AD role has been installed, the primary AD server has to be chosen and then from the flag icon within the dashboard the server needs to be promoted as a domain controller shown in figure 84.




Figure 79: Cardiff server promoted as a domain controller

Once the server has been promoted as a domain, a new forest can be created using the AD configuration wizard shown in figure 85. A forest within AD is considered the top of the tree within a hierarchical domain structure. Within a forest there can be many trees (sub domains) in which are controlled by the forest (domain controller).




Figure 80: New forest wizard: FQDN nwtraders.msft chosen

A Fully Qualified Domain Name (FQDN) then needs to be chosen. nwtraders.msft has been chosen for this lab. Windows server 2012 is chosen for minimum platform support, a restore password is also inputted. All other default configurations are used and then the forest is installed shown in figures 86-90.




Figure 84: Prerequisite check and new forest installation

After the forest is created and installed a restart is required, and the new forest domain can be logged into shown in figure 90.




Figure 81: Recovery Password inputted, and windows 2012 support selected




Figure 85: Server restart, welcome screen with new domain visible to log into




Figure 82: NETBIOS domain name default from forest: nwtraders

Once logged into the new domain, the DHCP service that was previously installed must be authorised and restarted in order for it to function within the domain shown in figures 91-93.




Figure 83: Default directories used for AD's stored files




Figure 86: DHCP authorisation wizard




Figure 87: New forest domain information inputted to be authorised within domain




Figure 88: DHCP service restarted at Cardiff server once authorised within domain

Now that AD has been correctly configured at Cardiff server, the second server Swansea and the client Lab need to join the domain. By adding/changing the workgroup in advanced system settings by clicking the computer name tab and then change. The domain option is then selected and new domain inputted. This is done within each of the OS's shown in figures 94-98




Figure 89: LAB client successfully joined new domain nwtraders.msft




Figure 90: Swansea server successfully joined new domain nwtraders.msft



Figure 91: Restart is needed after joining new domain on both OS's



Figure 92: After restart Lab is now part of domain. Cardiff credentials are used for administration domain login




Figure 93: After restart Swansea is now part of domain. Cardiff credentials are used for administration domain login

To verify that Swansea and Lab have joined domain successfully they can be viewed at Cardiff server's admin tools; AD Users and Computers shown in figure 99.




Figure 94: Verification that Swansea and Lab have successfully joined group

As discussed earlier in the report AD utilises DNS to locate a host within a domain using a FQDN for organisational and hierarchical considerations.

An FQDN is made up of a hierarchical namespace e.g. Lab.nwtraders.msft breaks down as; Lab being the host, nwtraders being the parent domain and msft being the top-level domain within this example.

To view the hostname of a Windows device the command; hostname is inputted into the Windows command prompt show in *figures 100-101*.




Figure 95: Hostname check at Cardiff; hostname confirmed as Cardiff




Figure 96: Hostname check at Swansea; hostname confirmed as Swansea

Next the FQDN needs to be verified after AD configuration at the systems screen > my computer > properties tab shown in *figures 102-103*.



Figure 97: FQDN for Swansea verified at the systems screen at the my computer properties name tab




Figure 98: FQDN for Cardiff verified at the systems screen at my computer properties name tab

This was also verified at the Lab client.

DNS now has to be installed from the management dashboard and configured at Swansea as this server is going to be the secondary DNS zone. The FQDN and DNS configuration is conducted within the DNS administration tools manager shown in *figure 104*.




Figure 99: DNS installed successfully at Swansea with correct FQDN

From the Cardiff server, Swansea's IP address and FQDN is added as a root hint which lists all other root DNS servers within the network. This is accomplished by selecting the Cardiff server from the DNS management console and then actions > root hint > add. The cache is also cleared which is referred to as DNS flushing which removes any old or conflicting DNS records, this is good practice when adding a

root DNS server to avoid a DNS conflict shown in *figures 105-106*.




Figure 100: Swansea IP and FQDN added as root hint at Cardiff DNS service




Figure 101: DNS cache cleared at Cardiff DNS server

The next step is to configure IP forwarding on the secondary DNS server (Swansea) to query records from the primary DNS server (Cardiff) to handle external lookup queries outside of the network or domain with a query timeout set to 10 seconds.

This is accomplished by highlighting Swansea server from the side panel then going to actions then properties shown in *figure 107*.




Figure 102: IP forwarder created to Cardiff primary DNS server. FQDN successfully resolved

Verification of the DNS IP forwarder is confirmed by pinging the hostname within the domain instead of the IP address of Cardiff to confirm that hostname to IP mapping is taking place shown in *figure 108*.




Figure 103: IP forwarder working correctly verified by pinging hostname

A forward lookup zone will now be created using the new zone wizard at Swansea DNS server by right clicking the Swansea FQDN within the DNS manager side panel and selecting new zone.

This lookup zone will be the secondary DNS zone and will be synchronised with the primary DNS server (Cardiff) to resolve host-to-IP addresses shown in figures 109




Figure 104: New zone selected from FQDN properties




Figure 105: Secondary zone selected




Figure 106: Forward lookup zone selected




Figure 107: Domain name added (nwtraders.msft) as secondary forwarding zone name




Figure 108: IP address of Master DNS added. FQDN resolved successfully



Figure 109: Configuration review and secondary forward zone created

So that the secondary DNS server can be updated from the Master DNS server, zone transfers must be permitted from both servers for security reasons to prevent a rogue DNS server being created within the domain by malicious actors and also to prevent malicious actors getting a full list of IP addresses of the servers within the network.

For periodic updates to take place after the Master transfer, dynamic updates need to be configured which will also be configured later in the report.

Zone transfers are used for load balancing in the case of the Master DNS server having a lot of requests, the jobs can be shared with the secondary DNS server for network efficiency. Also mirroring of the DNS records between the Master and Secondary DNS server creates redundancy in the event that the Master DNS server fails and/or the records got lost, the network could still function and the records could be retrieved once the Master server is operational.

This is configured by right clicking the DNS zone created and from the Zone transfers tab; allow zone transfers from 'any server' shown in figure 115.

'Any server' zone transfer was selected for this lab as it's a basic configuration, best practise would be to only allow known servers within the domain by filtering zone transfers by IP address of trusted DNS servers within the network.




Figure 110: Zone transfers permitted between master and secondary zones

Now that zone transfers between DNS servers has been permitted, the next step is to right click on the forward lookup zone at the secondary DNS server and transfer from 'Master', this will update the secondary DNS forward lookup zone from the Master/primary DNS servers' database shown in figure 116.




Figure 111: Master DNS record transfer to secondary forwarding zone

Now that a secondary forward lookup zone has been created, a reverse lookup zone is also required to resolve IP-address-to-hostnames queries.

This zone is configured similarly to the forward zone that was configured using the zone wizard, but instead of selecting forward zone, reverse lookup zone is selected instead.

Only the network ID is inputted so that all hosts within the subnet are encompassed within the zone. The Master DNS servers IP address (Cardiff) is also configured shown in figures 117-121.




Figure 112: Reverse lookup zone selected at zone configuration wizard




Figure 113: Subnet ID added at zone configuration wizard




Figure 114: Master DNS (Cardiff) IP and FQDN added at zone configuration wizard



Figure 115: Reverse lookup zone configuration confirmation and creation




Figure 116: Reverse lookup verification at the DNS manager screen at Swansea server. Confirmation zone has been loaded from Master DNS server

A primary zone will now be created at the Cardiff server called `swanseaserver.nwtraders.msft` alongside the default Master zone `Cardiff.nwtraders.msft` for added redundancy of the secondary DNS zone using periodic updates of zone transfer as discussed earlier in the report for redundancy and load balancing purposes shown in figures 122-126.

This configuration of the 'Primary' zone is used as a Lab example to demonstrate how DNS domains can be decoupled. Industry best practice would be to periodically update the secondary zone from the master zone. For added security in the unlikely event that both DNS servers were lost a 3rd DNS zone would be created remotely which only stores a backup copy of the DNS database's 'A' records and would not participate in the primary DNS operations.




Figure 117: Primary lookup zone created at Cardiff called `swanseaserver.nwtraders.msft`



Figure 118: Secure dynamic update selected (AD) for periodic update of secondary DNS zone (Swansea)




Figure 119: Confirmation and verification of new primary zone




Figure 120: Secure zone transfer configured for only Swansea server. IP and FQDN added




Figure 121: New zone verified at Swansea server. Zone transfer successful

DNS is the primary service utilised on the internet to translate website hostnames to web server IP addresses. To demonstrate this, a web server and index webpage will be configured at Swansea server using the IIS service installed earlier at the FTP section of the report.

Once it has been configured the web server will be accessible from the Lab clients web browser using Swansea.traders.msft FQDN shown in figures 127-132.




Figure 122: IIS service selected at administrative tools




Figure 123: Add Virtual directory selected at the default site icon




Figure 124: Physical path and alias (name) configured for virtual directory to wwwroot folder.




Figure 125: Permissions edited at c:\inetpub\wwwroot directory to allow applications (http) to read and execute scripts




Figure 126: Properties selected at the wwwroot directory to edit file permissions




Figure 127: Permissions edited at wwwroot directory to allow internet guests to modify files (http browsers)

Now that the directories and IIS web service has been configured the HTML file within the folder is changed to index.html as this is the first page (root file) the web browser looks for at a web server. The PNG file inside the folder has also been modified with the paint application to demonstrate how the HTML file utilises the image to present the page. 'Hello World' has been painted on the image to be displayed shown in figures 133.




Figure 128: Files inside wwwroot located. HTML file name changed to index.html and the PNG edited with the paint application to display 'Hello World'

Website access verified by accessing the Swansea server webserver through the Lab clients default browser using Swansea FQDN shown in figure 134.




Figure 129: Swansea webserver test site verified by accessing it via Lab clients default browser using Swansea's FQDN

Q1). What steps does a client take (what is the process involved) when trying to resolve a Name to an IP address?

The DNS process to find an internet websites IP address steps are as follows;

1. The operating System of the host sends a recursive query message to a local DNS server called the DNS resolver.
2. The DNS server then sends an iterative query to the root DNS server.
3. The root server checks its 'A' record database that maps domain names to IP addresses and then returns the IP address of the top-level domain DNS servers.
4. This process is repeated down the domain levels until the correct host IP address is found and returned to the client.

Within a LAN all domain levels are housed within the master DNS server. So, the host OS would send a recursive query to its DNS server which would search its database for the correct 'A' record. The DNS local server simply returns the correct IP address to the host OS.

Q2). Describe how the Recursive Query and Iterative Query Processes resolve DNS Client Requests?

A recursive query is the initial query sent by the client OS to its DNS server to resolve a domain name.

If the DNS server is unable to resolve the query from its own database of 'A' records it issues a query of its own to another DNS server, this request is called an iterative query.

The query is then returned with the correct IP address to the local DNS server who in turn sends the correct IP to the client who sent the recursive query.

Q3). Explain the significance of root hints and how this process enables us to resolve names that are not managed (or kept in an authoritative DNS server in our network)?

Root hints within a DNS server is a database/file that contains IP-addresses of DNS servers outside the domain of the local DNS server so that it can send iterative queries for traffic requests that it can't resolve from its own 'A' records.

Q4). We have a Server called Server1, this server will be hosted in a subdomain called 'subcontA' with a root domain called nwtraders.msft, answer the following questions:

(1). What is the exact FQDN for this host?

Server1.subcontA.nwtraders.msft

(2). Illustrate how the Domain Name System breaks up this FQDN into Host and Domain Namespace.

. = root domain (the dot is not visible within a browser/domain only to the machine)

msft = top level domain

nwtraders = parent domain (second level domain)

subcontA = subdomain

Server1 = host

DNS generalized breakdown depicted in figure 130.




Figure 130: FQDN breakdown [25]

Q5). Explain how DNS Zones are used in DNS, give examples to illustrate the use of multiple zones and how these can be used to manage resource records of domains.

What important factors/characteristics do we need to be aware of?

DNS zones are used for hierarchical querying of domains from the root domain that iteratively queries and finds top-level domains down to the second level and sub domains that receive recursive queries directly from the host and subsequently return the correct IP address.

DNS peer zones are also used for load balancing queries where there are many requests to a particular website; for example www.google.com has millions of requests per second so more than one (Many) DNS zones would be required to handle that amount of requests. These zones are usually spread out around the world to reduce latency and bandwidth congestion across the WAN for hosts within different countries.

In an enterprise LAN where bandwidth is limited secondary zones are also used to reduce bandwidth on the network.

Peered zones that are dynamically updated from other zones can also be used as a redundant DNS source as discussed earlier in the report.

VI. SERVER CLUSTERING

A server cluster is defined as a group of two or more physical servers which are referred to as nodes sharing resources to provide high availability network services that are viewed as one logical entity within a network's infrastructure.

Server clusters are primarily utilised to provide high availability (redundancy) of network services to end devices over a LAN or WAN topology thus they are referred to as failover clusters; they are also utilised for high performance Quality of Service (QoS) and load balancing. They most commonly share a bank of storage devices or are clones of each other's HDD's which are periodically updated for synchronisation. [26]

I. Quorum Logic

A server clusters architectural foundation is built upon the mathematical principles of a 'Quorum' which can be defined within computing as; a minimum number of members present within an autonomous system to reach a dynamic democratic decision. This principle is simple, but its mechanics depend on the redundancy of server clusters around the world. [27]

Each server cluster will contain 'quorum data' that encompasses the clusters configuration data, because clusters are primarily built for redundancy within high availability networks as discussed earlier in the report, each quorum data present within a cluster has its own dynamic 'vote', in the event that a node stops responding, the other nodes will send a ping for a set amount of time to determine if the node is offline.

If there is a majority of nodes still communicating which is determined by the quorum data that is still accessible by the remaining nodes, then a consensus is reached and traffic to that node is stopped and is distributed between the rest of the cluster. If there is not a majority of quorum data left online all nodes will shut down as each node can't be sure if the other is available because there is not a majority present, so a logical democratic decision cannot be made to shut down an individual node.

If all nodes weren't taken offline when there was not a majority present it could result in more than one node processing the same job and could result in data corruption which would have larger implications to the network than if the cluster was offline, for this reason quorum logic is used within failover-clusters. [28] [29]

II. Cluster Configurations

There are many types of server cluster architectures within modern day networks depending on the geographical topology i.e. Campus Area Network (CAN), Enterprise

network, Small Medium Business (SMB) network, Wide Area Network (WAN). [30]

Three cluster configurations are most prevalently used within modern day network cluster architectures and are considered industry standards, they are defined by how the quorum data configuration is distributed within a cluster. These are; [31]

1. **Single Quorum Device Cluster;** A single quorum device cluster is a cluster of two or more servers connected within a bus topology configuration each governed by a single quorum. This configuration is considered the most robust in terms of redundancy and is the most widely used. Within this model the cluster can still function with only one node as there is only one quorum governing the autonomous system, therefore a majority is not needed shown in *figure 136*.

The disadvantage of this cluster configuration is that because there's only one copy of quorum data, a single point of failure is created which defeats the point of a failover cluster. If the quorum data got corrupted or went offline the entire cluster would fail, that's why redundant storage is also used within this configuration which is referred to as cluster storage. Cluster storage will be discussed later in the report.




Figure 131: Single Quorum Device Cluster [31]

2. **Majority node set cluster;** A majority node set cluster is a cluster of three or more servers which have to be an odd number to maintain a majority vote connected within a 'bus' network topology similarly to a single quorum device cluster, the difference is that each node has its own quorum data which is stored locally at the nodes Hard Disk Drive (HDD) or at a local Storage Area Network (SAN). The clusters quorum data are all kept consistent across the nodes by the quorum resources synchronisation mechanism. Within this configuration the cluster can only operate with a majority of nodes still functioning which provides less redundancy than a single quorum device cluster. In a scenario where the majority of the nodes go offline the entire cluster would shutdown even the healthy nodes.

The advantages of this cluster is that its scalable across a WAN because each node has its own quorum data. Also quorum data for each node can be temporarily taken offline for maintenance without effecting the cluster as a whole as long as the majority of quorum data remains online within the cluster. Shown in *figure 137*.




Figure 132: Majority node set cluster topology [31]

3. **Local quorum cluster;** A local quorum cluster consists of only one node and one storage disk that contains the quorum data. Theoretically this is not a cluster because there is only one server within the configuration. It is referred to as a local quorum cluster because it is used in a lab environment to test new cluster technologies, cluster architectures and configurations by engineers before exposing them to a live production network shown in figure 138.




Figure 133: Local quorum cluster topology [31]

III. Storage Redundancy

As discussed earlier in the report storage redundancy is also important within a single quorum device cluster to prevent the quorum data becoming lost which would result in the cluster failing.

The most common redundant cluster storage [32] used within a single quorum device cluster is a; Redundant Array of Independent Disks (RAID) 10 configuration also known as RAID 1+0 [33]. This configuration consists of a minimum of four HDD's or Solid-State Drives (SSD) configured for mirroring and striping of data at the same time across an array of disks shown in figure 139.




Figure 134: RAID 10 configuration [34]

Data Mirroring (RAID 1); Data mirroring creates two or more copies of the same data at the same time. It is economically expensive as it only provides 50% of storage in comparison to the cost of two disks, the advantage is that the data is in more than one place providing a redundant backup copy. This is a small price to pay to secure the company's important information. [35]

Data Striping (RAID 0); Data Striping is the process of sharing blocks of data between drives for higher throughput of data, utilising the read speeds of the disks, because the data is spread over two or more disks the drives speeds are combined for higher read speeds for the data being accessed by the client device. [36]

IV. Cluster Load Balancing

Server clusters are utilised for many different business solutions around the world not just for their primary function of fail-over redundancy. Within industry, production servers are very expensive; having server hardware sitting idle just for backup purposes does not make economic sense for companies. [37]

To utilise failover clusters system administrators configure them for load balancing of jobs to evenly distribute requests from clients; the most common load balancing method used is a 'round robin' format shown in figure 140. [38]




Figure 135: Load balancing; Round robin format diagram [39]

Load balancing/failover clusters are most commonly used within data centres where vast amount of data is being accessed by millions of end users around the world at a given time, load balancing within a cluster has two advantages; it makes computational expensive data accessible to vast amounts of end devices providing higher throughput i.e. video streaming and also provides a safety net if a server went down or needed to be taken offline for maintenance shown in figure 141.




Figure 136: Datacentre load balancing example [40]

Windows server provides a failover cluster service role [41] that can be used to cluster a group of servers, it also includes load balancing which will be utilised in the next report; *Cluster services in a virtual environment.*

VII. CONCLUSION

Network services have evolved dramatically within the last 20 years to keep up with the ever-growing demand for data

consumption around the world. Network services have become more complex as more end users demand access to massive computationally expensive data such as video, voice, social media etc within LAN and WAN's every day.

What this report has briefly outlined is how virtualisation and server clustering have been intelligently developed and adopted to provide affordable, efficient, and high availability services to keep up with the challenges the information age we live in today presents.

Within a future report virtualised cluster services will be demonstrated within a Lab environment similar to the one used in this report to illustrate how the cluster technologies discussed earlier are utilised practically within industry today.

VIII. REFERENCES

- [1] cisco, "Understanding LANs & Cisco ICND1 Foundation Learning Guide: LANs and Ethernet," *ciscopress.com*, 2013. [Online]. Available: <http://www.ciscopress.com/articles/article.asp?p=2092245&seqNum=2>. [Accessed: 14-Feb-2018].
- [2] Cisco, "Introduction to WAN Technologies - DocWiki," *cisco.com*, 2017. [Online]. Available: http://docwiki.cisco.com/wiki/Introduction_to_WAN_Technologies. [Accessed: 14-Feb-2018].
- [3] Microsoft, "Windows Server 2012 Administration Overview | Microsoft Press Store," *micorsoft press store*, 2012. [Online]. Available: <https://www.microsoftpressstore.com/articles/article.aspx?p=2201311>. [Accessed: 14-Feb-2018].
- [4] VMware, "Hypervisor Virtualization Overview," *vmware.com*, vol. 1, no. 1, p. 11, 2017.
- [5] Microsoft, "Overview of System Center 2012 - Virtual Machine Manager," *microsoft.com*, 2016. [Online]. Available: [https://technet.microsoft.com/en-us/library/gg671827\(v=sc.12\).aspx](https://technet.microsoft.com/en-us/library/gg671827(v=sc.12).aspx). [Accessed: 14-Feb-2018].
- [6] Cisco, "Packet Tracer - Network Simulator."
- [7] Microsoft, "Understanding Generation 1 and Generation 2 Virtual Machines," *microsoft.com*, 2016. [Online]. Available: [https://technet.microsoft.com/en-us/library/dn440675\(v=sc.12\).aspx](https://technet.microsoft.com/en-us/library/dn440675(v=sc.12).aspx). [Accessed: 14-Feb-2018].
- [8] vmware, "ESXi | Bare Metal Hypervisor | VMware," *vmware.com*, 2017. [Online]. Available: <https://www.vmware.com/products/esxi-and-esx.html>. [Accessed: 14-Feb-2018].
- [9] virtualization practice, "2016-2017 Hypervisor Cost Comparison," *virtualizationpractice.com*, 2017. [Online]. Available: <https://www.virtualizationpractice.com/2016-2017-hypervisor-cost-comparison-39502/>. [Accessed: 14-Feb-2018].
- [10] Microsoft, "Hyper-V Technology Overview | Microsoft Docs," *microsoft.com*, 2016. [Online]. Available: <https://docs.microsoft.com/en-us/windows-server/virtualization/hyper-v/hyper-v-technology-overview>. [Accessed: 14-Feb-2018].
- [11] Citrix, "Citrix XenServer."
- [12] SANS, "The OSI Model: An Overview The OSI Model: An Overview Rachelle L. Miller GSEC Practical Assignment Version 1.2e."
- [13] Cisco, "Core Routing and Routers," *cisco.com*, p. 12, 2017.
- [14] Cisco, "Multiprotocol Label Switching (MPLS) on Cisco Routers - Cisco," *cisco.com*, 2002. [Online]. Available: https://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/fs_rtr22.html. [Accessed: 14-Feb-2018].
- [15] J. Ungerman and C. #6167, "ISP Edge design," *cisco.com*, vol. 1, no. 1, p. 61, 2015.
- [16] microsoft, "Hyper-V Network Virtualization - Routing and Remote Access services," *microsoft.com*, 2017. [Online]. Available: <https://docs.microsoft.com/en-us/windows-server/networking/sdn/technologies/hyper-v-network-virtualization/hyperv-network-virtualization-technical-details-windows-server>. [Accessed: 14-Feb-2018].
- [17] ciscopress, "Dynamic Routing Protocols (3.1) & Cisco Networking Academy's Introduction to Routing Dynamically," *ciscopress.com*, 2014. [Online]. Available: <http://www.ciscopress.com/articles/article.asp?p=2180210&seqNum=4>. [Accessed: 14-Feb-2018].
- [18] Omnisecu, "Difference between Interior Gateway Protocol (IGP) and Exterior Gateway Protocol (EGP)," *omnisecu.com*, 2015. [Online]. Available: <http://www.omnisecu.com/cisco-certified-network-associate-ccna/difference-between-igp-and-egp.php>. [Accessed: 14-Feb-2018].
- [19] juniper, "RIP Overview - Technical Documentation - Support - Juniper Networks," *juniper.net*, 2017. [Online]. Available: https://www.juniper.net/documentation/en_US/junos/topics/concept/routing-protocol-rip-security-overview.html. [Accessed: 14-Feb-2018].
- [20] Cisco, "Enhanced Interior Gateway Routing Protocol - Cisco," *cisco.com*, 2017. [Online]. Available: <https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/16406-eigrp-toc.html>. [Accessed: 14-Feb-2018].
- [21] cisco, "Open Shortest Path First (OSPF) - Cisco," *cisco.com*, 2017. [Online]. Available: <https://www.cisco.com/c/en/us/products/ios-nx-os-software/open-shortest-path-first-ospf/index.html>. [Accessed: 14-Feb-2018].
- [22] cisco, "Static Route Information," *cisco.com*, 2015. [Online]. Available: https://www.cisco.com/c/en/us/td/docs/switches/data-center/sw/5_x/nx-os/unicast/configuration/guide/13_cli_nxos/13_route.html. [Accessed: 14-Feb-2018].
- [23] cisco, "Cisco BGP overview," *cisco.com*, 2015. [Online]. Available: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/15-mt/irg-15-mt-book/irg-overview.html. [Accessed: 14-Feb-2018].
- [24] juniper, "ARP Overview - Technical Documentation - Support - Juniper Networks," *juniper.net*, 2012.


- [Online]. Available: https://www.juniper.net/documentation/en_US/junos/e13.3/topics/concept/ip-arp-overview.html. [Accessed: 14-Feb-2018].
- [25] ccm, "DNS (Domain Name System)," *cccm.net*, 2018. [Online]. Available: <http://ccm.net/contents/271-dns-domain-name-system>. [Accessed: 14-Feb-2018].
- [26] Microsoft, "Failover Clustering | Microsoft Docs," *Microsoft.com*, 2016. [Online]. Available: <https://docs.microsoft.com/en-us/windows-server/failover-clustering/failover-clustering-overview>. [Accessed: 14-Feb-2018].
- [27] Oxford English Dictionary, "quorum | Definition of quorum in English by Oxford Dictionaries," *oxforddictionaries.com*, 2018. [Online]. Available: <https://en.oxforddictionaries.com/definition/quorum>. [Accessed: 14-Feb-2018].
- [28] Microsoft.com, "Understanding Quorum in a Failover Cluster | Clustering and High-Availability," *Microsoft.com*, 2011. [Online]. Available: <https://blogs.msdn.microsoft.com/clustering/2011/05/27/understanding-quorum-in-a-failover-cluster/>. [Accessed: 14-Feb-2018].
- [29] S. Perriman, "Introduction to the Cluster Quorum Model | Clustering and High-Availability," *microsoft.com*, 2010. [Online]. Available: <https://blogs.msdn.microsoft.com/clustering/2010/05/13/introduction-to-the-cluster-quorum-model/>. [Accessed: 14-Feb-2018].
- [30] Cisco, "Network Architectures for the Enterprise & Network Architectures for the Enterprise Overview," *ciscopress.com*, 2011. [Online]. Available: <http://www.ciscopress.com/articles/article.asp?p=1718490>. [Accessed: 14-Feb-2018].
- [31] Microsoft, "What Is a Server Cluster?: Server Clusters (MSCS)," *microsoft.com*, 2003. [Online]. Available: [https://technet.microsoft.com/en-us/library/cc785197\(v=ws.10\).aspx](https://technet.microsoft.com/en-us/library/cc785197(v=ws.10).aspx). [Accessed: 14-Feb-2018].
- [32] microsoft, "How to Configure a Clustered Storage Space in Windows Server 2012 | Clustering and High-Availability," *microsoft.com*, 2012. [Online]. Available: <https://blogs.msdn.microsoft.com/clustering/2012/06/02/how-to-configure-a-clustered-storage-space-in-windows-server-2012/>. [Accessed: 14-Feb-2018].
- [33] IBM, "IBM Knowledge Center - RAID 10 overview," *ibm.com*, 2015. [Online]. Available: https://www.ibm.com/support/knowledgecenter/ST5GLJ_8.3.0/com.ibm.storage.ssc.help.doc/f2d RAID-10implementation_1tkkta.html. [Accessed: 14-Feb-2018].
- [34] cyberciti, "RAID 5 vs RAID 10: Recommended RAID For Safety and Performance - nixCraft," *cyberciti.biz*, 2018. [Online]. Available: <https://www.cyberciti.biz/tips/raid5-vs-raid-10-safety-performance.html>. [Accessed: 14-Feb-2018].
- [35] oracle, "Overview of RAID 1 (Mirror) Volumes (Solaris Volume Manager Administration Guide)," *oracle.com*, 2010. [Online]. Available: <https://docs.oracle.com/cd/E19683-01/817-5776/about-metadevices-9/index.html>. [Accessed: 14-Feb-2018].
- [36] oracle, "Chapter 7 RAID 0 (Stripe and Concatenation) Volumes (Overview) (Solaris Volume Manager Administration Guide)," *oracle.com*, 2010. [Online]. Available: <https://docs.oracle.com/cd/E19683-01/817-2530/6mi6gg86c/index.html>. [Accessed: 14-Feb-2018].
- [37] Microsoft, "Load-Balanced Cluster," *microsoft.com*, 2000. [Online]. Available: <https://msdn.microsoft.com/en-us/library/ff648960.aspx>. [Accessed: 14-Feb-2018].
- [38] Microsoft, "Network Load Balancing | Microsoft Docs," *microsoft.com*, 2017. [Online]. Available: <https://docs.microsoft.com/en-us/windows-server/networking/technologies/network-load-balancing>. [Accessed: 14-Feb-2018].
- [39] Adobe, "High availability clustering ColdFusion: Using ColdFusion Enterprise Manager | Adobe Developer Connection," *adobe.com*, 2008. [Online]. Available: https://www.adobe.com/devnet/coldfusion/articles/clustering_cf8.html. [Accessed: 14-Feb-2018].
- [40] Microsoft, "Clustering Technologies: Server Clusters (MSCS); Network Load Balancing (NLB)," *microsoft.com*, 2003. [Online]. Available: [https://technet.microsoft.com/en-us/library/cc784941\(v=ws.10\).aspx](https://technet.microsoft.com/en-us/library/cc784941(v=ws.10).aspx). [Accessed: 14-Feb-2018].
- [41] Mono Inc, "Windows Server 2012 Hyper-V Failover Clustering - Part 3: Installation | Mono Software," *mono*, 2016. [Online]. Available: <https://mono.software/2016/01/28/Hyper-V-Failover-Cluster-3/>. [Accessed: 14-Feb-2018].

IX. BIBLIOGRAPHY

- Dr Bashford, Tim Moodle:
<https://moodle.uwtsd.ac.uk/course/view.php?id=3239>
- Stanek, William 'Windows Server 2012: Inside Out'
ISBN-10: 0735666318
- Odom, Wendell CCNA 'Routing and Switching'
200-105 official cert guide cisco press (2016)
ISBN-10: 1-58720-580-7
- Robinson, Tony Building Virtual Machine Labs: A Hands-On Guide
ISBN-10: 1546932631

X. APPENDIX-A

Full size Hyper-V cluster environment topology on next page.



- 100GB of partitioned space containing Lab virtual images; Swansea, Cardiff and Lab client
- Storage Area Network containing virtual images for Hyper-V
- Two Hyper-v physical servers running 'Windows 2012 RS'
- Tesla Virtual Gateway (IP Address/domain) Configured at Router depicted as server
- Local Area Network LAN
Gateway Load Balancing Protocol (GLBP) used at router to create virtual gateway to load balance and create redundancy

Virtualized Network Services

GRADEMARK REPORT

FINAL GRADE

GENERAL COMMENTS

Instructor

86/100

PAGE 1



Comment 1

Take a look at how/where semicolons should be used: <https://writing.wisc.edu/Handbook/Semicolons.html>



Comment 2

e.g.



Comment 3

Hyphenate



Comment 4

Applied - hypervisors existed for quite a few decades before this point, and this almost implies they were developed specifically to solve this problem

PAGE 2



Comment 5

Cabling



Comment 6

No apostrophe here - because if you were typing it in full, you'd say "virtual machines" not "virtual machine's"

PAGE 3

PAGE 4

PAGE 5

PAGE 6

PAGE 7

PAGE 8

PAGE 9

PAGE 10

PAGE 11

PAGE 12



Comment 7

opposing

PAGE 13

PAGE 14

PAGE 15

PAGE 16

PAGE 17

PAGE 18

PAGE 19

PAGE 20

PAGE 21

PAGE 22

PAGE 23

PAGE 24

PAGE 25

PAGE 26

PAGE 27

PAGE 28

PAGE 29

GRADING FORM: SMU GENERIC_GJ

RESEARCH-INFORMED LITERATURE



An excellent number of sources used, with referencing standards correctly and consistently applied. There is a heavy weighting towards web sources, many of them of questionable credibility. You should aim to predominantly rely on peer-reviewed work if you can (i.e. content from academic journals). Authoritative web sources can be acceptable, though should be considered in context - for example, Microsoft's TechNet would be a good source for discussing the technical aspects of Hyper-V... but should be taken with a pinch of salt if being used to describe its benefits, or compare it against competitors. For the next assignment, try to aim for a higher portion of peer-reviewed material - journal articles or similar authoritative sources (as discussed above) are always preferable.

7

KNOWLEDGE AND UNDERSTANDING



Some excellent knowledge and understanding displayed. You've obviously done a huge amount of additional reading in the completion of this report - very well done, Paul.

18.5

ANALYSIS

Very good synthesis of sources to form cogent



arguments. As a suggestion, you might want to try to include a few more interim conclusions explicitly throughout the report (conclusion subsections) - this'd help tie your report together, and smooth the transition between sections.

16

PRACTICAL APPLICATION AND DEP.



An excellent practical application. I especially enjoyed the paint "Hello World". A few marks lost simply because I couldn't read the text in a few of the screenshots - not that this meaningfully detracts from an excellent piece of work.

37

SKILLS FOR PROFESSIONAL PRAC.



Generally good spelling, punctuation and grammar. Be careful with your use of apostrophes and semicolons - there're a few systematic mistakes throughout the report. I'm happy to go through their use with you at some point, if you'd like. I love the IEEE formatted document, though it has resulted in some of the screenshots being a tad small, which isn't a major issue, though some of the text is hard to read, which matters in a few places - you could consider spanning them across the whole page, where the text really matters (e.g. ipconfig dumps).

7.5