NOS Worksheets 1-5

by George Searle

FILE	2092_GEORGE_SEARLE_NOS_WORKSHEETS_1-5_100686_1936660944.DOCX (15.34M)		
TIME SUBMITTED	11-FEB-2018 06:12PM (UTC+0000)	WORD COUNT	15685
SUBMISSION ID	83057495	CHARACTER COUNT	105187

Network Services Worksheets 1-5

George Searle – Computer Networks & Cyber security (P148812)
Network Services (M5X01050)
UWTSD School of Applied Computing
Dr. Tim Bashford
Submitted: 11/02/18

Contents

Section 1: Facilitating a Virtualised Remote Network Infrastructure utilising Type 1 Hypervisor software, and the RRAS (Routing and Remote Access Service) 3

1.1. Verifying Network Connections 3
1.2. Enabling the Routing and Remote Access Service on a Server 9
1.3. Add Routing Protocols on the Server and Add Two Routing Interfaces to the Routing Protocol on the Server 11
1.4. Learning about Local Area Network Routing 13
1.5. Type 1 Hypervisor and Routing Observational Tasks 14

1.5.1. The benefits of Type 1 Hypervisor software services and available products 14
1.5.2. The role of Routers in Network Architecture and types available 15
1.5.3. Routing Protocol types, usage and benefits 15
1.5.4. The purpose of Routing Tables and the information they provide 16
1.5.5. The purpose and significates of a subnet mask in the Routing Table 16

Section 2: Routing and Remote Access Services in Windows Server: ARP, Managing TCP/IP & ICMP Operations using Packet Filtering 17

2.1. Using the ARP Protocol to obtain Physical/Logical Network Node information 17
2.2. Permitted TCP and ICMP Operations 18
2.3. Restricting ICMP Operations Using IP Packet Filtering with the Windows 2012 R2 Server Routing and Remote Access Service 24
2.4. Restricting TCP Operations Using IP Packet Filtering with the Windows 2012 R2 Server Routing and Remote Access Service 26
2.5. Allowing only Cardiff to FTP to Swansea Server while blocking all other traffic from Cardiff and the Lab network 29
2.6. ARP, Packet Filtering and Server Roles Observation Tasks 30

2.6.1. ARP Process 30
2.6.2. Packet Filtering 31
2.6.3. Contradictory filters 31
2.6.4. Server Roles and Services 31

Section 3: TCP/IP Configuration using DHCP (Dynamic Host Control Protocol) 32

3.1. Understanding Different TCP/IP Configurations 32
3.2. Installing the DHCP Server 34
3.3. Authorising the DHCP Server 34
3.4. DHCP Scopes 35
3.5. Configuring DHCP reservations 40
3.6. DHCP Observational Tasks 42

3.6.1. Role of the DHCP Server 42
3.6.2. Troubleshooting DHCP 42
3.6.3. The DHCP Lease Process 43
3.6.4. DHCP Relay Agent 43

Section 4: Name Resolution using the Domain Name System (DNS) 45

4.1. Understanding DNS operation; Planning and deploying DNS 45
4.2. Storing and Maintaining DNS Data 52
4.3. DNS Application 68
4.4. DNS Observational Tasks 73

4.4.1. Resolving a Name to IP address 73
4.4.2. Recursive Query and Iterative Query processes 73
4.4.3. Root Hints 74
4.4.4. Question 74
4.4.5. DNS Zones 75
Section 5: Evaluation of clustering strategies for hosting a 3-tier web development system 78
5.1. Introduction 78
5.2. Web Architecture 78
5.2.1. Physical Tiers and Distributed Deployment of Applications 78
5.2.2. Logical Web Application Tiers 79
5.3. Application Clustering 80
5.4. Failover Clustering 81
5.4.1. Cluster Topologies 81
5.4.2. Failover Cluster and Network Load Balancing 82
5.5. Tools and Technologies 83
5.6. Discussion 84
5.7. Conclusion 85
References 85

Section 1: Facilitating a Virtualised Remote Network Infrastructure utilising Type 1 Hypervisor software, and the RRAS (Routing and Remote Access Service)

1.1. Verifying Network Connections

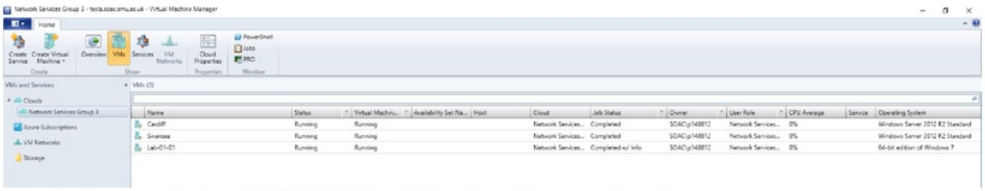


Fig. 1. System Center 2012 R2 Virtual Machine Manager interface

System Center 2012 R2 Virtual Machine Manager is a Virtual Machine (VM) management software based on Microsoft Virtual Server and Hyper-V [1].
Fig. 1. shows the three VMs Cardiff, Swansea and Lab-01-01 that are part of the Network Services Group 3 private cloud on the tesla.soac.smu.ac.uk cluster. This is the environment that will be used for the tasks presented in this report.

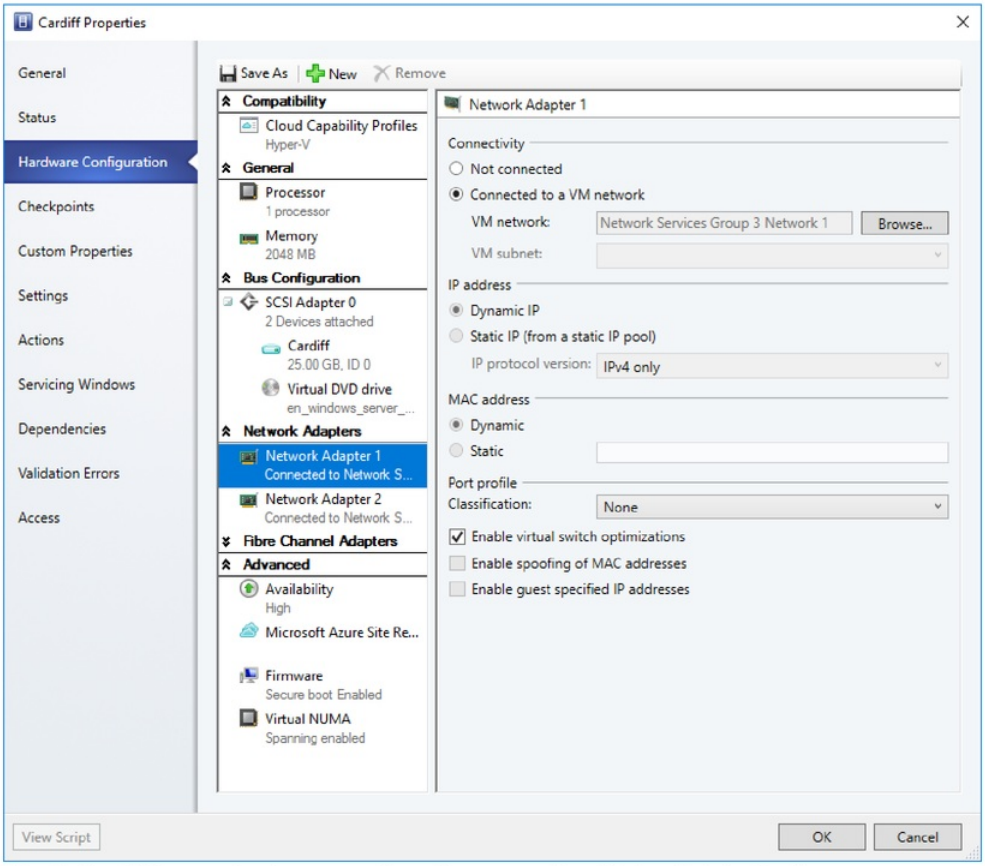


Fig. 2. Cardiff server hardware configuration properties (Network Adapter 1)

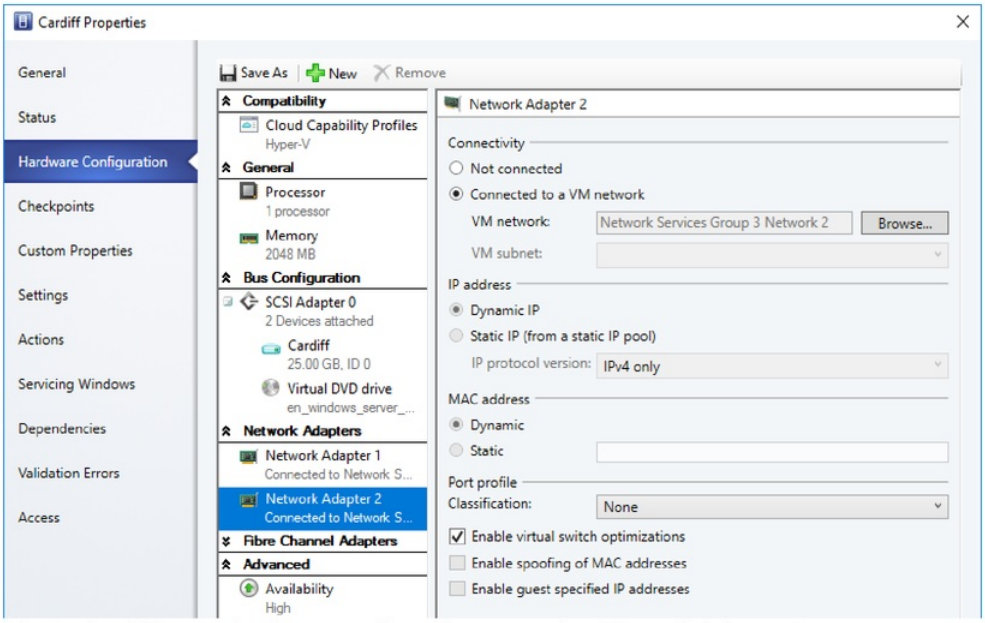


Fig. 3. Cardiff server hardware configuration properties (Network Adapter 2)

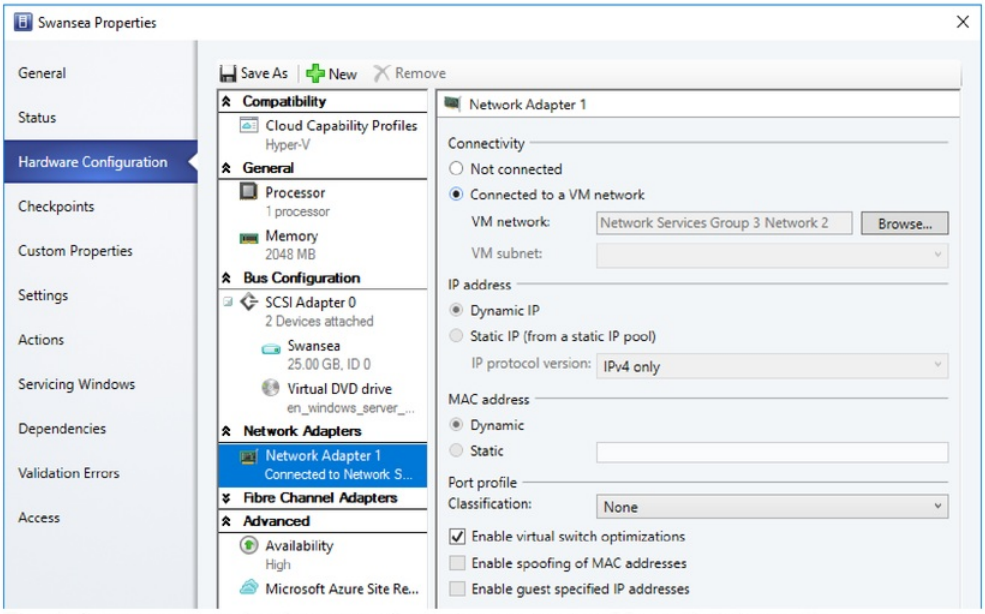


Fig. 4. Swansea server hardware configuration properties (Network Adapter 1)

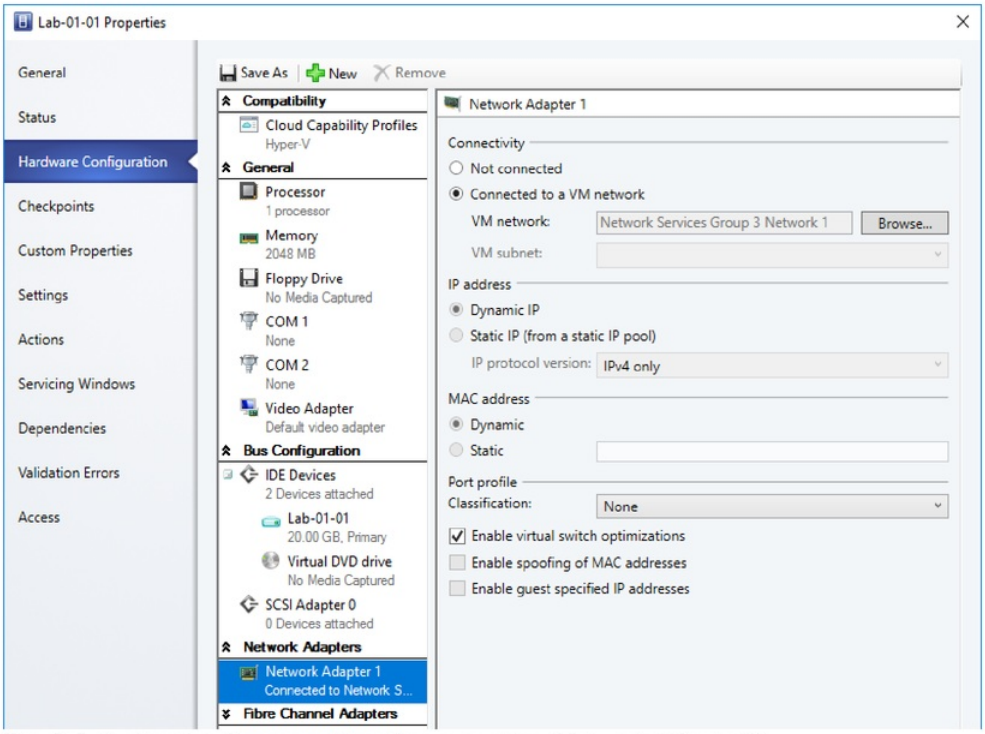


Fig. 5. Lab client hardware configuration properties (Network Adapter 1)

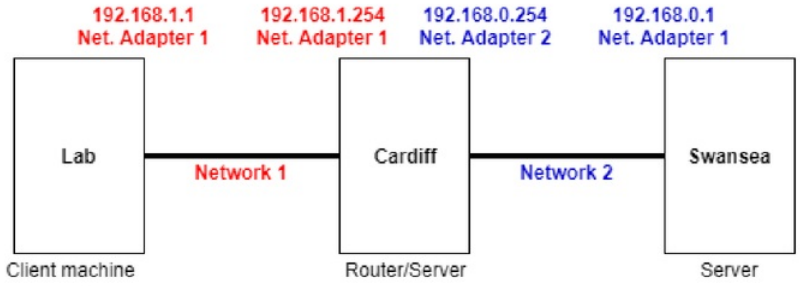


Fig. 6. The virtual network connectivity between servers and client

Fig. 2-3. show the Cardiff server's hardware configuration and Network Adapter properties. Cardiff's Network Adapter 1 (Fig. 2) is connected to the VM network of Network Services Group 3 Network 1 and will act as a virtual cable connecting Cardiff's Network Adapter 1 to the Lab's Network Adapter 1 (Fig. 5). Cardiff's Network Adapter 2 (Fig. 3) is connected to the VM network of Network Services Group 3 Network 2. It will act as a virtual cable connecting Cardiff's Network Adapter 2 to Swansea server's Network Adapter 1 (Fig. 4). The network connectivity of the Lab, Cardiff and Swansea can be visualised in Fig. 6.

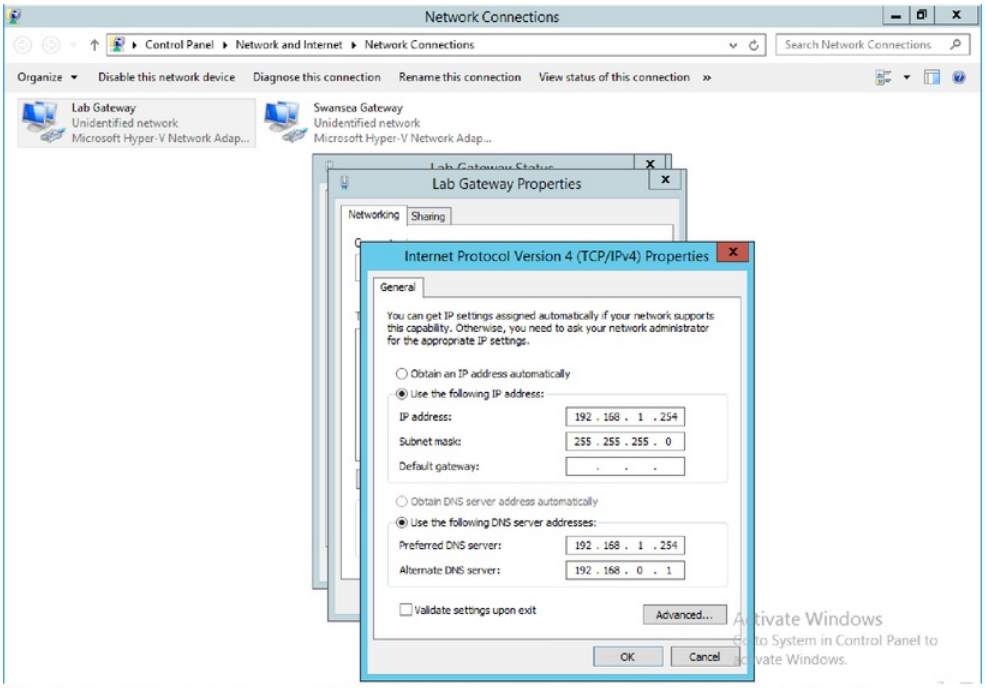


Fig. 7. Cardiff's static Internet Protocol (IP) address for Network Adapter 1 (Fig. 2).

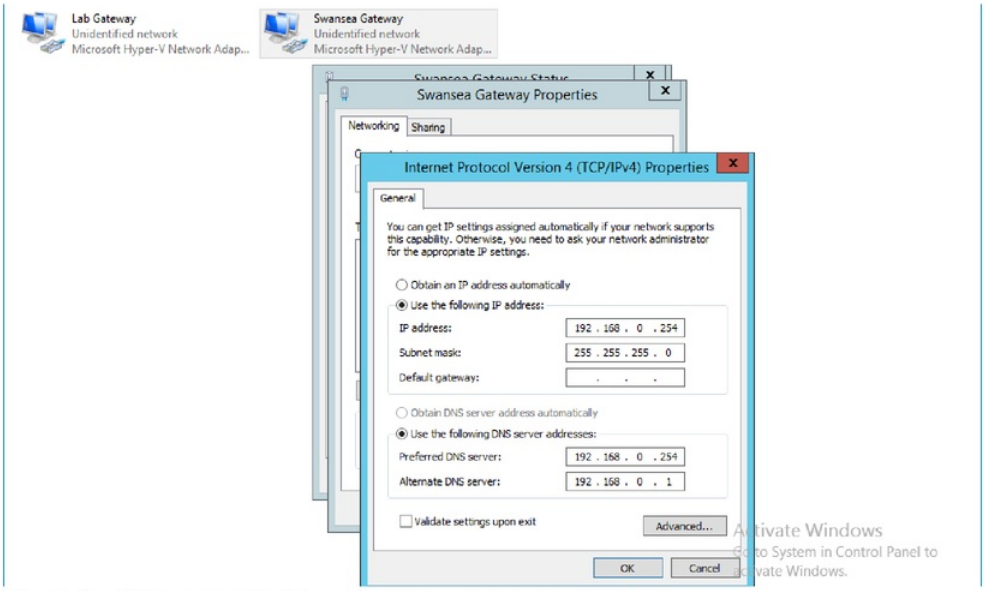


Fig. 8. Cardiff's static IP address for Network Adapter 2 (Fig. 3).

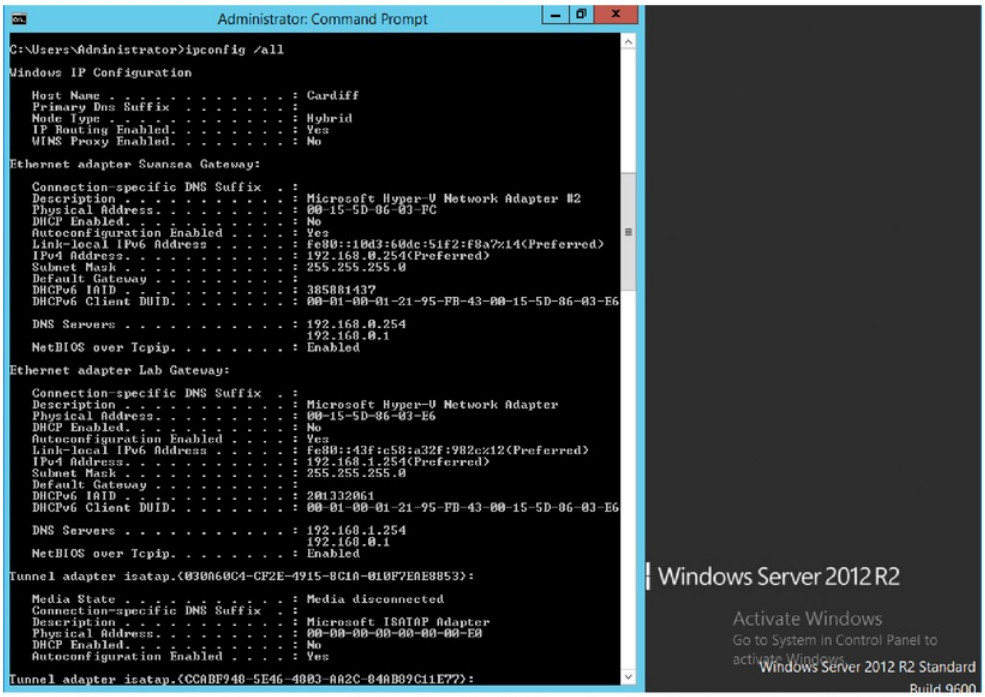


Fig. 9. The ipconfig /all command showing Cardiff's network adapter settings for Network Adapter 1 (Lab Gateway) and 2 (Swansea Gateway).

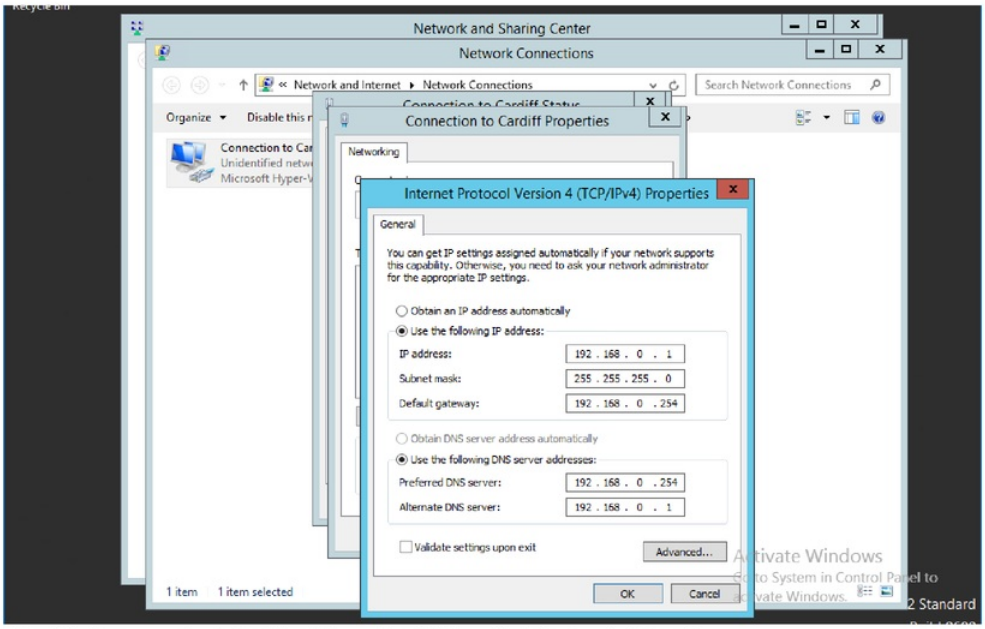


Fig. 10. Swansea's static IP address for Network Adapter 1 (Fig. 4).

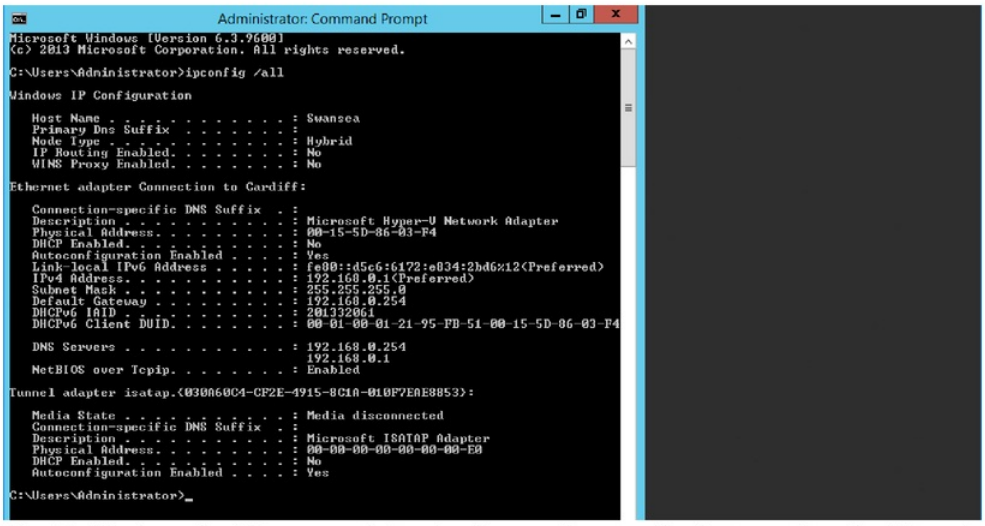


Fig. 11. The ipconfig /all command showing Swansea's network adapter settings for Network Adapter 1 (Connection to Cardiff).

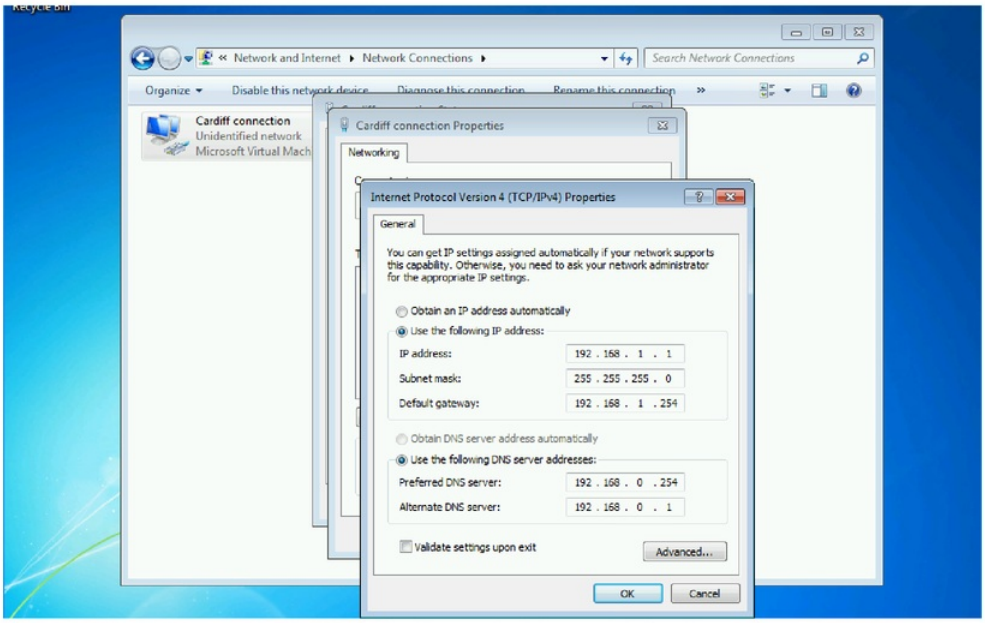


Fig. 12. The Lab's static IP address for Network Adapter 1 (Fig. 5).

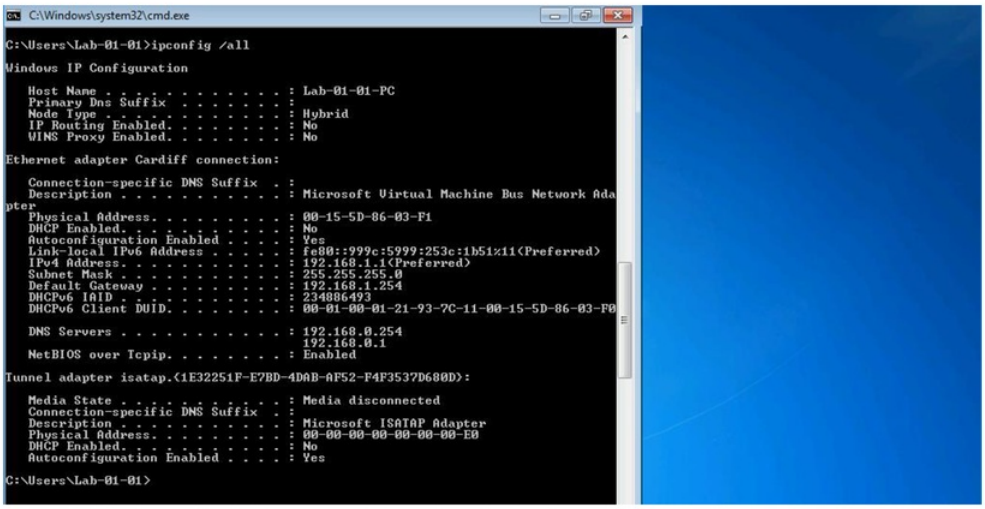


Fig. 13. The ipconfig /all command showing the Lab's network adapter settings for Network Adapter 1 (Cardiff connection).

For the Lab, Cardiff and Swansea to be able to communicate with each other they require Internet Protocol (IP) addresses. IPv4 addresses have an address structure length of 32-bits expressed as four numbers separated by dots [2]. To facilitate communication between the Lab, Cardiff and Swansea each adapter must be assigned a unique IP address (in this case IPv4). The IPv4 address acts as a unique identifier for a device on the Internet and local intranet. In this case, local intranet as the Lab, Cardiff and Swansea are interconnected local area networks (LANs) [3], [4]. The IPv4 addressing scheme can be seen in Fig. 6.

1.2. Enabling the Routing and Remote Access Service on a Server

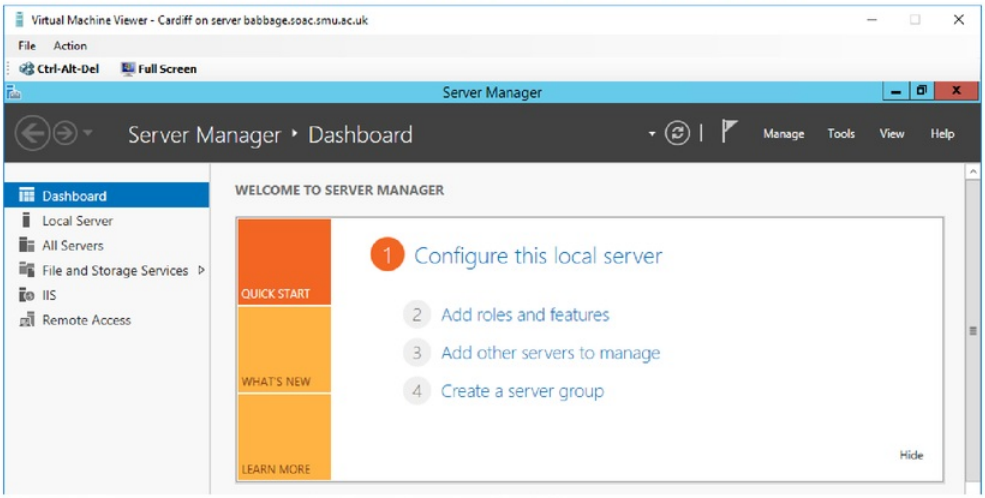


Fig. 14. The server manager dashboard on the Cardiff server

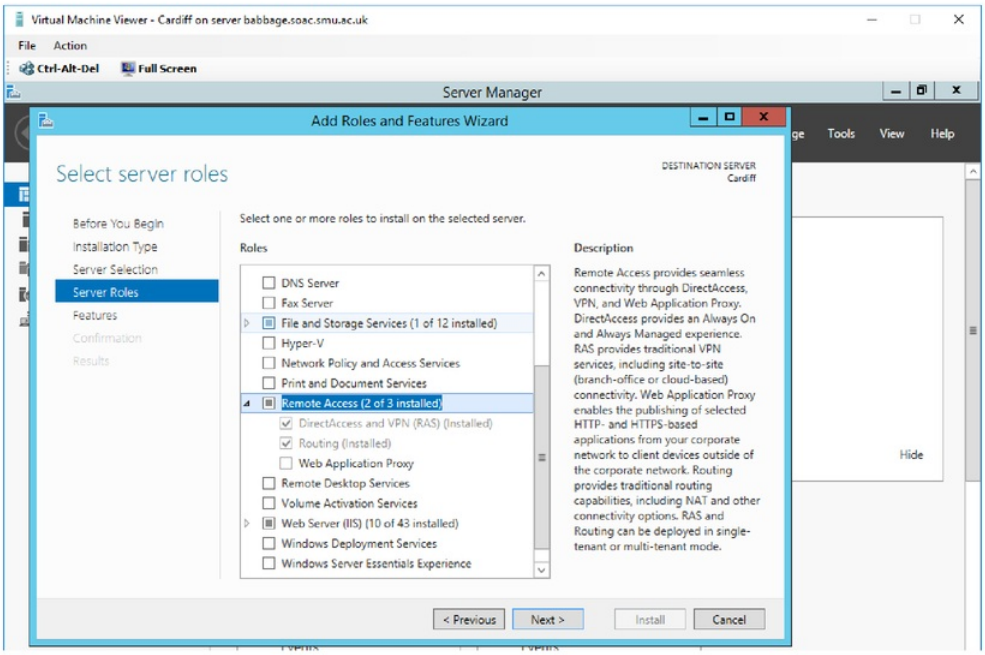


Fig. 15. The Routing server role installed on Cardiff

Fig. 14. shows the server manager dashboard that is used to add new roles and features on Windows Server 2012 R2 running on the Cardiff server. Fig. 15. shows that the Routing role (part of Remote Access) has been installed. The Routing role is needed to allow the Cardiff server to act as a router to route network traffic between the 192.168.0.0/24 network and 192.168.1.0/24 network (Fig. 6).

1.3. Add Routing Protocols on the Server and Add Two Routing Interfaces to the Routing Protocol on the Server

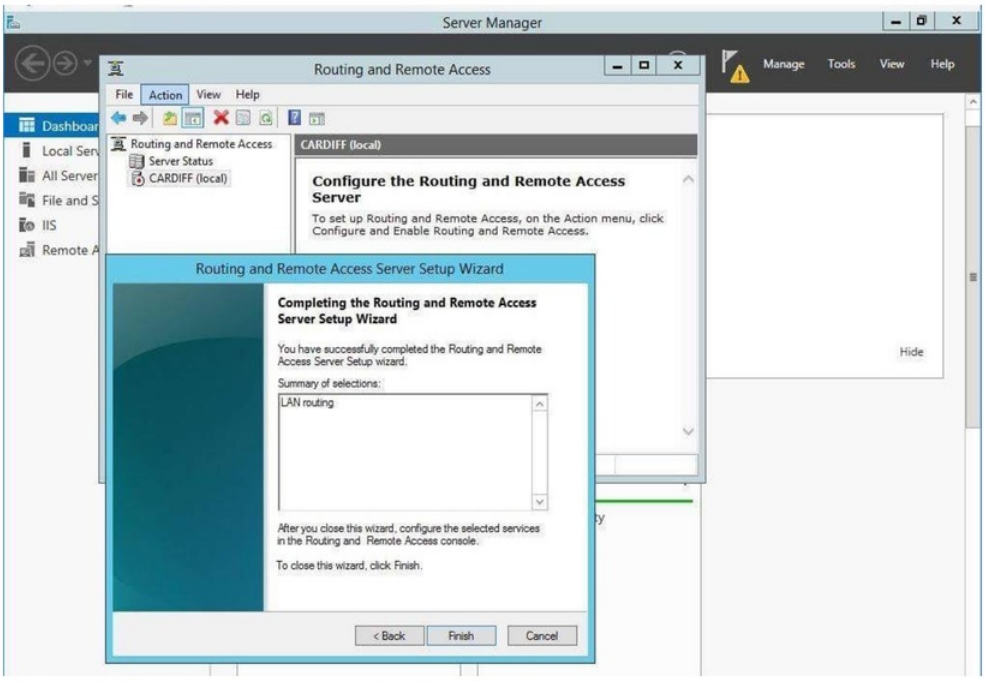


Fig. 16. Setting up LAN routing on the Cardiff server

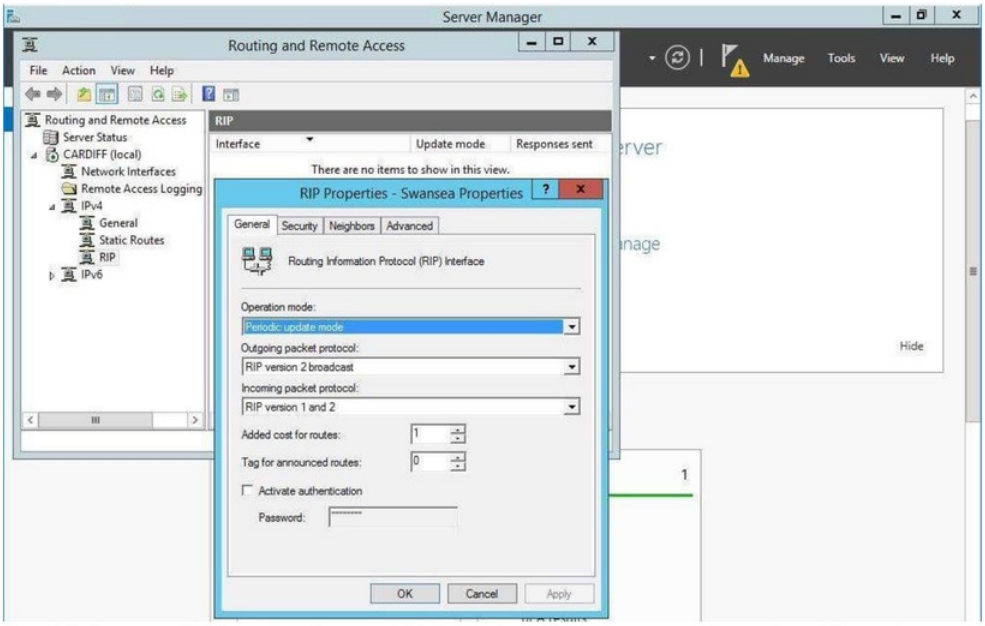


Fig. 17. Routing Information Protocol Version 2 (RIPv2) being configured on Network Adapter 2 (connection to Swansea server)

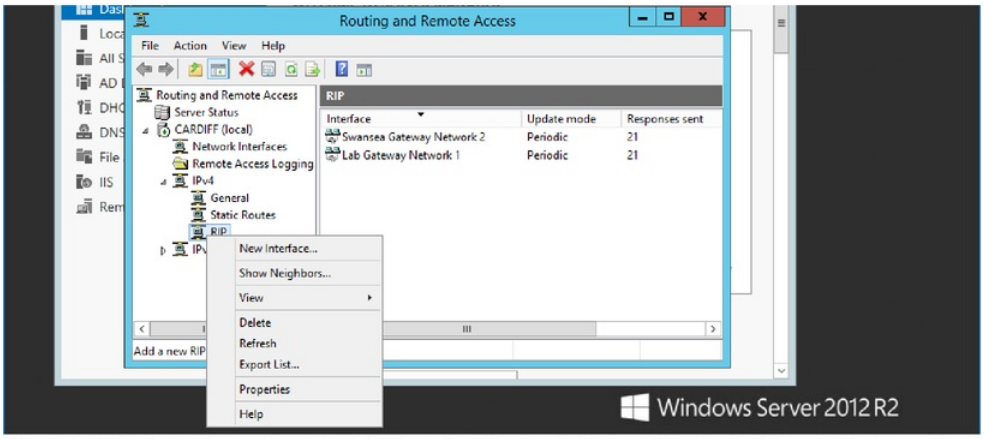


Fig. 18. RIPv2 configured on both Network Adapter 1 (connection to Lab) and 2 (connection to Swansea)

Fig. 16-18. shows the process of configuring RIPv2 on both Network Adapters on the Cardiff server (right click 'RIP' in the panel to add a new interface). RIPv2 is a distance-vector routing protocol that uses router hop count as its metric, with a maximum hop count of 15 [5], [6]. RIPv2 sends out routing table updates every 30 seconds and supports variable-length subnet mask (VLSM) which is an added improvement over RIPv1 [6], [7]. RIPv2 is a widely supported and easy to configure routing protocol making it suited to small networks [6], [8]. However, the protocol does not scale well for large networks due to a slow convergence time [6], [8]. It should be noted that a routing protocol is required for there to be communication between the Lab and Swansea as they are on different networks (Fig. 6). RIPv2 is not the only routing protocol that can be used, there are many others including static routes, Enhanced Interior Gateway Routing Protocol (EIGRP) and Open Shortest Path First (OSPF).

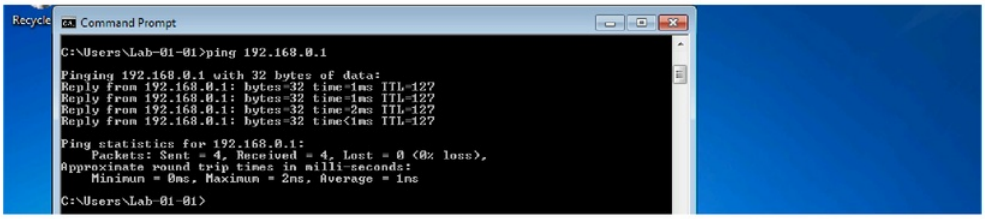


Fig. 19. Ping successful from Lab client to Swansea server

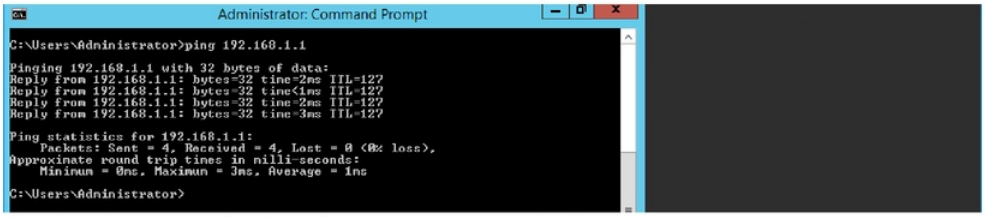


Fig. 20. Ping successful from Swansea server to Lab client

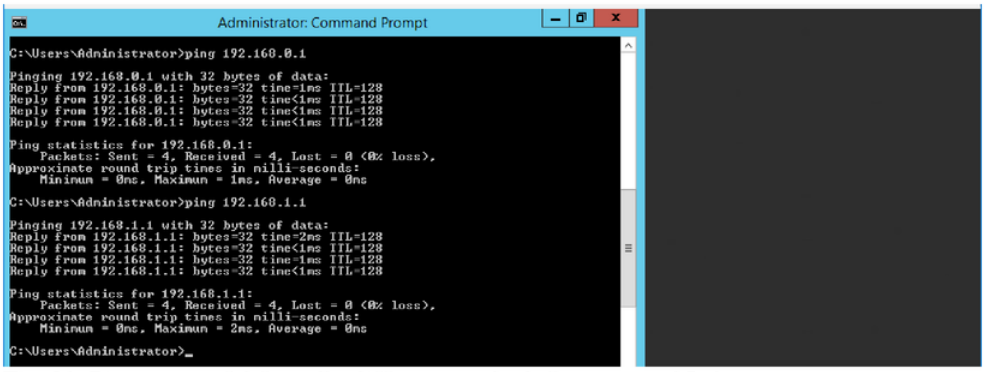


Fig. 21. Ping successful from Cardiff server to both Swansea and Lab

Fig. 19-21. shows the use of the Ping network utility and demonstrates how there is successful communication between the Lab, Cardiff and Swansea (seen by 0% loss).

1.4. Learning about Local Area Network Routing

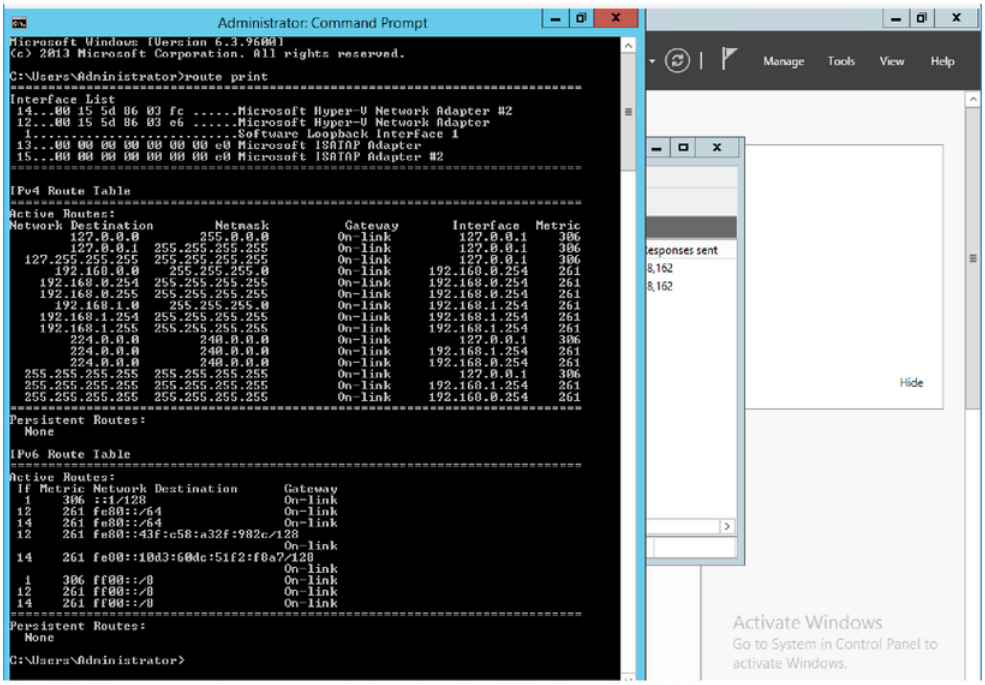


Fig. 22. The route print command showing Cardiff's Routing Table

Fig. 22. shows information about IP networks and how they are reachable (directly or indirectly) [9]. For example in the routing table, the network destination 192.168.0.0 can be seen as an active route. This shows that the 192.168.0.0/24 network is reachable (which includes the 192.168.0.1 address of Swansea) by the Cardiff server via the Network Adapter interface assigned with the IPv4 address of 192.168.0.254 (directly connected). The 192.168.1.0 network is also an active route and the same theory applies allowing the Lab client to be reached by Cardiff. The address 127.0.0.0 is a loopback network address, 127.0.0.1 is a localhost loopback

address, 224.0.0.0 is a multicast address and 255.255.255.255 a limited broadcast address [9], [10]. The interface addresses of 192.168.0.254 and 192.168.1.254 with the subnet mask 255.255.255.255 are host routes that allow for more efficient processing of packets that are directed to Cardiff (router) [11].

1.5. Type 1 Hypervisors and Routing Observational Tasks

1.5.1. The benefits of Type 1 Hypervisor software services and available products

A Hypervisor is software that manages multiple virtual machines and allows multiple Operating Systems (OSs) to share the same hardware. It controls, allocates and virtualises hardware resources for every virtual system [12].
There are Type 1 Hypervisors that run directly on a system's hardware and Type 2 that run on a host's OS (Fig. 23). The host OS can then provide virtualisation services e.g. I/O device support and memory management [13].

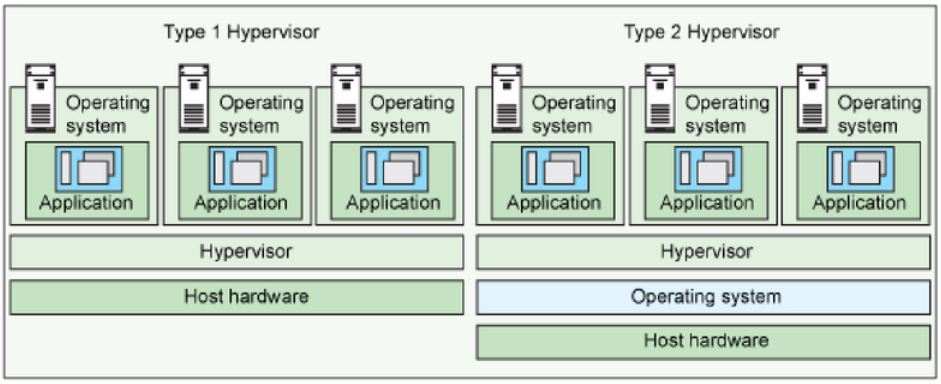


Fig. 23. The difference between Type 1 and Type 2 Hypervisors [13]

The benefits of Type 1 Hypervisors include better performance and flexibility as they run directly on the hardware and have less overhead [14], [15]. Another advantage is that if a single VM crashes it will not affect any of the other VMs running on the hypervisor [16].

There are a number of Type 1 Hypervisors available detailed below [12], [17], [18]:

- VMware ESX and ESXi – Offers advanced features and saleability but requires higher cost licensing. VMware is the market leader in Type 1 hypervisors. Their vSphere/ESXi product is available in a free edition and 5 commercial editions.
- Microsoft Hyper-V – Lacks many advanced features offered by VMware ESX and ESXi e.g. fault tolerance and power distribution management. Available free (no GUI & virtualisation rights) and in 4 commercial editions. Microsoft is aggressively improving Hyper-V to catch up with vSphere and releasing their own utilities.
- Citrix XenServer – Linux based, lacks a lot of features that ESX/ESXi has. Management and administration are more difficult than Hyper-V and ESX/ESXi. Citrix is lagging behind in developing and improving their product compared to Microsoft and VMware.

- Oracle VM – Based on the open source Xen project. Lacks many advanced features in other Type 1 Hypervisors. Free, however support and updates will cost.

1.5.2. The role of Routers in Network Architecture and types available

A router is a layer 3 networking device that is responsible for forwarding packets from network to network, from the source to final destination. A router connects multiple networks together. This means that the multiple interfaces on a router will belong to a different IP network. The two primary functions of a router are to determine the best path to send a packet and then forward packets towards their destination. A router uses a routing table to determine the best path to a destination network [19].

There are many different types of routers available [20]–[22]:

- Consumer routers – Used in the home or by small businesses. They normally act as a wireless access point, switch, Dynamic Host Configuration Protocol (DHCP) server and modem.
- Enterprise routers – Used by medium to large businesses. They allow for advanced device configurations and support many routing protocols. Interfaces are also customisable using interface cards (for Cisco routers) [23].
- Wireless routers – Instead of using wires to transmit packets, wireless routers use radio waves (transmitted via one or more antennae). Wireless routers would form a wireless local area network (WLAN) as opposed to LAN.
- Core routers – Forwards packets within a network, but no between networks.
- Edge routers – Is a router located at the edge of a network and used to connect to external networks e.g. wide area networks (WANs) and the Internet.
- Bridge routers – Is a device that acts as both a router and bridge.
- Virtual router – Is software that replicates the functionality of hardware-based layer 3 routers. Virtual routing is a form of Network Function Virtualisation (NFV) [24]. Hyper-V uses network virtualisation and can perform virtual routing as seen in section 1.3..

1.5.3. Routing Protocol types, usage and benefits

Routing protocols are used to allow the exchange of routing information between routers. The purpose of dynamic routing protocols is to discover remote networks, maintain up-to-date routing information, choose the best path to destination networks and find a new best path if the current path becomes unavailable. Dynamic routing protocols are made up of data structures, routing protocol messages and an algorithm for calculation of the best path [19].

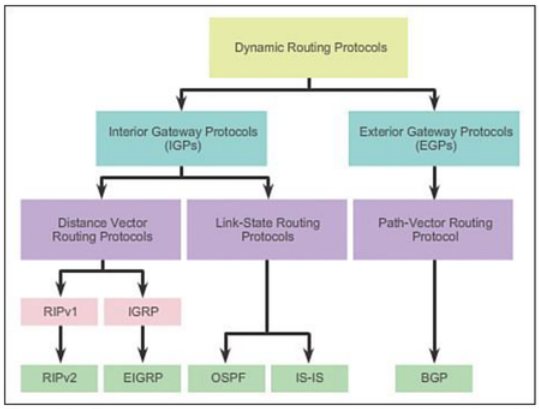


Fig. 23. Dynamic routing protocol classification [11]

There are many classifications of different dynamic routing protocols as seen in Fig. 23, with RIPv2, EIGRP, OSPF, Intermediate System to Intermediate System (IS-IS) and Border Gateway Protocol (BGP) used commonly today [11].

There are a number of advantages and disadvantages associated with dynamic routing protocols [19]:

Advantages:	**Disadvantages:**
Appropriate in all topologies where multiple routers are required.	More complex to implement.
Normally independent of network size.	Less secure by default – addition configurations are required to secure.
Automatically changes topology to reroute traffic if possible.	Route depends on the current topology.
	Can be resource heavy on CPU, RAM and link bandwidth.

1.5.4. The purpose of Routing Tables and the information they provide

The purpose of the routing table is to store route information about directly connected and remote networks. The routing table contains IPv4 and IPv6 network addresses and next hop address associations [19]. In Windows routing table the subnet mask is given and the Gateway column is the next hop address with 'on-link' meaning directly connected [25]. The network address and next hop address (Gateway) tells the router how to optimally reach a particular destination [19]. The metric column shows a route's preference if there are multiple routes to a destination. For example, the shortest route will have the lowest metric and be the preferred route [9].

1.5.5. The purpose and significates of a subnet mask in the Routing Table

A subnet mask is a 32-bit number that normally accompanies an IP address. Its purpose is to divide an IP address into the network address and host address. For example, 192.168.1.0 with the subnet mask 255.255.255.0 means that 192.168.1 is the network portion of the address and .0 is the host portion of the address. The host part of the address can then be assigned to devices (.1-.254) except for .0 which is reserved as the network address and .255 reserved as the broadcast address. Subnetting can be used to

further divide the host part of the address into different subnets. For example, 192.168.1.0 with subnet mask 255.255.255.252 allows 2 hosts to be assigned 192.168.1.1 and 192.168.1.2. This means that 192.168.1.4 255.255.255.252 would now be a different subnet and not part of the 192.168.1.0/30 network.

The subnet mask is significant in the routing table as the mask determines the network address of its connected interface [26]. The subnet mask will also highlight how many hosts will be part of that subnet. It is worth noting that the subnet mask 255.255.255.255 that accompanies the 192.168.0.254 and 192.168.1.254 interface addresses make them host routes [9].

Section 2: Routing and Remote Access Services in Windows Server: ARP, Managing TCP/IP & ICMP Operations using Packet Filtering

2.1. Using the ARP Protocol to obtain Physical/Logical Network Node information

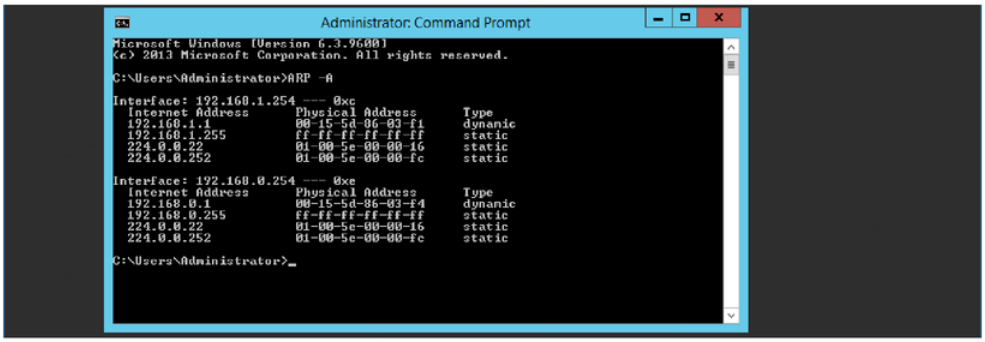


Fig. 24. Cardiff's Address Resolution Protocol (ARP) table

Fig. 24. shows Cardiff's Address Resolution Protocol (ARP) table which maps logical IP addresses to physical Media Access Control (MAC) addresses of network interfaces. In the table, the Lab's IP address 192.168.1.1 can be seen mapped to the MAC address 00-15-5d-86-03-f1. Swansea's IP address 192.168.0.1 is mapped to the MAC address 00-15-5d-86-03-f4. The IP address starting with 224.0.0 are multicast addresses [27] and the MAC address ff-ff-ff-ff-ff-ff is a broadcast [28].

2.2. Permitted TCP and ICMP Operations

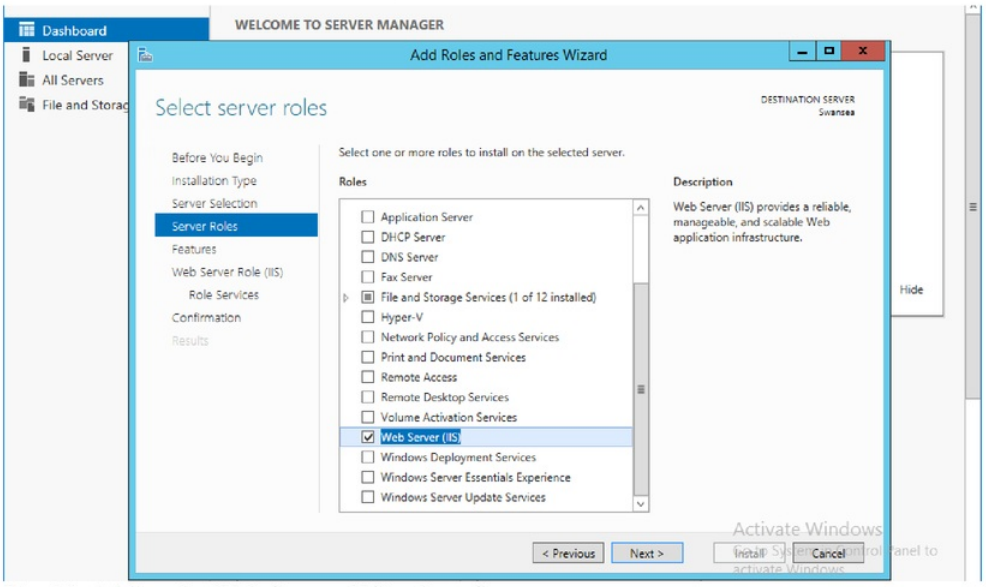


Fig. 25. Adding the Web Server (IIS) role to Swansea server

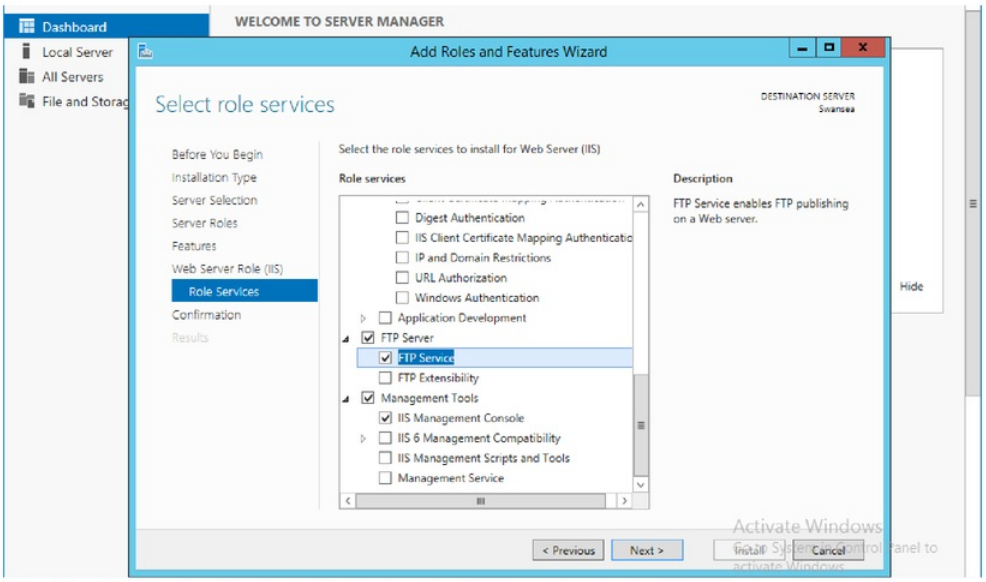


Fig. 26. Adding the File Transfer Protocol (FTP) service to the Web Server (IIS) Role

Fig. 26-27. show the Web Server (IIS) role with File Transfer Protocol (FTP) service being added to the Swansea server. This is because the Swansea server will act as an FTP server that Cardiff and the Lab will connect to.

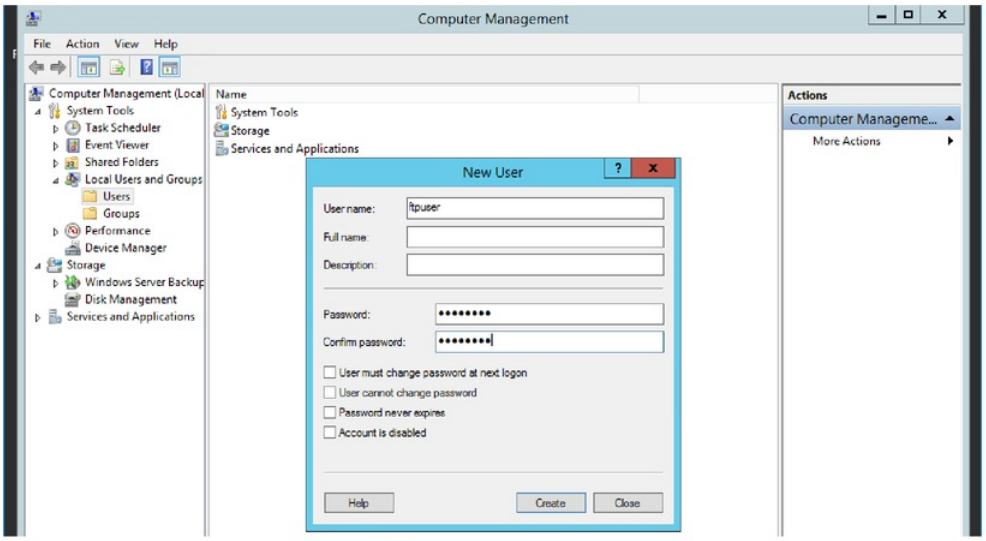


Fig. 27. Adding a new user to Local Users

Fig. 27. shows the user 'ftpuser' with password 'P@ssw0rd' being added to the Users folder of Local Users and Groups. This will be used as the username and password to access to Swansea FTP server.

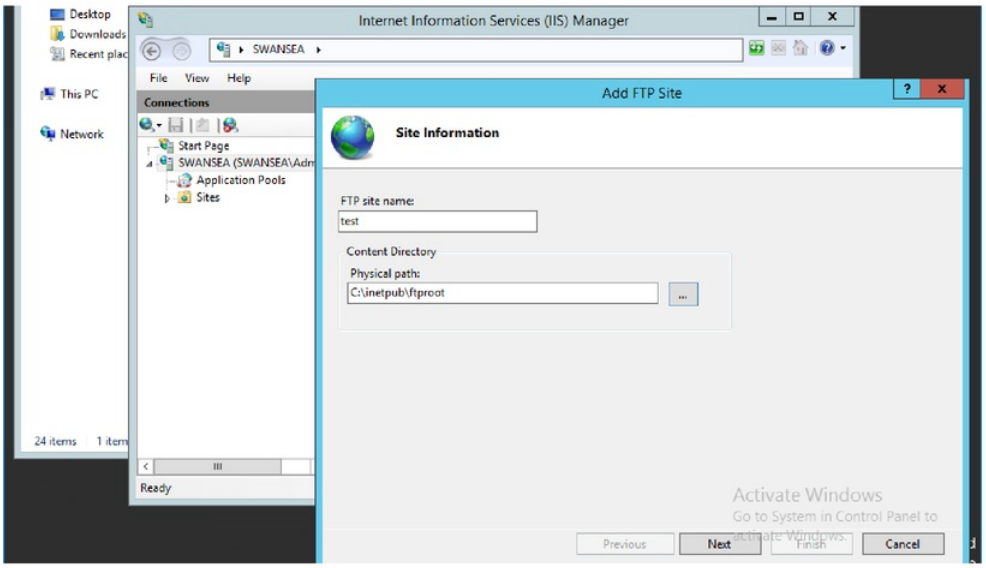


Fig. 28. Adding a new FTP site

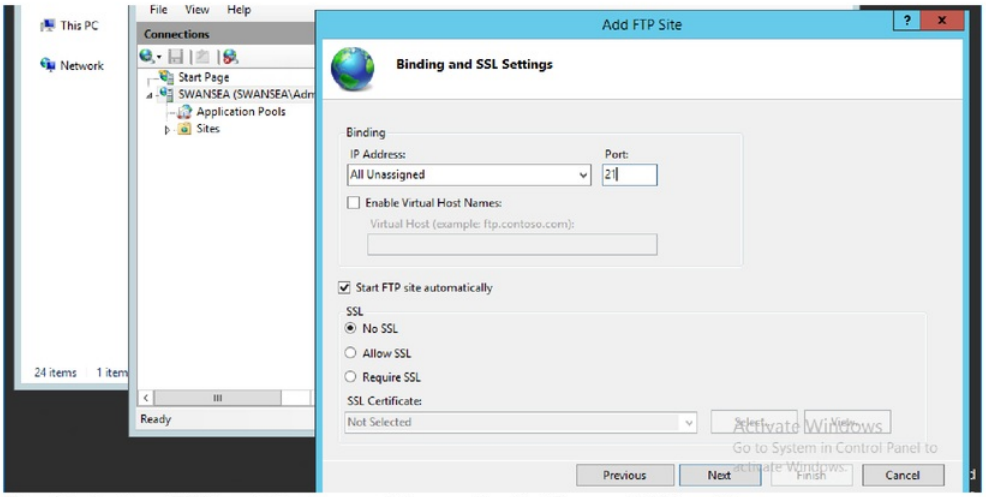


Fig. 29. Setting FTP port binding and Secure Socket Layer (SSL) settings

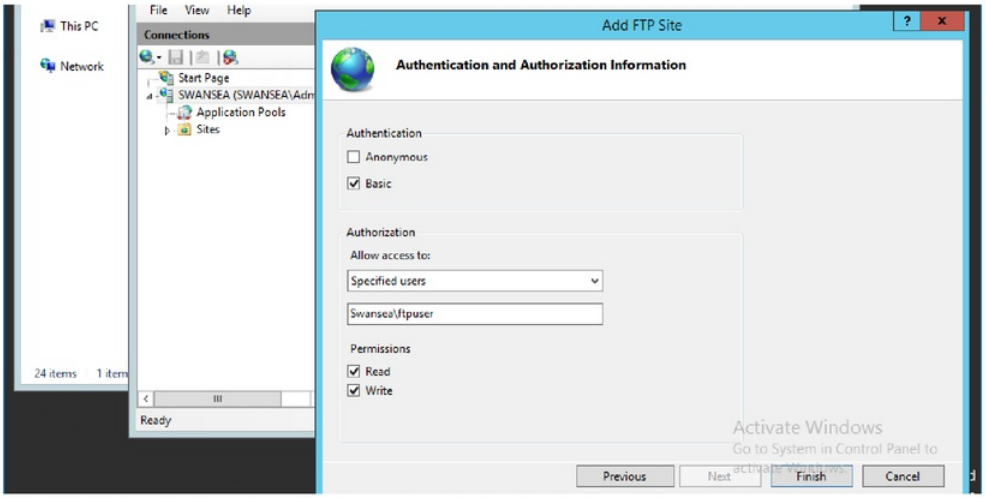


Fig. 30. Setting FTP authentication and authorisation information

Fig. 28-30. shows the process of setting up an FTP site on Swansea. Fig. 29. shows that port 21 is used as the binding. This is because port 21 is a well-known port number that is assigned to FTP (Control Channel) which will listen for incoming client connections [29], [30].

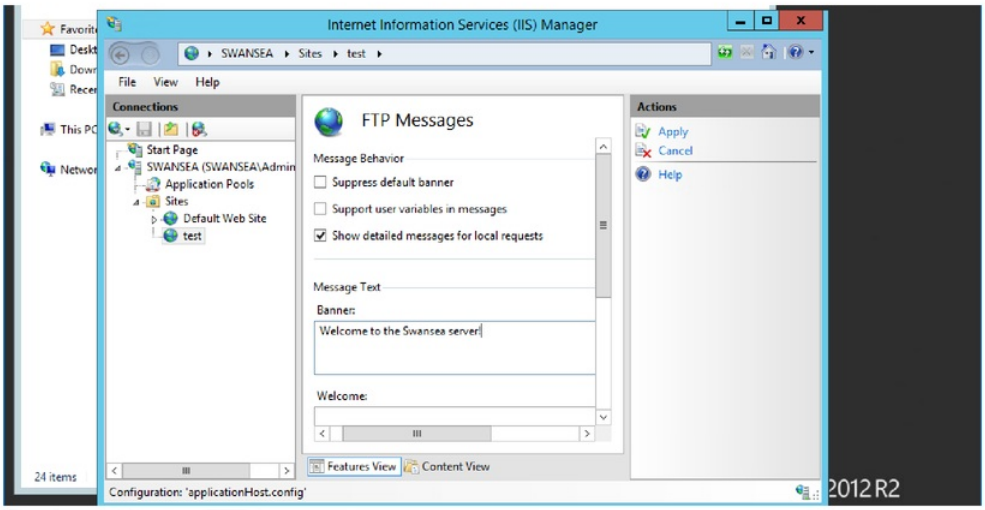


Fig. 31. Adding an FTP banner message

Fig. 31. shows the creation of a banner message that will display when a user connects to the FTP service. Banners can be used to display legal disclaimers against unauthorised access.

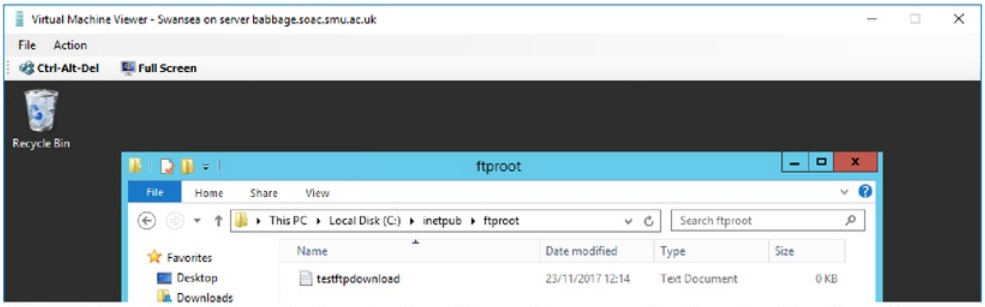


Fig. 32. Creating a test testftpdownload.txt file on Swansea to be download by Cardiff

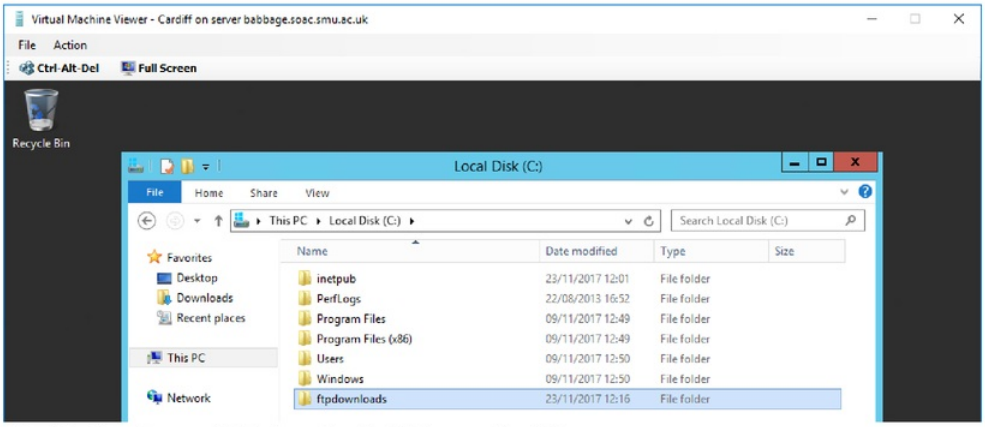


Fig. 33. Creating an FTP downloads folder on Cardiff

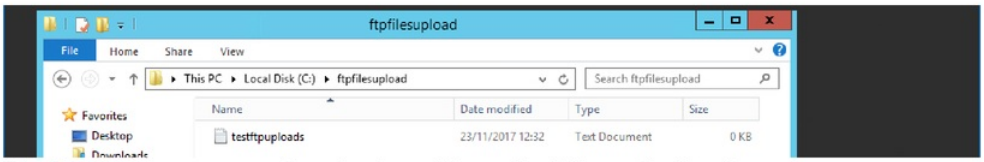


Fig. 34. Creating a test testftpuploads.txt file on Cardiff to upload to Swansea

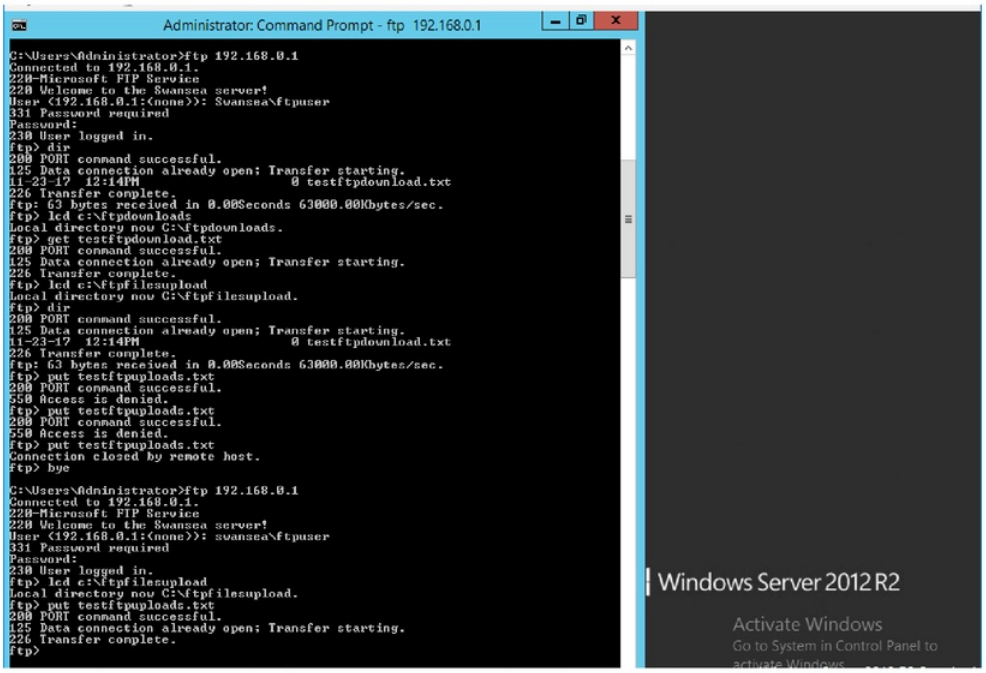


Fig. 35. Connecting to the Swansea FTP server from Cardiff to download and upload the .txt files

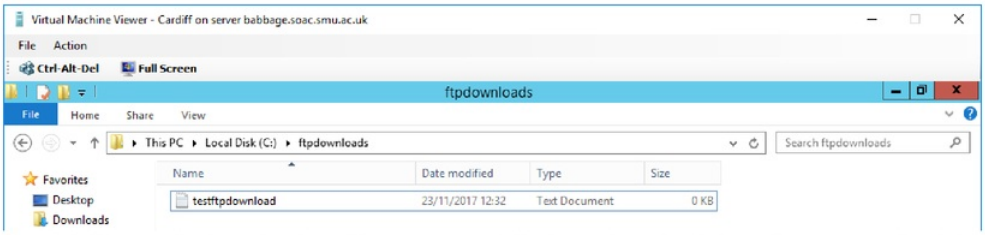


Fig. 36. The testftpdownload.txt file as successfully been download from Swansea to Cardiff

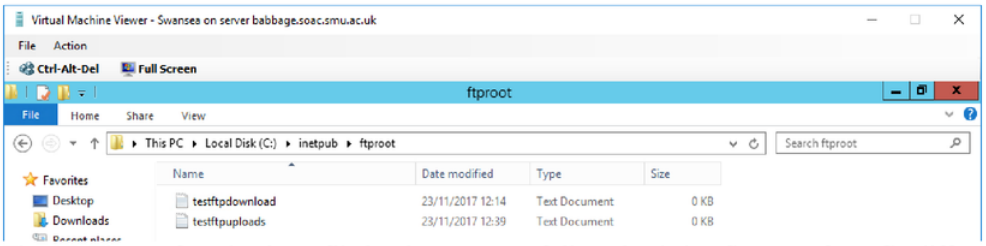


Fig. 37. The testftpuploads.txt file has been successfully uploaded to Swansea from Cardiff

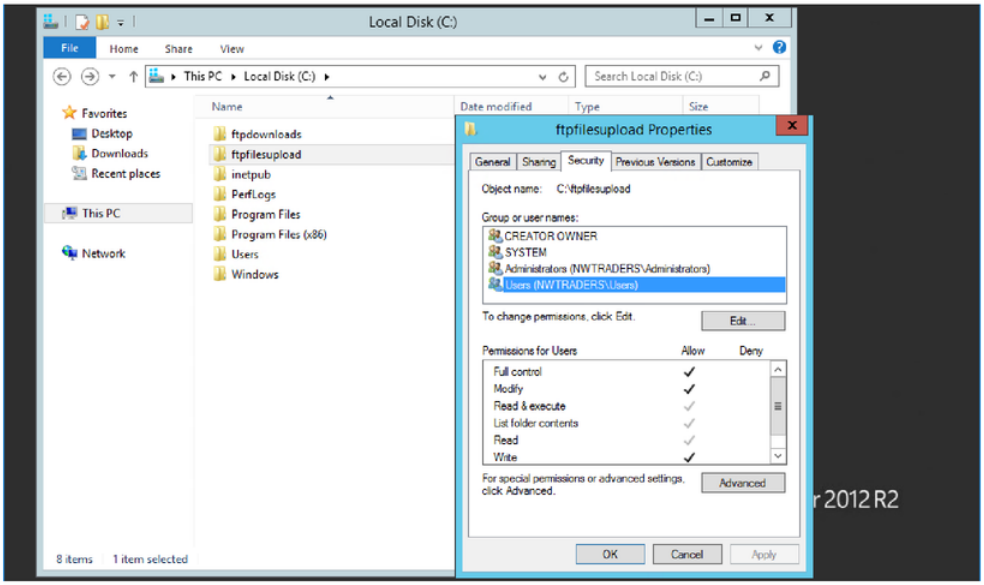


Fig. 38. ftpfilesupload folder permissions

Fig. 32-34. shows the process of creating test .txt files to test the ability to download and upload files to and from the FTP service on the Swansea server. Fig. 35. shows that the Cardiff server can successfully access the FTP service on Swansea (using username 'Swansea\ftpuser' and password 'P@ssw0rd'). The 'lcd c:\ftpdownloads' command is used to set the local directory to the ftpdownloads folder (Fig. 33). The 'get testftpdownload.txt' is used to download the testftpdownload.txt file (Fig. 32) from Swansea into the ftpdownloads folder on Cardiff. The testftpuploads.txt file is uploaded to Swansea using a similar process. First, the local directory of the Cardiff server is changed to the location where testftpuploads.txt is saved (lcd c:\ftpupload). The 'put testftpuploads.txt' command is then used to upload the file to Swansea from Cardiff. Fig. 36-37. show that the respective files have in fact been downloaded and uploaded successfully. It is worth noting the 'ACCESS IS DENIED' (Fig. 35) when first trying to upload. This is because the folder permissions of ftpfileupload need to be changed on Cardiff (Fig. 38).

2.3. Restricting ICMP Operations Using IP Packet Filtering with the Windows 2012 R2 Server Routing and Remote Access Service

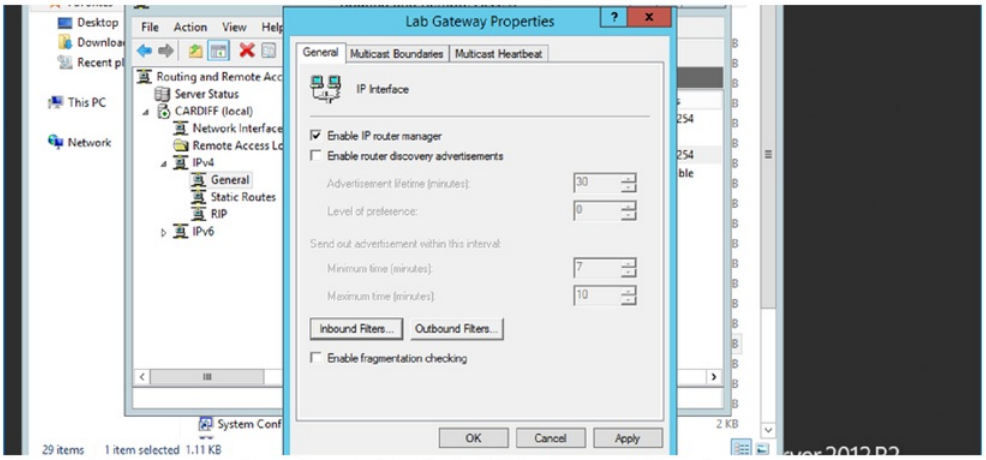


Fig. 39. Network Adapter 1 (connection to Lab) Routing and Remote Access properties

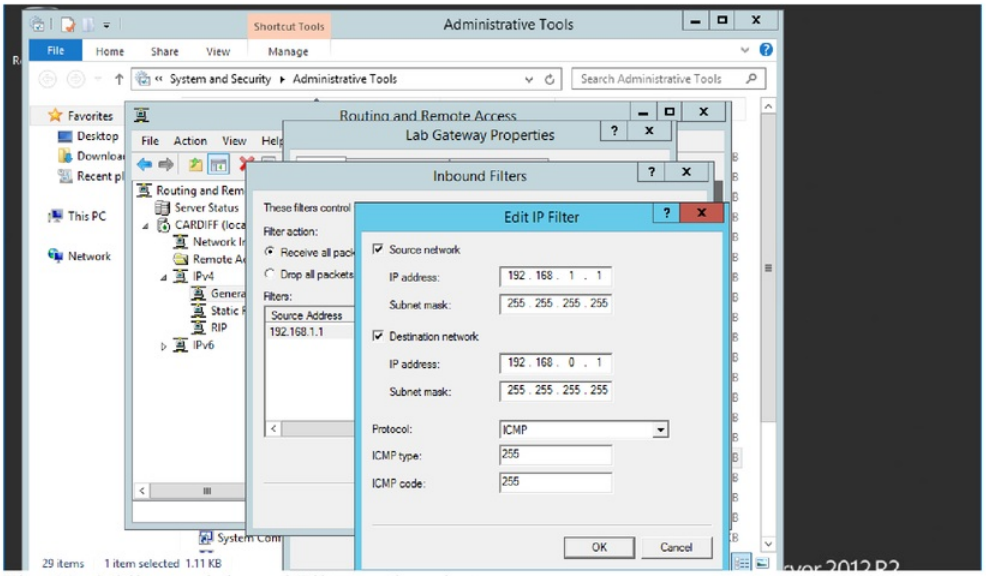


Fig. 40. Adding an inbound Filter to the adapter

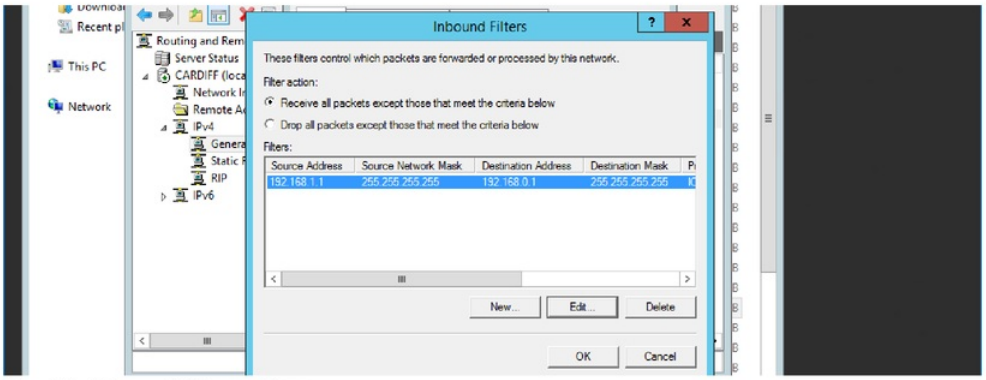


Fig. 41. Inbound filter actions

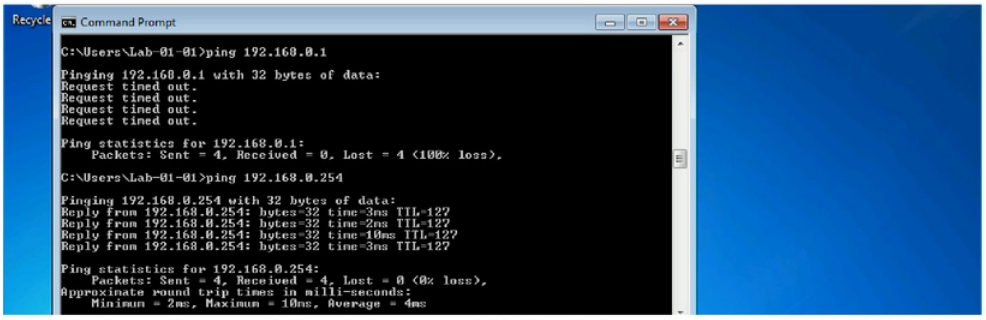


Fig. 42. Ping to Swansea fails as it is successfully blocked

Fig. 39-41. shows the process of configuring an inbound Internet Control Message Protocol (ICMP) filter that will block all ICMP (ping) packets to the destination 192.168.0.1 (Swansea server) from the source 192.168.1.1 (Lab client). ICMP filtering is a useful practice as the ICMP poses security risks. For example, ICMP flood Denial of service (DoS) attacks exploit ICMP vulnerabilities and incorrect network configurations [31]. Fig. 42. shows that the ICMP ping to Swansea from the Lab fails and has been successfully blocked. Note the ping to 192.168.0.254 (Network Adapter 2 – connection to Swansea) shows that ICMP is still working correctly.

2.4. Restricting TCP Operations Using IP Packet Filtering with the Windows 2012 R2 Server Routing and Remote Access Service

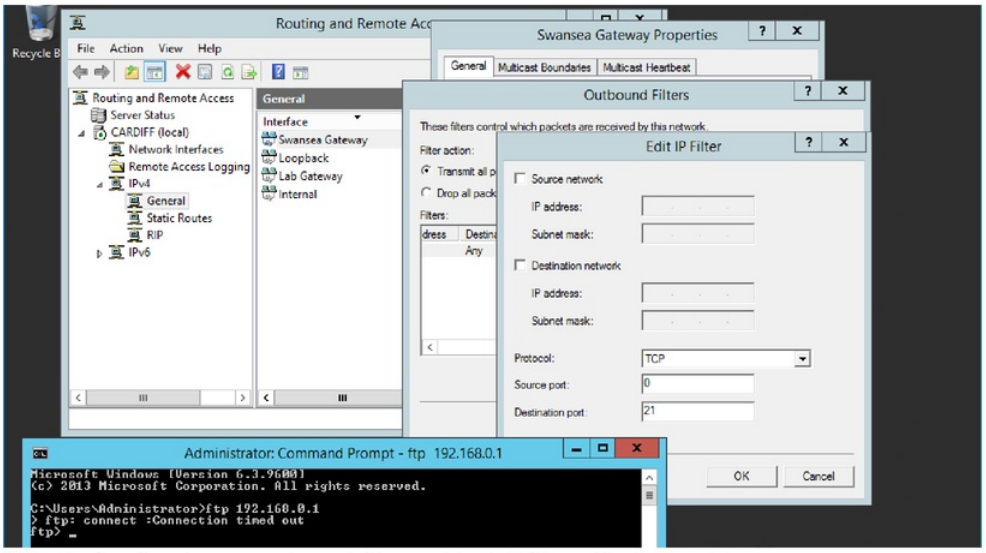


Fig. 43. Configuring an outbound filter to block TCP traffic to port 21

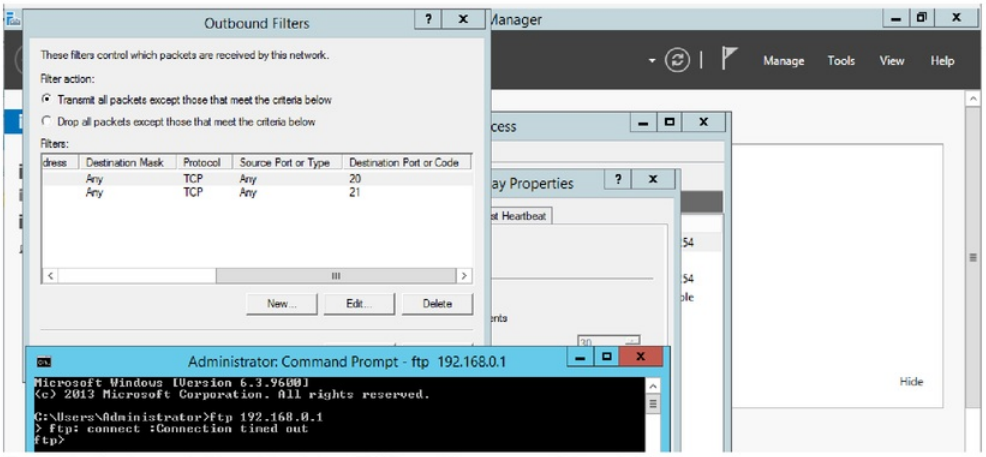


Fig. 44. Outbound filter blocking TCP traffic to port 20 and 21 from any source to any destination

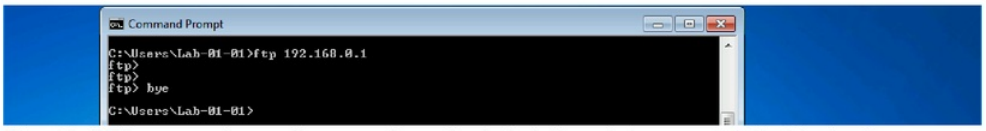


Fig. 45. FTP connection to Swansea from the Lab fails as it is successfully blocked

Fig. 43-44. show the configuration of an outbound filter to block every network (Lab and Cardiff) from accessing the FTP service running on the Swansea server. The FTP service uses the TCP protocol over port 20/21 [29] which is why destination ports 20 and 21 are blocked.

Port 21 is the FTP (Control Channel) which will listen for incoming client connections, filtering only port 21 should be sufficient to block the FTP service [30]. Fig. 44-45. shows the attempts to access the FTP service on the Swansea server from Cardiff and the Lab which are blocked (as it fails to connect). Blocking the FTP service is useful to control who can and cannot access the FTP service which is an important security consideration [31].

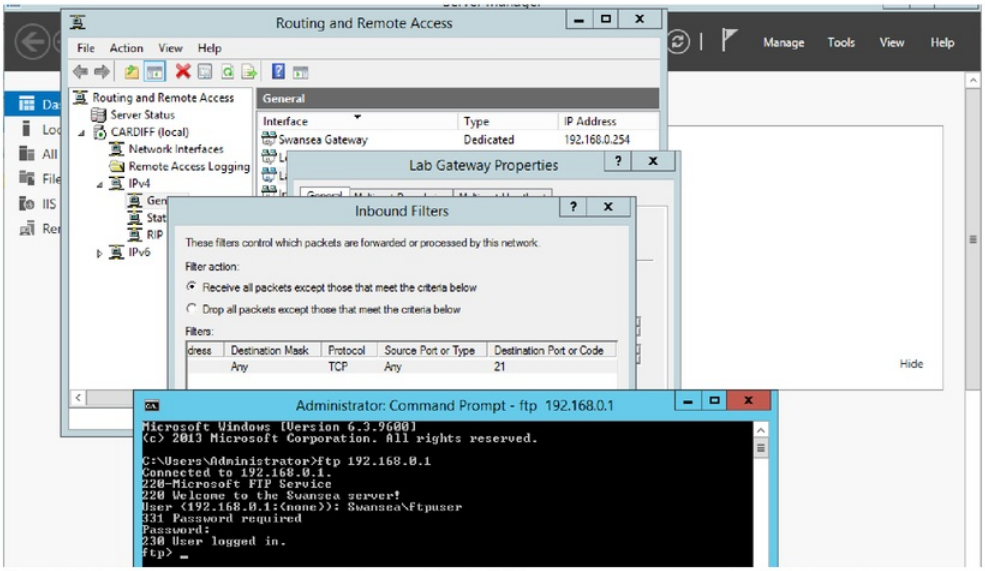


Fig. 46. Configuring an inbound filter to block FTP connections from the Lab, but allow access from Cardiff



Fig. 47. FTP connection to Swansea from the Lab fails

Fig. 46-47. shows the configuration and evidence of an inbound filter to block FTP connections from the Lab, but allow the Cardiff server to access the FTP service. Note how the inbound filter is applied to Network Adapter 1 (connection to Lab) which is why a source address is not needed as all FTP connection attempts through that adapter will be blocked. Cardiff is able to access the FTP service because its connection attempt will be sent out of Network Adapter 2 (connection to Swansea) which has no filter.

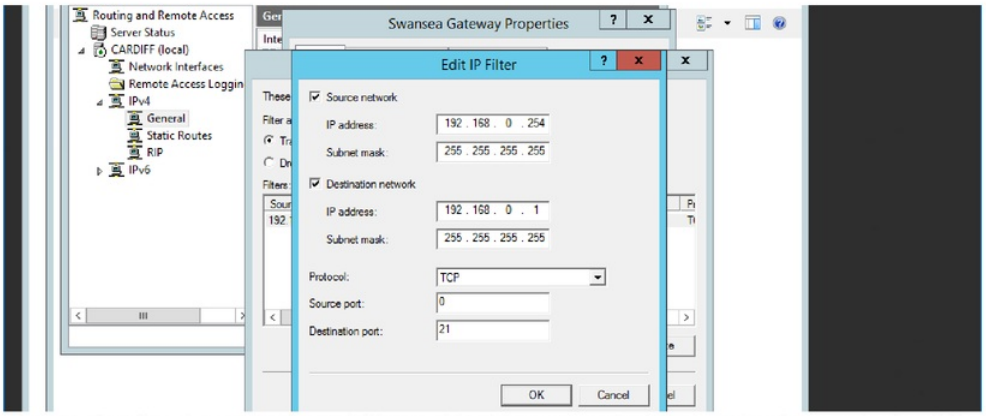


Fig. 48. Configuring an outbound filter to block access to the FTP service from only Cardiff

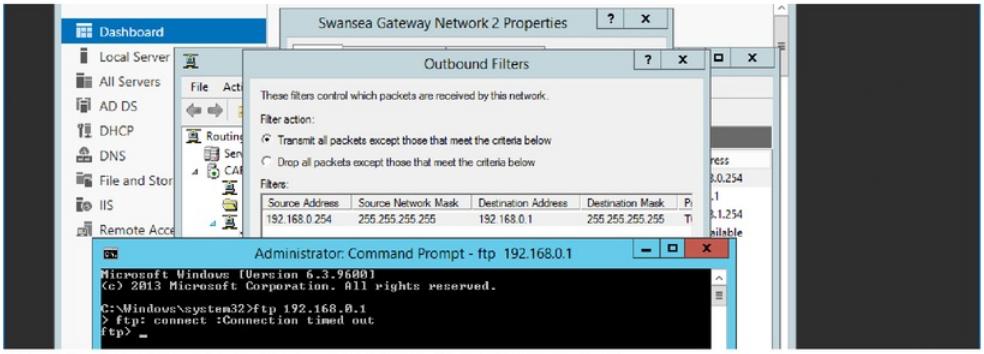


Fig. 49. Access to the FTP services from Cardiff is blocked

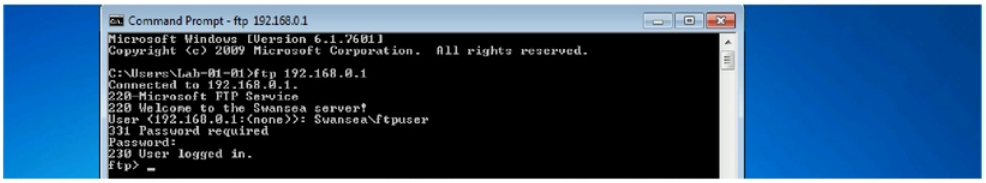


Fig. 50. The FTP service can successfully be accessed from the Lab

Fig. 48-50. shows the configuration and evidence of an outbound filter on Network Adapter 2 (connection to Swansea) to block FTP attempts from only the Cardiff server, allowing the Lab to connect. The 192.168.0.254 source address is significant as that will be the source address when connecting to the FTP service from Cardiff.

2.5. Allowing only Cardiff to FTP to Swansea Server while blocking all other traffic from Cardiff and the Lab network

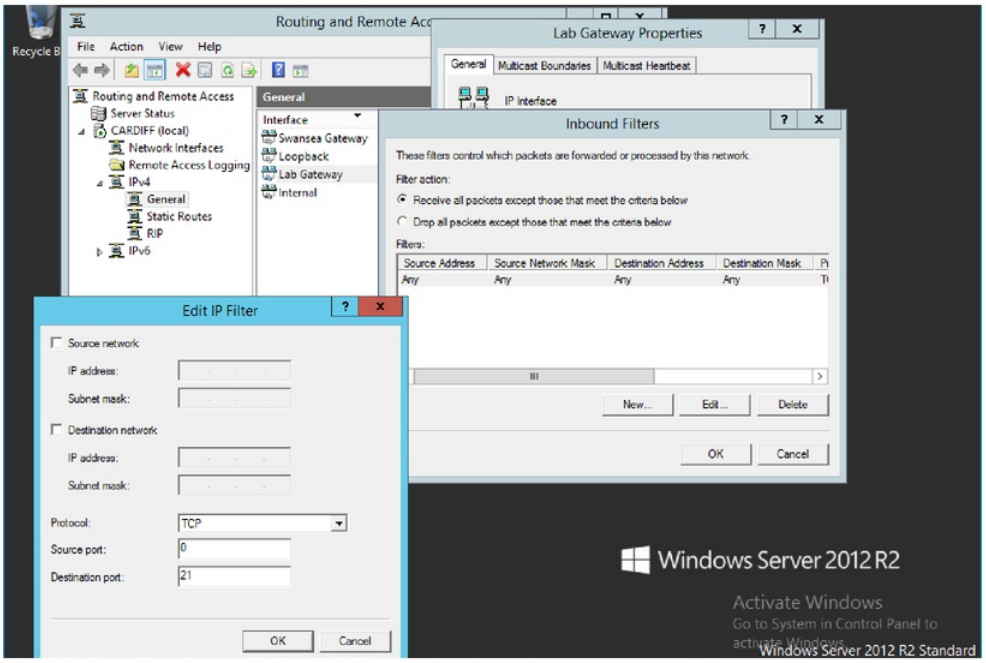


Fig. 51. Configuring an inbound filter to block all FTP connection attempts from the Lab

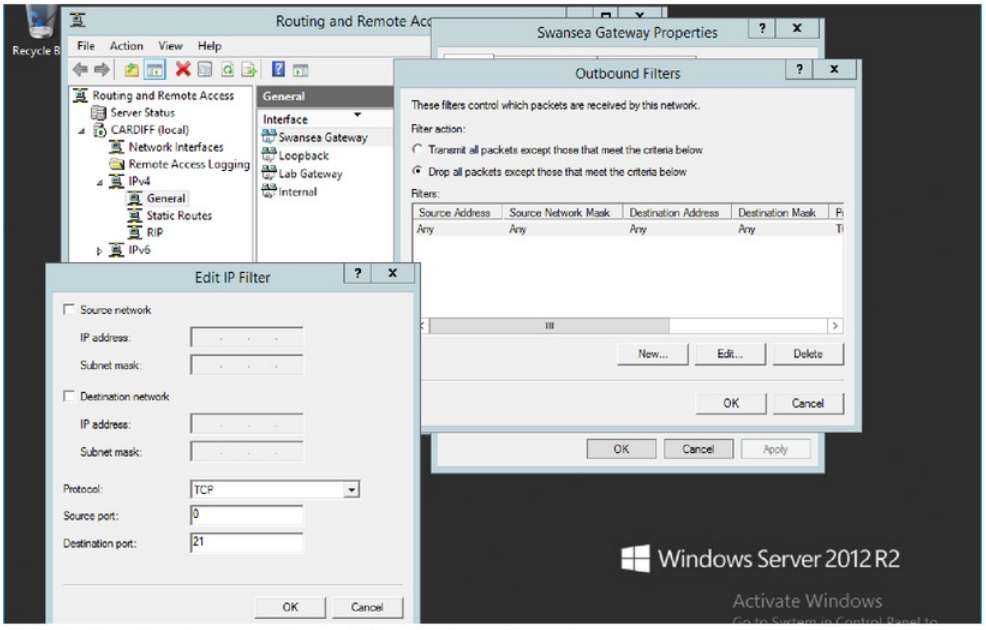


Fig. 52. Configuring an outbound filter to drop all packets that are not FTP connections

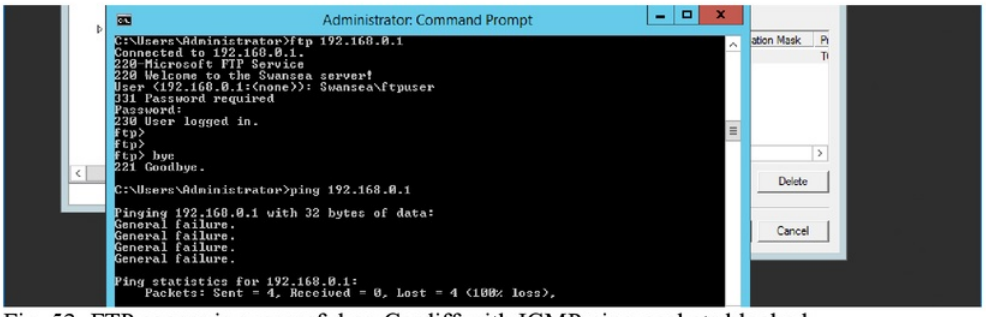


Fig. 53. FTP access is successful on Cardiff with ICMP ping packets blocked

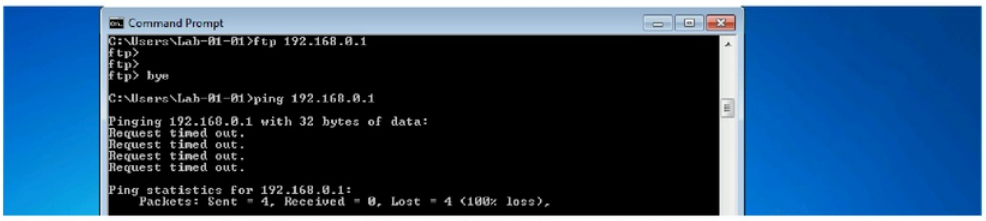


Fig. 54. FTP access denied on the Lab with ICMP ping packets blocked

Fig. 51. shows the configuration of an inbound filter on Network Adapter 1 (connection to Lab) that will block all FTP connections entering through that adapter (Lab). Fig. 52. shows the configuration of an outbound filter on Network Adapter 2 (connection to Swansea) that will drop all packets (block) that are not for TCP port 21. These two filters work together so that only FTP traffic from Cardiff can be successfully received by Swansea. Fig. 53. shows how the Cardiff server can FTP to Swansea but is unable to send ICMP packets. Fig. 54. shows how FTP attempts and ICMP packets are both blocked. This is an example of a good security policy by only allowing the traffic that is required to Swansea.

2.6. ARP, Packet Filtering and Server Roles Observation Tasks

2.6.1. ARP Process

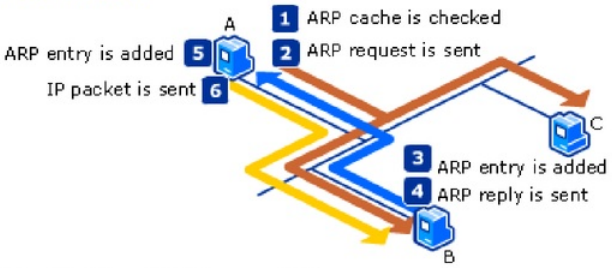


Fig. 55. The ARP Process [32]

The Address Resolution Protocol (ARP) is a process that is used by a host on a LAN to resolve logical IPv4 addresses to physical Media Access Control (MAC) addresses [33]. This is a necessary process as the underlying ethernet hardware communicates using MAC addresses [34]. The ARP process stores the IPv4 and MAC mappings in an ARP cache (table) [33].

The ARP process is comprised of a number steps (Fig. 55) [32], [33]:

- A host will check its ARP cache for an entry for the destination IP address or gateway (next-hop) address in the case where the destination IP address is located on a remote network. If there is an entry ARP skips to step 6 in Fig. 55.
- If no entry is found an ARP Request frame is broadcast out of the appropriate interface. The ARP frame contains the MAC address and IP address of the interface sending the ARP request as well as the destination or gateway IP address.
- All local hosts will receive and process the broadcast ARP frame. The receiving hosts will check to see if the requested IP address (destination/gateway) matches their own IP address. If there is a match the host updates its ARP cache with the address mapping of the sender of the ARP request.
- If there is not a match, the ARP request is dropped.
- The (matching) receiving host sends an ARP Reply containing its requested MAC address and sends it directly back to the sender of the ARP request.
- When the ARP reply is received by the sender, it updates its ARP cache with the address mapping.
- Both hosts involved in the ARP request and reply now have up-to-date address mappings in their ARP caches.
- ARP now sends the IP packet to the destination address or gateway by addressing it to the resolved MAC address in its cache.

2.6.2. Packet Filtering

Packet filtering is a technique used to control network access by allowing or disallowing the flow of specific types of IP traffic [35], [36]. Packet filtering is normally conducted on a router (which Cardiff is acting as). In our context packet filtering is termed IP packet filtering as we are filtering traffic based on source and destination IP address. Packet filtering is a security measure and can be used in networks as a first-line of defence to block specific inbound packets from reaching a protected network. The benefit of this is that it can be used to reduce the load on an application firewall or proxy [37]. Other advantages include application independence, high performance and scalability. Disadvantages include lack of security as packet filtering is not application aware and offers no monitoring above the network layer [38]. Outbound filtering can be used to limit the number of users accessing a service e.g. mail server or the Internet [37].

2.6.3. Contradictory filters

Contradictory filters can occur as inbound and outbound filters can be configured for each interface (adapter) [35]. For example, Network Adapter 1 (connection to Lab) could have an inbound filter that allows traffic but the outbound filter on Network Adapter 2 (connection to Swansea) might not allow the same traffic to be sent. This would mean that traffic would not pass across the Cardiff router.

2.6.4. Server Roles and Services

There are many roles and services that can be configured on Windows Server 2012 R2 [39]. For example, there are many Active Directory services that can be used to automate network management of user data, security and distributed resources [40], [41]. Active Directory can also control access for users based on security policy [42]. When using Active Directory machines would be connected to a domain where information is stored in a central location [43]. Such features would be advantageous

in a local and remote context to control access and manage user data. A global administrator would also be able to oversee everything that happens in the domain [43]. There are many Networking services such as Dynamic Host Configuration Protocol (DHCP) to automate IP address assignment to clients and Domain Name System (DNS) to map domain names to IP addresses [44]. These services will be explored in sections 3 and 4.

Other roles include an Application Server for running custom server-based business applications [39]. The Failover Clustering service can be used to increase the availability of applications and other services e.g. local web server [45]. The File and Storage Services could be used to provide centralised network storage for local users [46]. The Remote Desktop Service could also be especially useful to allow remote access e.g. virtual desktop infrastructure (VDI) and session-based desktops [47].

Section 3: TCP/IP Configuration using DHCP (Dynamic Host Control Protocol)

3.1. Understanding Different TCP/IP Configurations

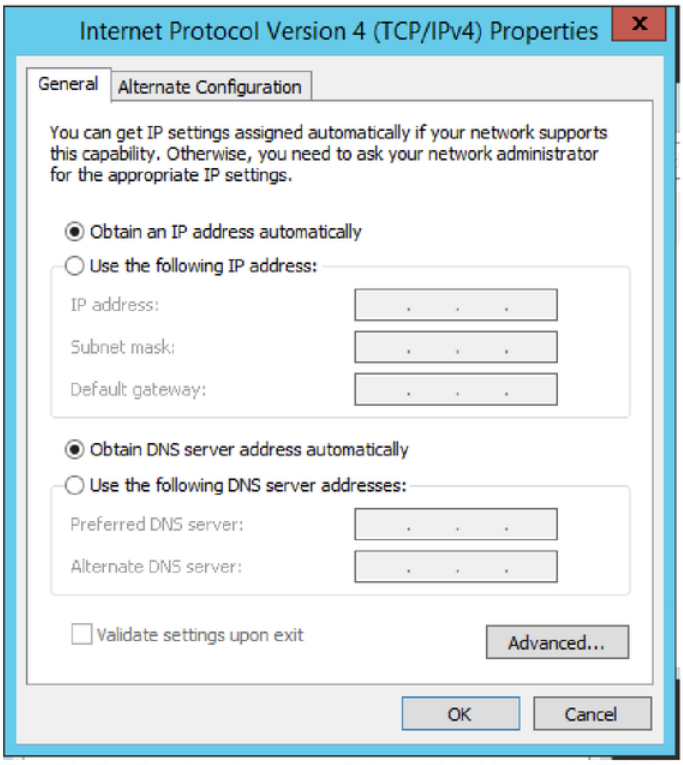


Fig. 56. Configuring Automatic Private IP Addresses (APIPA) on the Lab and Cardiff before setting up a DHCP server

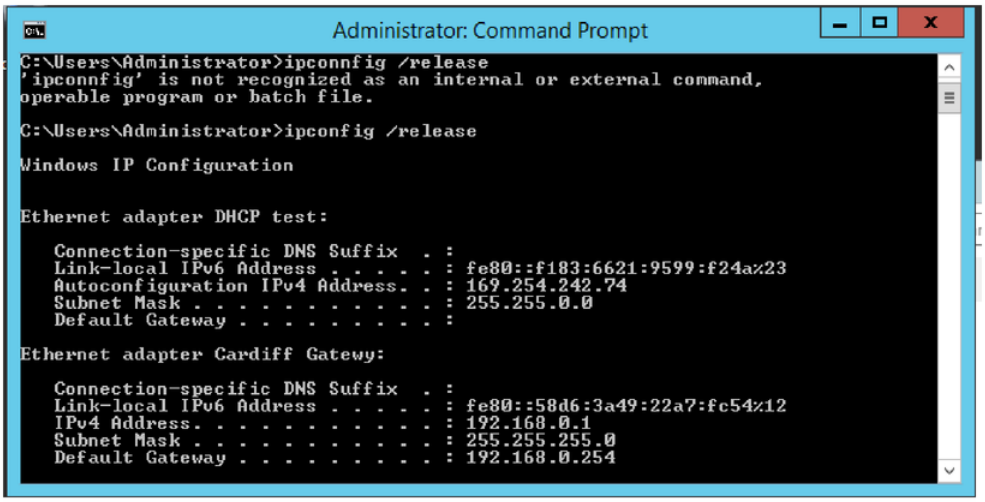


Fig. 57. The ipconfig /release command on the Lab machine

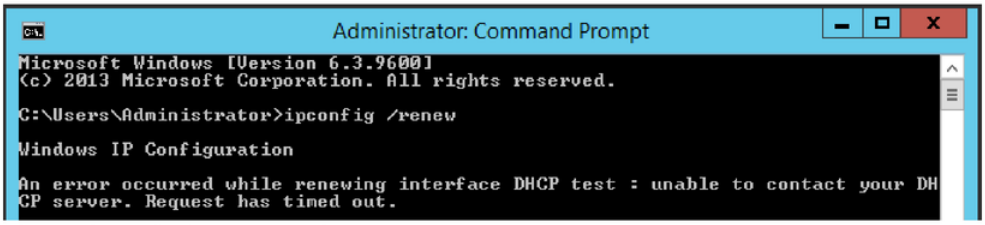


Fig. 58. ipconfig /renew on Lab

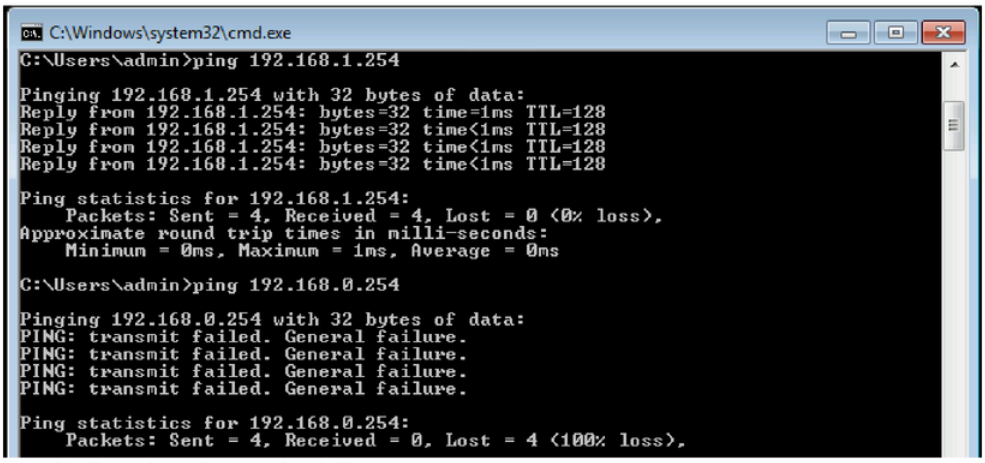


Fig. 59. Ping from Lab to Cardiff

Fig. 56. shows the configuration of the Lab and Cardiff network adapters to automatically obtain private IP address (APIPA) before setting up the DHCP server. This is a test to show that APIPA will fail without the DHCP server. Fig. 57. shows the use of ipconfig /release command to make sure the previous static IP address is erased. Fig. 58. shows the ipconfig /renew command to try and obtain a new IP address automatically from a DHCP server. There

is no DHCP server at this stage causing the process to fail and time out. Fig. 59. shows that pings from the Lab to the different adapters on Cardiff fails. This is because as the APIPA process failed the Lab network adapter no longer has a default gateway and can no longer contact Cardiff (default gateway) or remote networks (Swansea).

3.2. Installing the DHCP Server

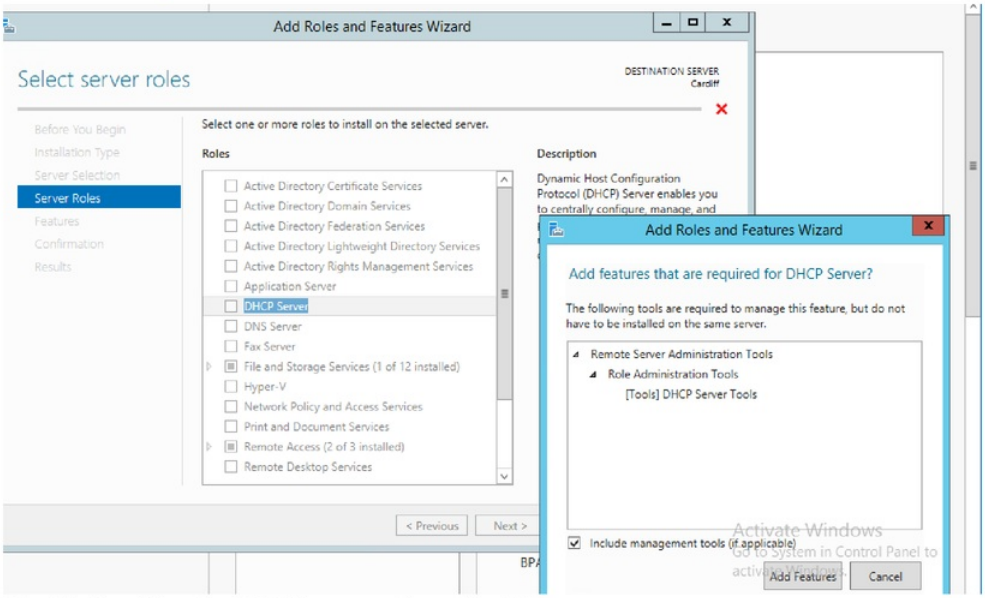


Fig. 60. Installing the DHCP server role on Cardiff

Fig. 60. shows the installation of the DHCP server role and related management tools on the Cardiff server.

3.3. Authorising the DHCP Server

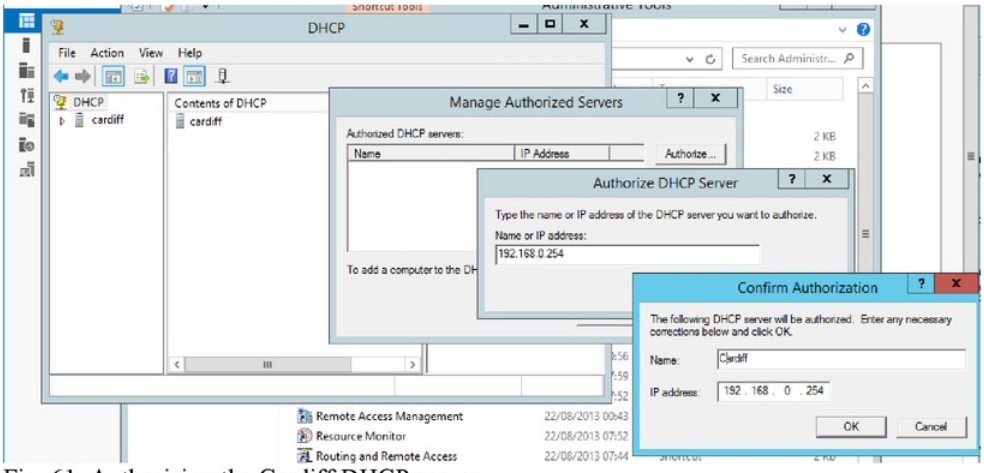


Fig. 61. Authorising the Cardiff DHCP server

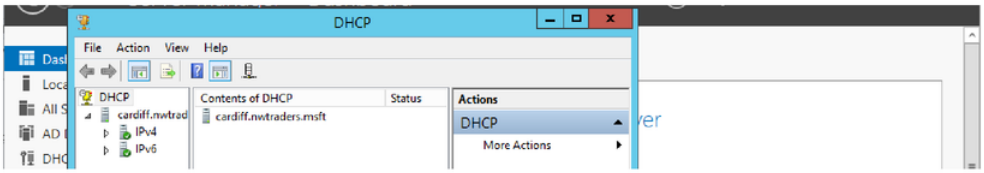


Fig. 62. Confirmation of successful authorisation

Fig. 61. shows the Windows Active Directory (AD) Domain authorisation for the Cardiff DHCP server. This is because Windows AD Domains protect against rogue DHCP servers that may be added to the network. This is a precaution as DHCP clients are not concerned about which DHCP server offer they respond to [48]. The green tick in Fig. 62. shows that the Cardiff server has been successfully authorised.

Note: This step is only completed after adding the servers/client to a domain and configuring Active Directory.

3.4. DHCP Scopes

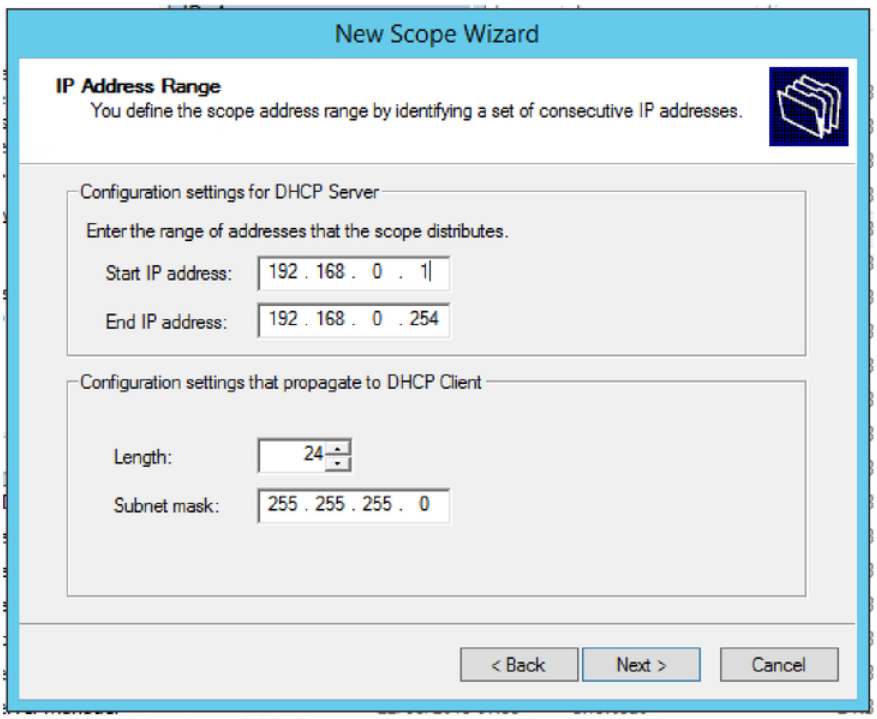


Fig. 63. Configuring a DHCP scope for the Swansea network

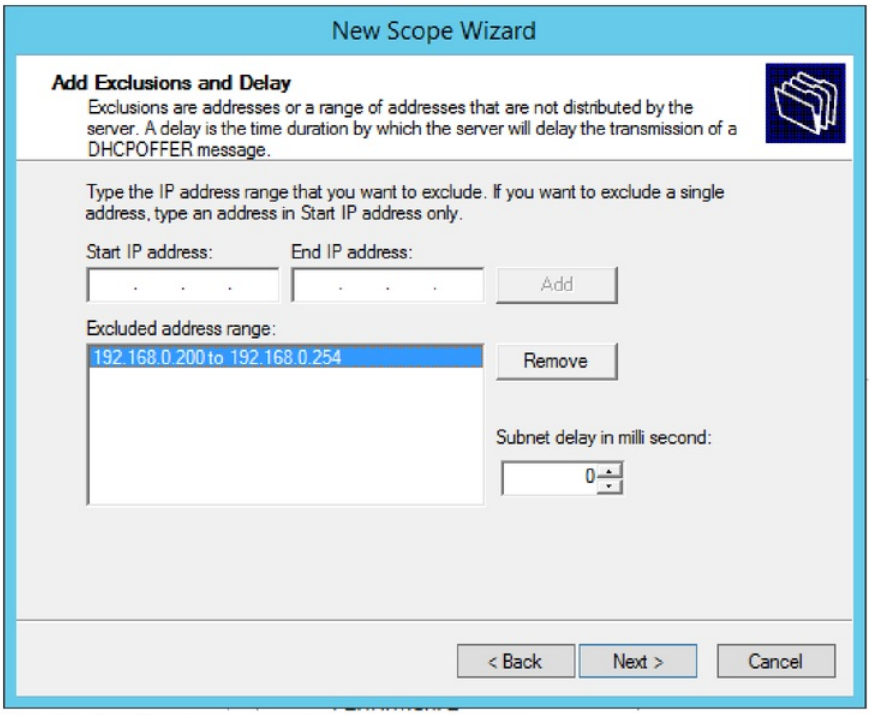


Fig. 64. Configuring an IP address exclusion range

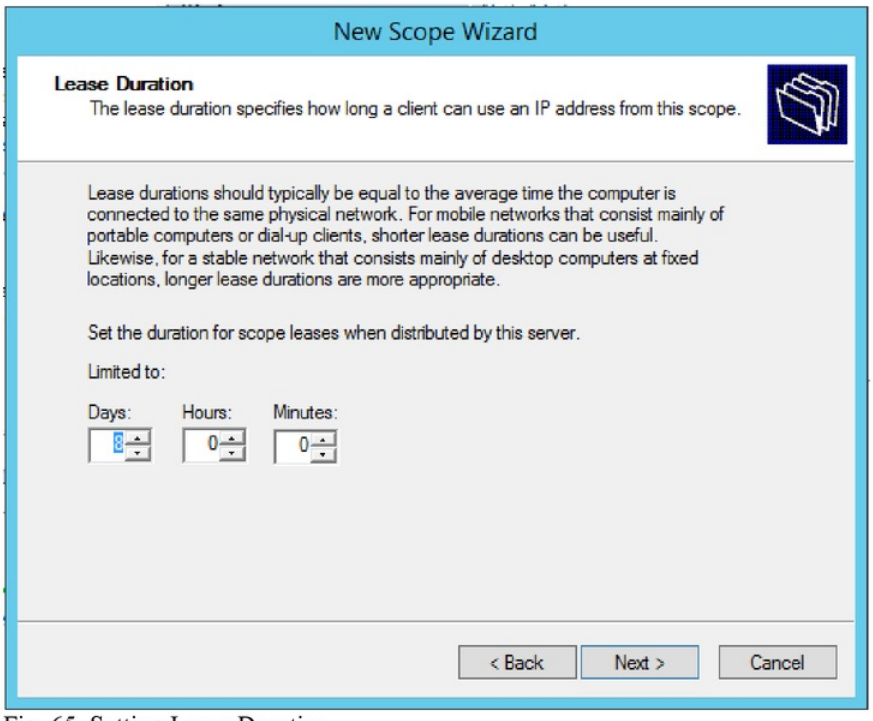


Fig. 65. Setting Lease Duration

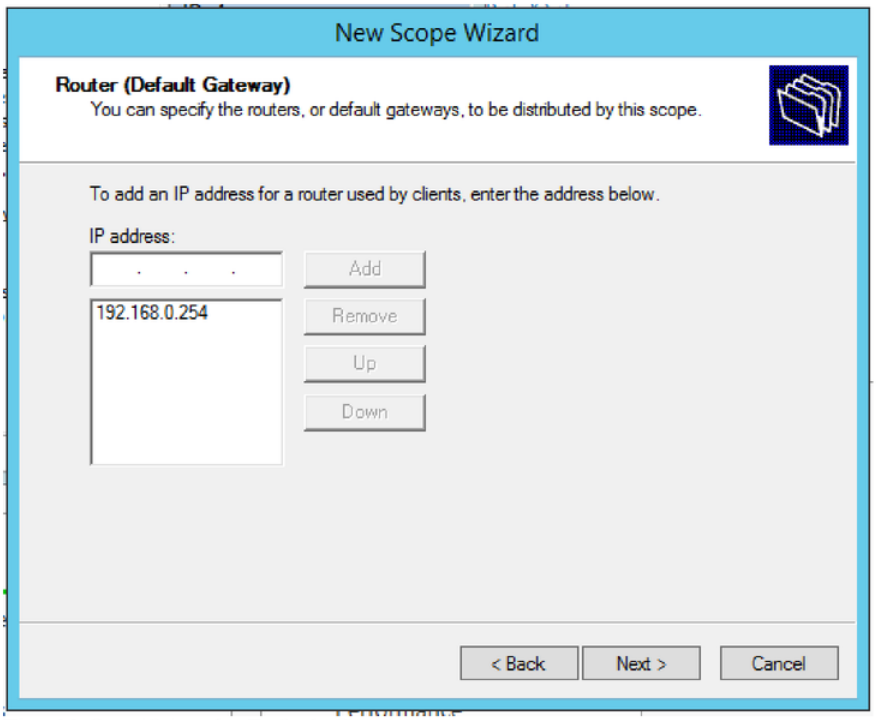


Fig. 66. Specifying the Default Gateway

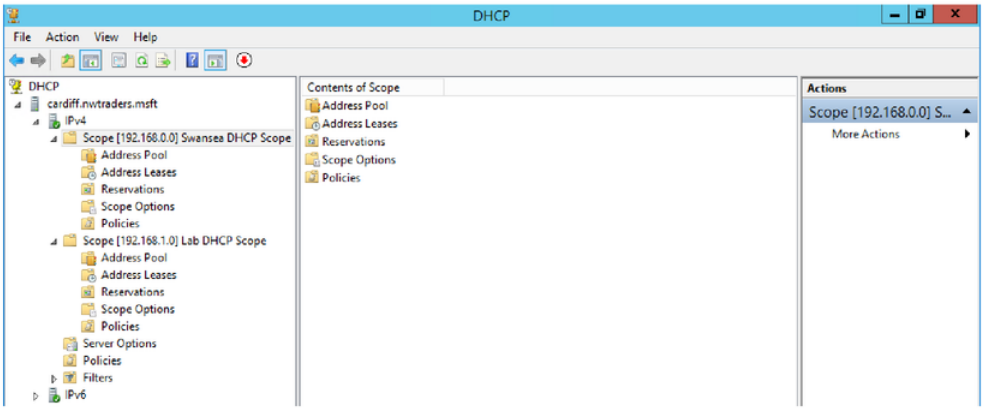


Fig. 67. Scopes for the Swansea and Lab network added to DHCP server

Fig. 63-66. shows the step-by-step process of configuring a DHCP scope (right clicking IPv4 and selecting 'New Scope' in the DHCP Administration Tools window). A scope is a range of IP addresses that are valid to be leased to DHCP clients within the network that the range is part of [49]. A scope must be set for DHCP to work. Fig. 64. shows an excluded address range for IP addresses that will not be distributed by the server. This is useful as these excluded addresses can be used as static addresses for Web servers, Print servers etc. where you do not want the IP address to possibly change if a DHCP lease expires [50]. Fig. 65. shows that the lease duration is set to 8 days. This is the length of time a DHCP client will have the IP address

it has been assigned. A client must renew the lease before it expires, or it will be assigned a new IP address from the scope. Fig. 66. shows the default gateway address that the DHCP scope will be distributed from.

IMPORTANT: A new scope will be now be made for the Lab network using the IP address range 192.168.1.1-192.168.1.254, exclusion range 192.168.1.200-192.168.1.254, lease duration of 8 days and default gateway address 192.168.1.254.
Two DHCP scopes will now be visible in the DHCP Management Console as seen in Fig. 67.

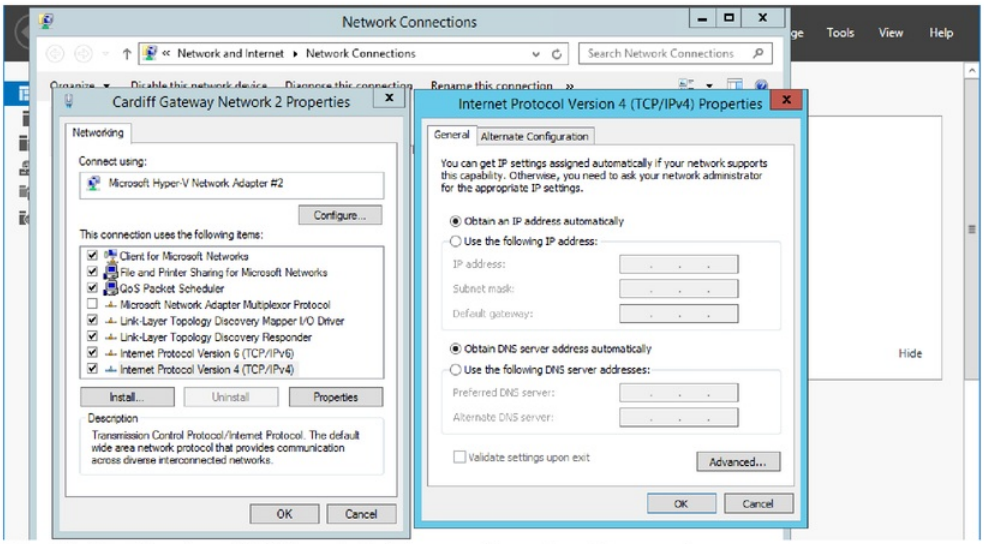


Fig. 68. Configuring APIPA on Lab (also configured on Swansea)

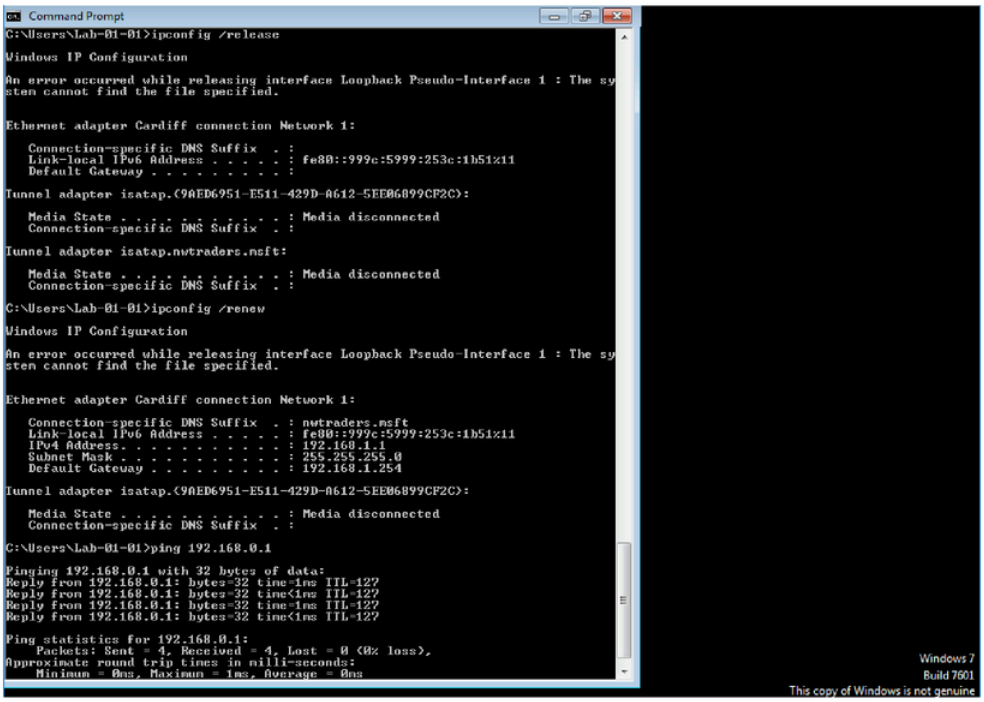


Fig. 69. Using the ipconfig /release and ipconfig /renew commands on the Lab

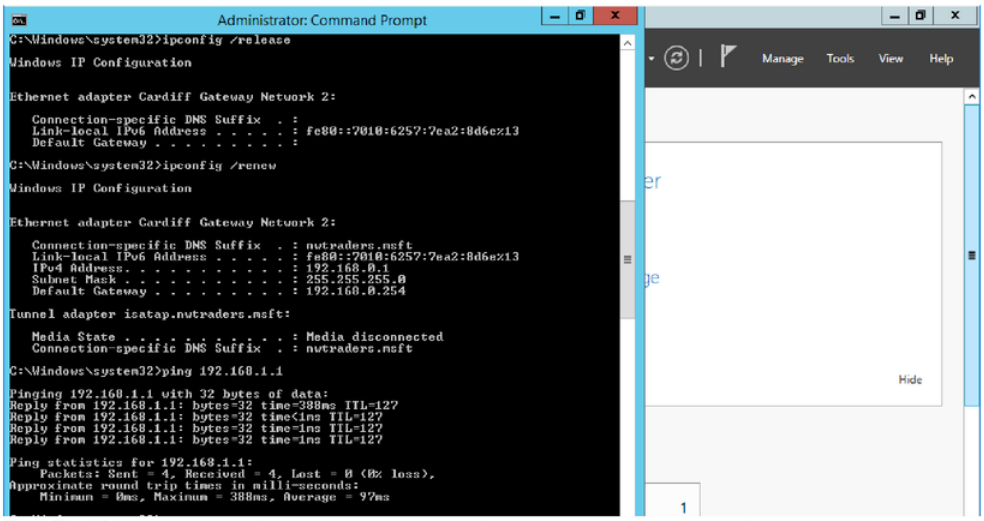


Fig. 70. Using the ipconfig /release and ipconfig /renew commands on Swansea

Fig. 68-70. show the final steps conducted on the DHCP clients (Swansea and the Lab) to configure and verify the success of DHCP address allocations. Firstly, the APIPA option is selected on the network adapter connected to Cardiff on both Swansea and the Lab (Fig. 68). The ipconfig /release and ipconfig /renew commands are then used on both Swansea and the Lab (Fig. 69-70). A ping test is then connected on both the Lab and Swansea to verify connectivity.

3.5. Configuring DHCP reservations

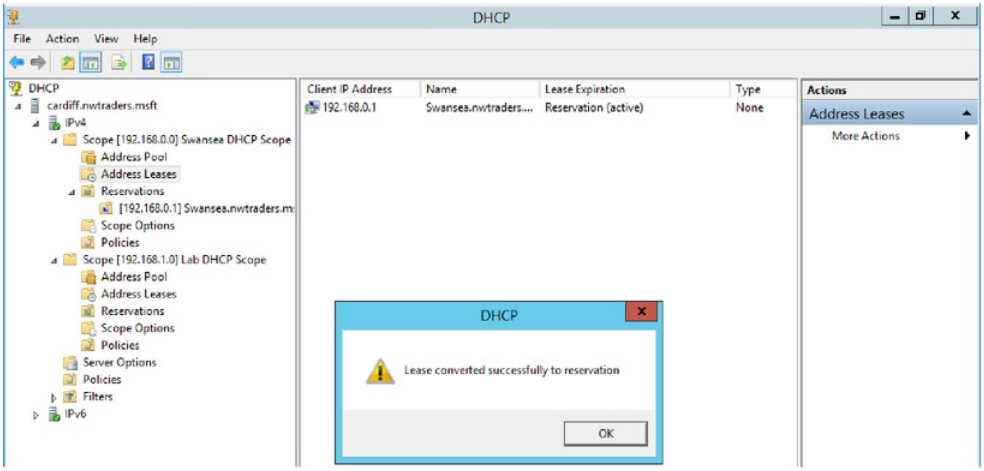


Fig. 71. Converting the Swansea server's lease into a reservation

Fig. 71. shows the easiest way to create a DHCP reservation (right click lease and select 'Add to reservation'). A DHCP reservation means that a client will always have the same reserved IP address [50]. It is important to use reservations on any clients that need a permanent IP address assignment (e.g. Print servers, File servers, Web servers etc.) [50]. These clients require a permanent IP address to avoid possible connection problems that may occur when using an unreserved address that could change [51]. Such clients could also be assigned IP addresses statically using an IP address in the DHCP exclusion range.

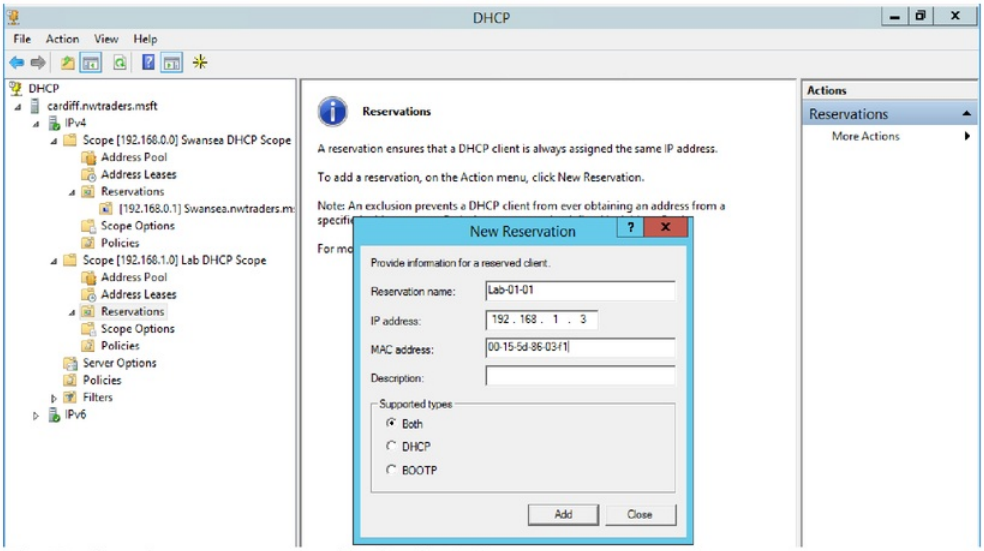


Fig. 72. Creating a new reservation for the Lab

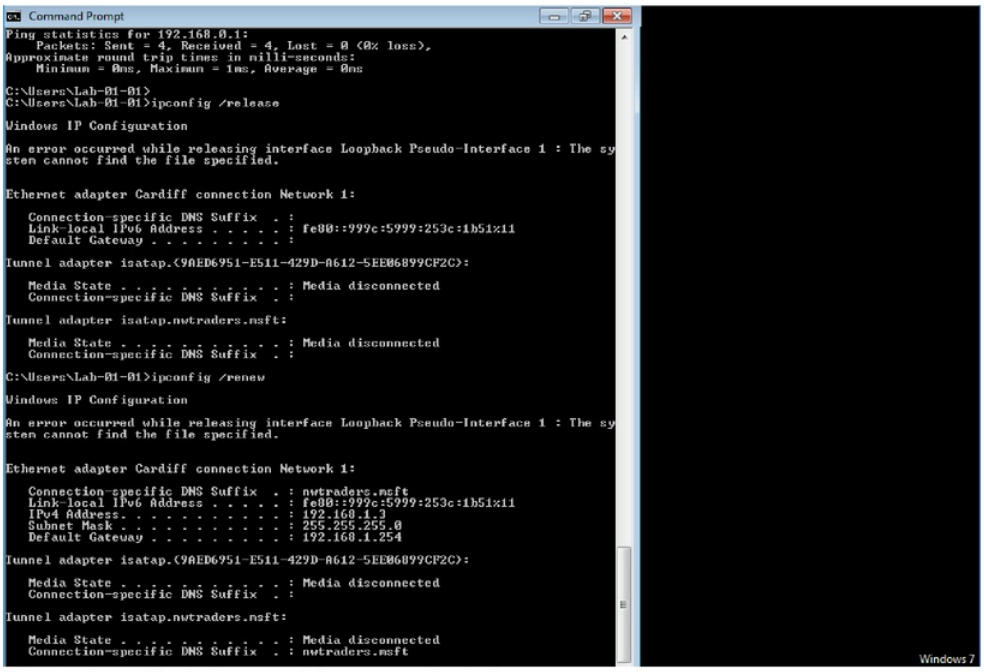


Fig. 73. ipconfig /release and /renew on the Lab

Fig. 72. shows the manual configuration of a reservation for the Lab using its MAC address. A MAC address is unique to a machine allowing the DHCP server to identify the client it is reserving the IP address for. Fig. 73. shows that the IP address of the Lab is the reserved IP address of 192.168.1.1 after ipconfig /release and /renew.

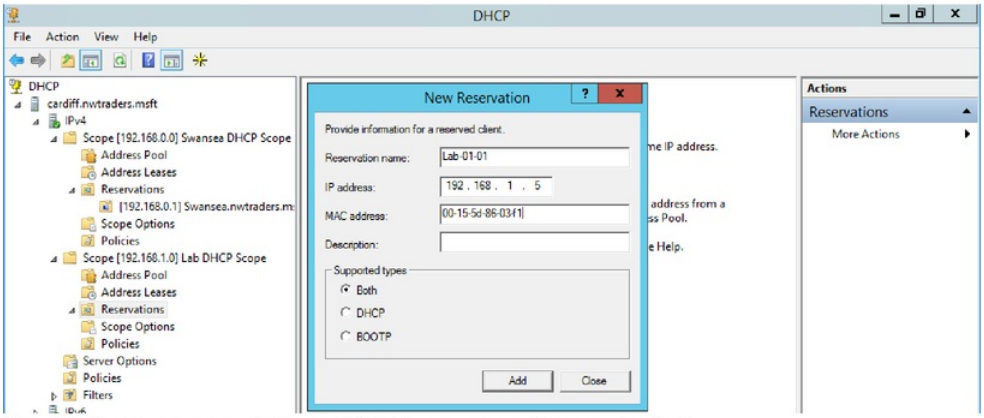


Fig. 74. Testing a new, different IP address reservation on the Lab

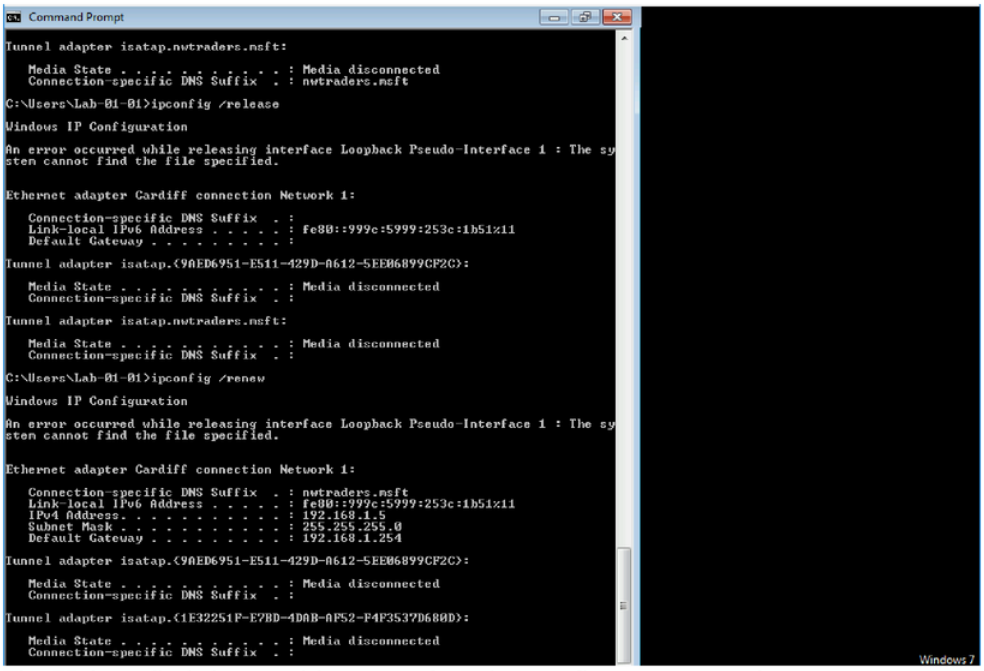


Fig. 75. ipconfig /release and /renew on the Lab to confirm new reservation

Fig. 74-75. shows another example of a different reservation (192.168.1.5) on the Lab machine.

3.6. DHCP Observational Tasks

3.6.1. Role of the DHCP Server

Dynamic Host Configuration Protocol (DHCP) is a client/server protocol that automatically assigns IPv4 addresses and other information such as subnet mask and default gateway to IP hosts on a network [52]. The DHCP server will also ensure that all assigned IP addresses are unique [33]. In the author's presented network, the Cardiff server acts as the DHCP server to provide the Lab and Swansea server Network Interface Cards (NICs) with IPv4 addresses, subnet mask and default gateway configuration information.

3.6.2. Troubleshooting DHCP

If the DHCP server cannot be contacted while the Lab and Swansea are configured to obtain their IP address automatically. The Lab and Swansea will not have an IP address configured or have the IP address 0.0.0.0 and as result network connectivity will be lost [53]. For this reason, it may be worth considering clustering DHCP servers to provide high availability. This means that when a failure is detected the DHCP service will quickly be restarted on a reserve server [54].

3.6.3. The DHCP Lease Process

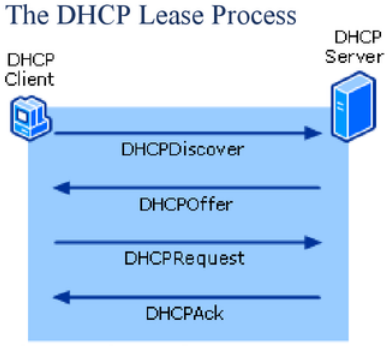


Fig. 76. The DHCP 4 stage process [55]

The DHCP lease process is comprised of 4 stages (Fig. 76.) [33]:

1. When a configured DHCP device (client) boots or connects to the network it will broadcast a DHCP discover (DHCPDISCOVER) message to identify any available DHCP servers on the network.
2. A DHCP server will then reply with a DHCP offer (DHCPOFFER) message, offering a lease to the client. The DHCP offer contains the IPv4 address and subnet mask as well as the default gateway and DNS address that will be assigned. The duration of the lease will also be included.
3. The client then sends a DHCP request (DHCPREQUEST) message that identifies the specific DHCP server and lease offer that the client is accepting. A client may also request an address it has been allocated before by the DHCP server. Note a client may receive multiple DHCP offers if there are multiple DHCP servers on the LAN, meaning the client will have to choose which offer to accept.
4. The DHCP server will then send a DHCP acknowledgment (DHCPACK) message to acknowledge to the client that the lease has been finalised. If the DHCP offer is no longer valid or available for the client, the DHCP server will reply with a negative acknowledgment (DHCPNAK) message. If this occurs the DHCP process must start again with a new DHCPDISCOVER message. When a client has a lease, it must renew it before the lease expires with another DHCPREQUEST message.

3.6.4. DHCP Relay Agent

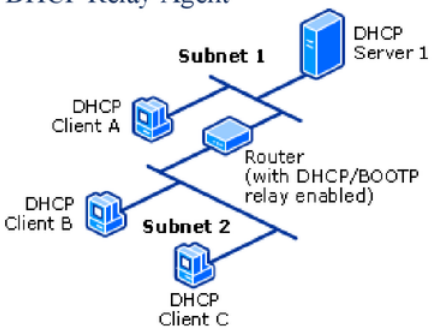


Fig. 77. DHCP Relay Agent on a router

A DHCP Relay Agent is a Bootstrap Protocol (BOOTP) that relays DHCP messages between DHCP clients and DHCP servers on different IP networks [56]. DHCP is a broadcast based protocol meaning DHCP messages will not normally pass through routers [56]. A DHCP Relay Agent on a router will receive DHCP broadcast messages (the DHCP protocol allows every message to be broadcast in certain specific cases [57]) on the network/subnet. The relay agent will then forward them to the stated IP address on a different network/subnet [56]. The advantage of having relay agent forwarding on a router is that a single DHCP server will be able to support many networks/subnets [58]. This will eliminate the need to have a DHCP server on every subnet in an internetwork [59]. Low management overhead is also an advantage as a single DHCP server is simpler to manage [58]. Fig. 77. shows how Client B and C on Subnet 2 can obtain a DHCP address lease from the DHCP Server 1 on Subnet 1.

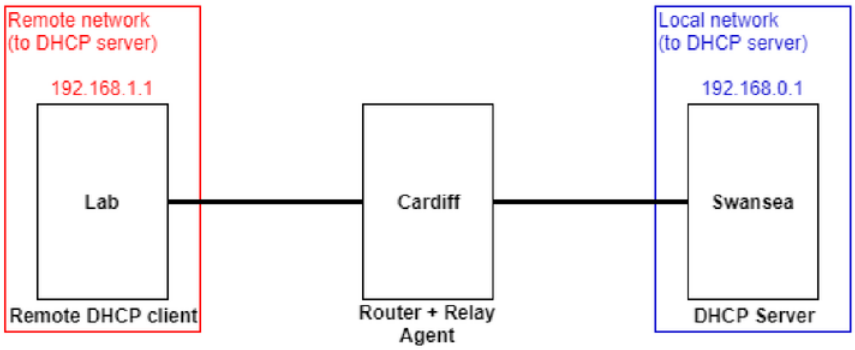


Fig. 78. An example topology of how a relay agent would work in the author's network (if the DHCP server was configured on Swansea)

A relay agent can be configured in Windows Server 2012 R2 under the Routing and Remote Access Microsoft Management Console (MMC) under IPv4 [56]. A relay agent is not needed in the author's presented network. This is because the DHCP service is running on Cardiff which is acting as the router and is directly connected to both the 192.168.0.0/24 and 192.168.1.0/24 networks. This means that the DHCP service will send and receive DHCP messages through the corresponding local interfaces to Swansea and the Lab. Note it is also not possible to configure a DHCP relay agent on a machine that is running the DHCP service (which Cardiff is) [56]. The author's network could be modified by making Swansea the DHCP server and configuring a DHCP relay agent on the router Cardiff. This would forward DHCP messages to and from the remote network (Lab), see Fig. 78.

Section 4: Name Resolution using the Domain Name System (DNS)

4.1. Understanding DNS operation; Planning and deploying DNS

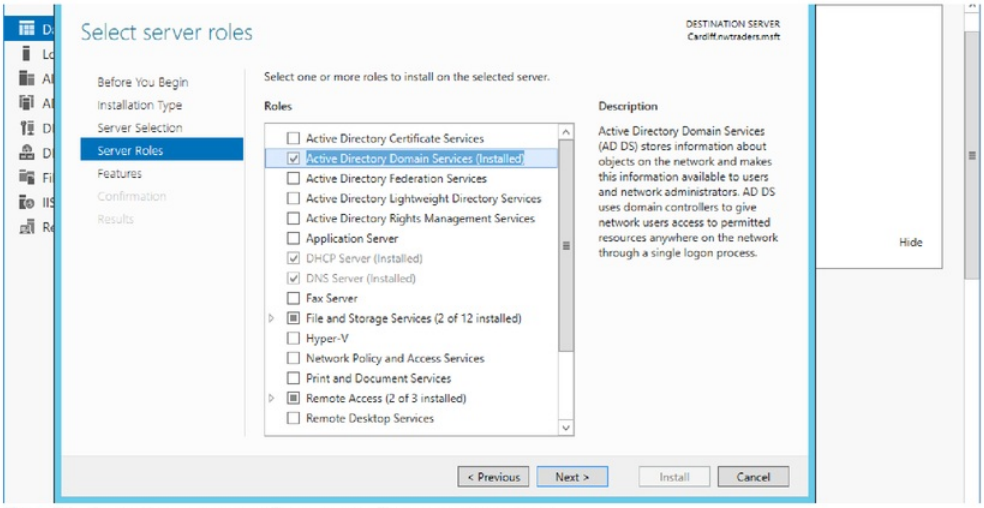


Fig. 79. Installing Active Directory Domain Services

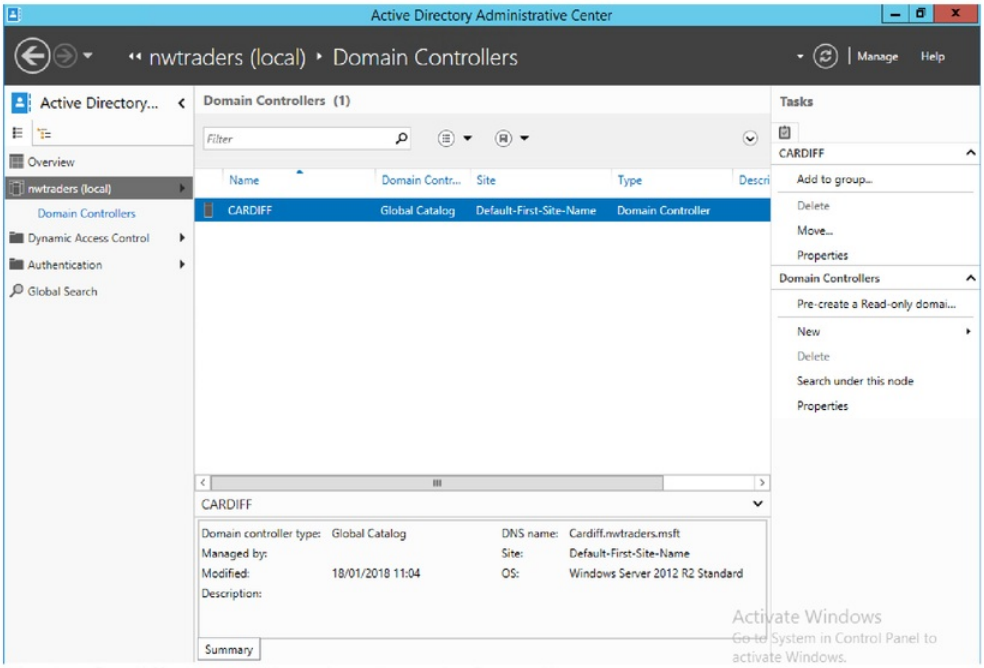


Fig. 80. Cardiff assumes the role as Domain Controller

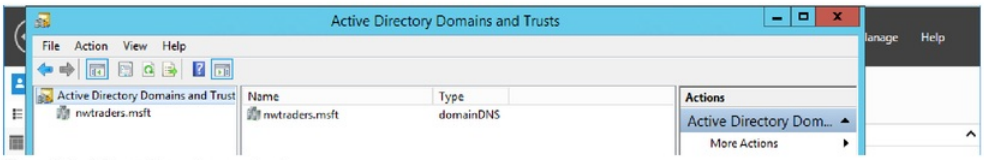


Fig. 81. New forest created

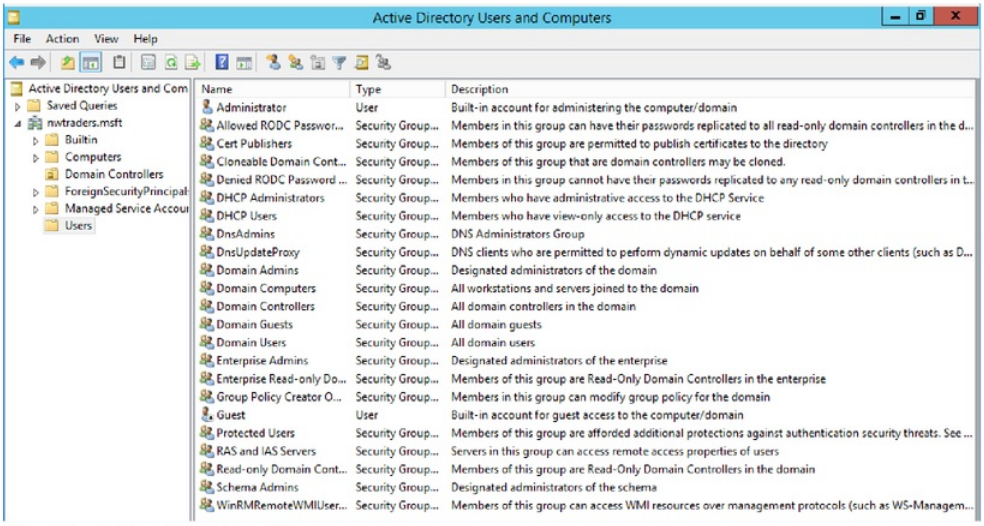


Fig. 82. Active Directory Users

Fig. 79-82. shows Active Directory (AD) successfully set up with nwtraders.msft as the Fully Qualified Domain Name (FQDN) for the forest (a single instance of AD at the highest level which groups domain trees) [60]. The process of configuring AD is conducted through a wizard after the Cardiff server is promoted to a domain controller via a notification in the Server Manager. After a restart, the administrative account should be used to log in (to allow configurations to be made).
Active Directory (AD) is a centralised Windows OS directory service that is used to automate and streamline network management of user data, security (e.g. authentication and authorisation) and distributed resources (see section 2.6.4.) [61], [62].

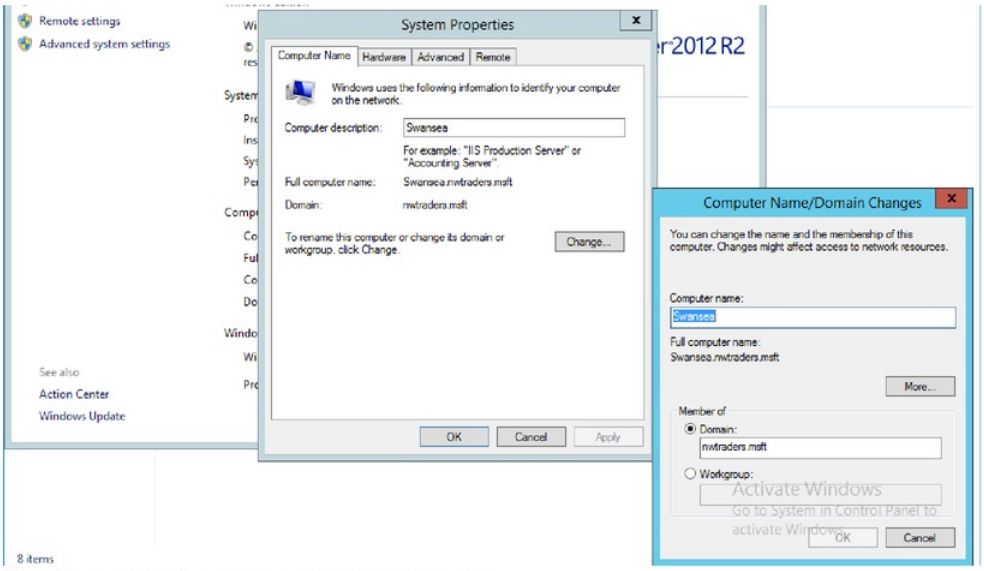


Fig. 83. Adding Swansea server to the domain

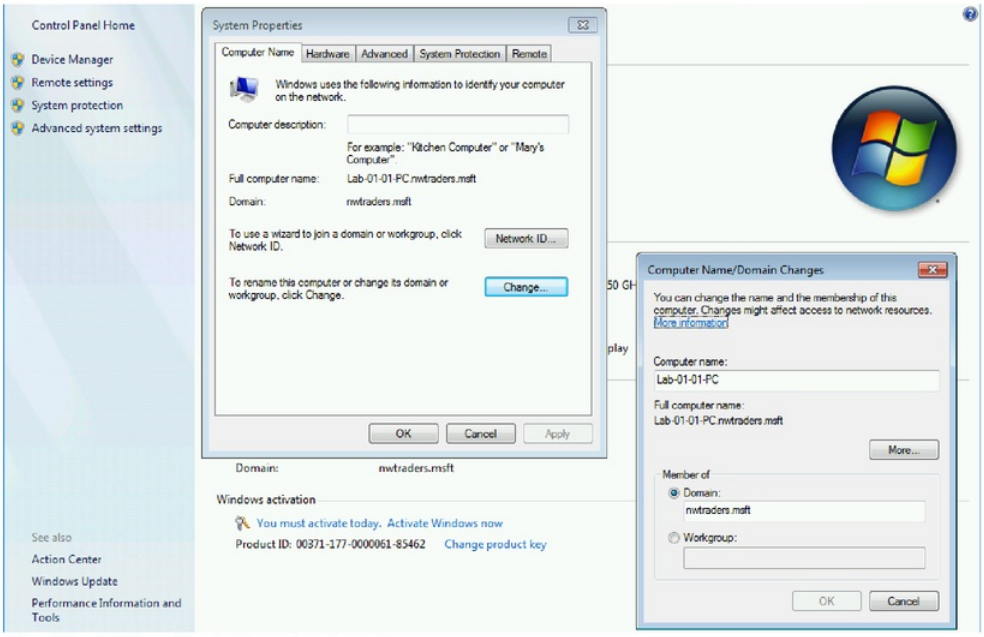


Fig. 84. Adding the Lab to the domain

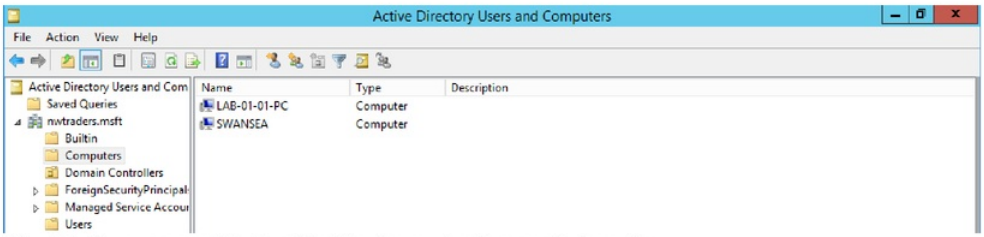


Fig. 85. Swansea and Lab added to the nwtraders.msft domain

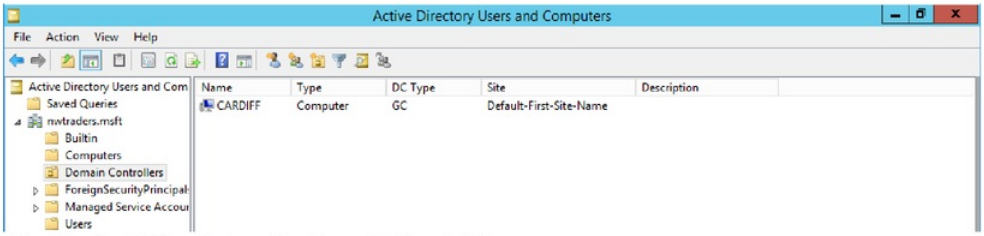


Fig. 86. Cardiff acting as the Domain Controller

Fig. 83-84. shows the process of adding Swansea and the Lab to the nwtraders,msft domain. Fig. 85-86. shows that Swansea and the Lab are successfully part of the nwtraders.msft forest and now under the jurisdiction of the AD service running on Cardiff (Fig. 86).

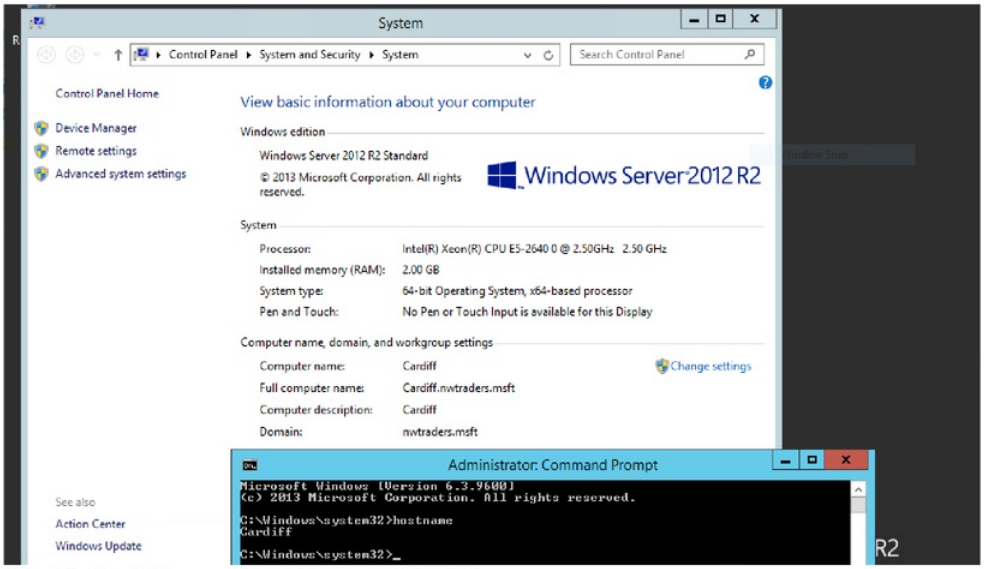


Fig. 87. Cardiff server Host Name and FQDN

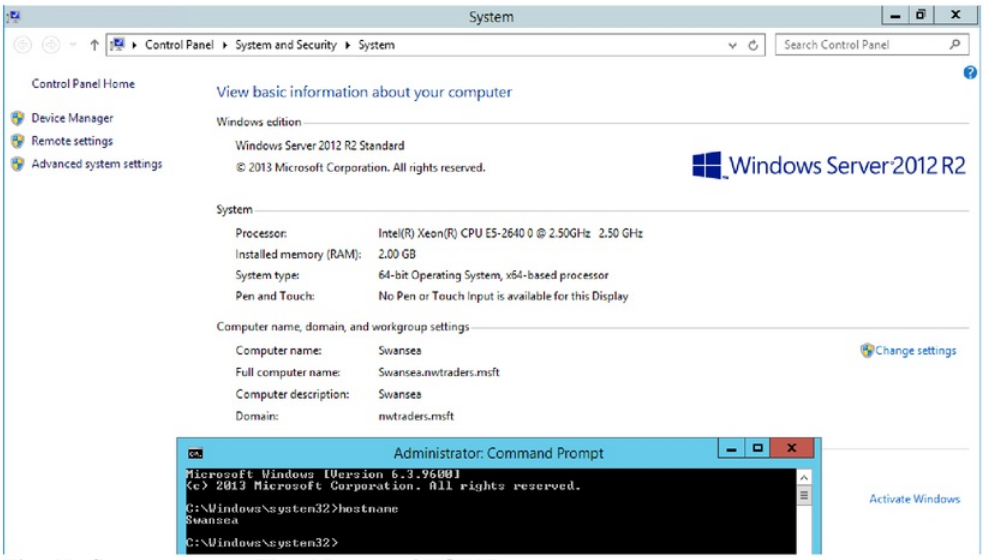


Fig. 88. Swansea server Host Name and FQDN

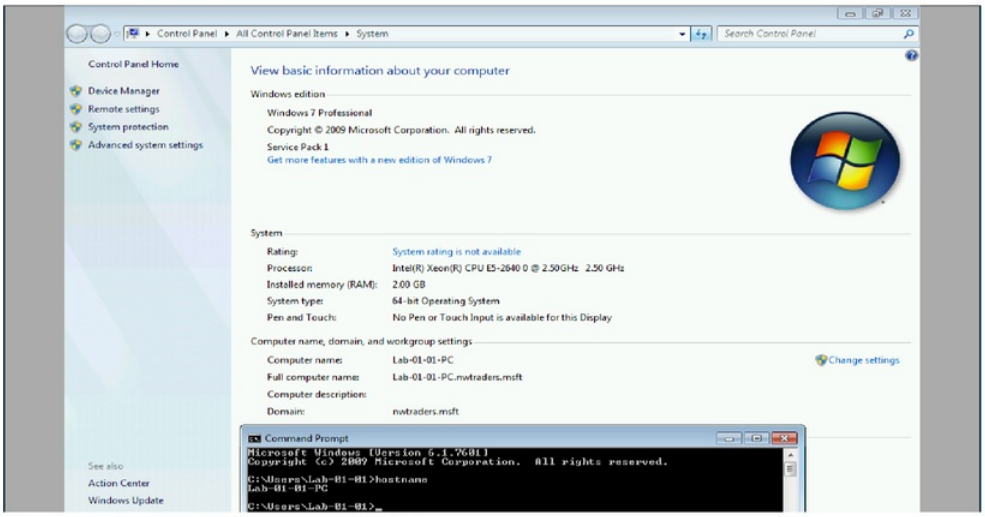


Fig. 89. Lab client Host Name and FQDN

Fig. 87-89. shows confirmation of Host Names and FQDNs.

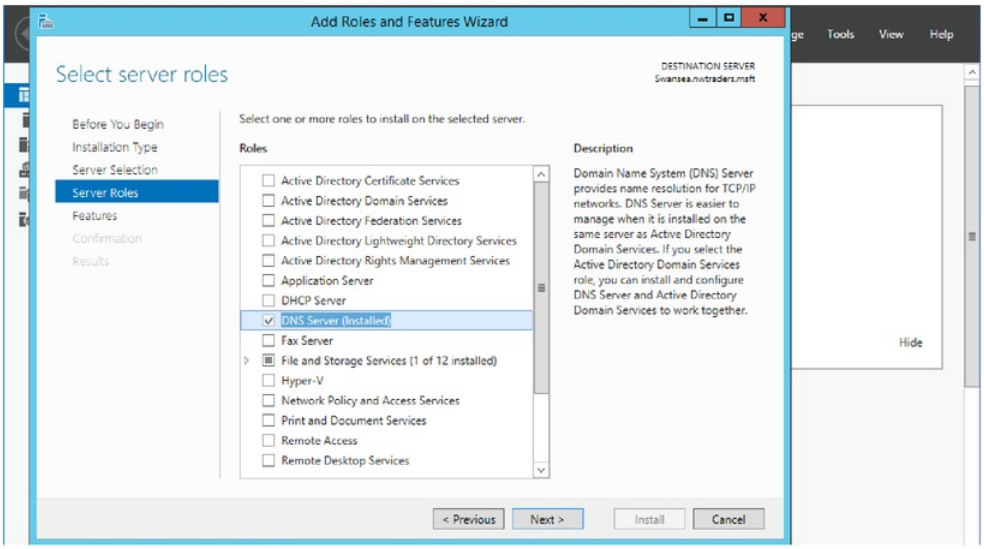


Fig. 90. Installing the DNS server on the Swansea server

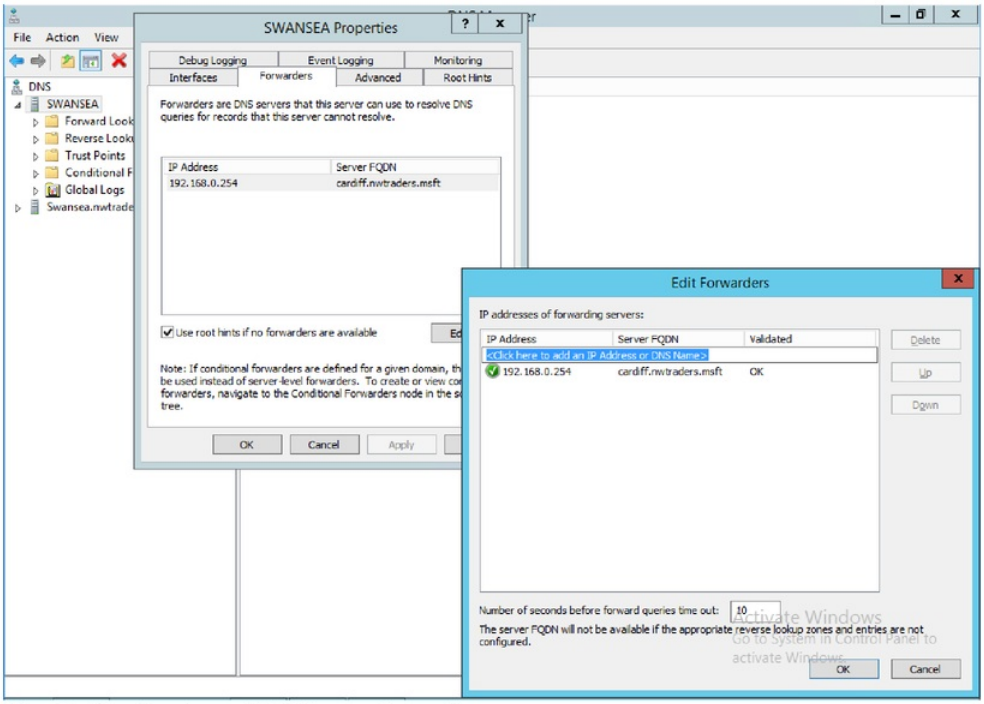


Fig. 91. Configuring a DNS Forwarder on Swansea

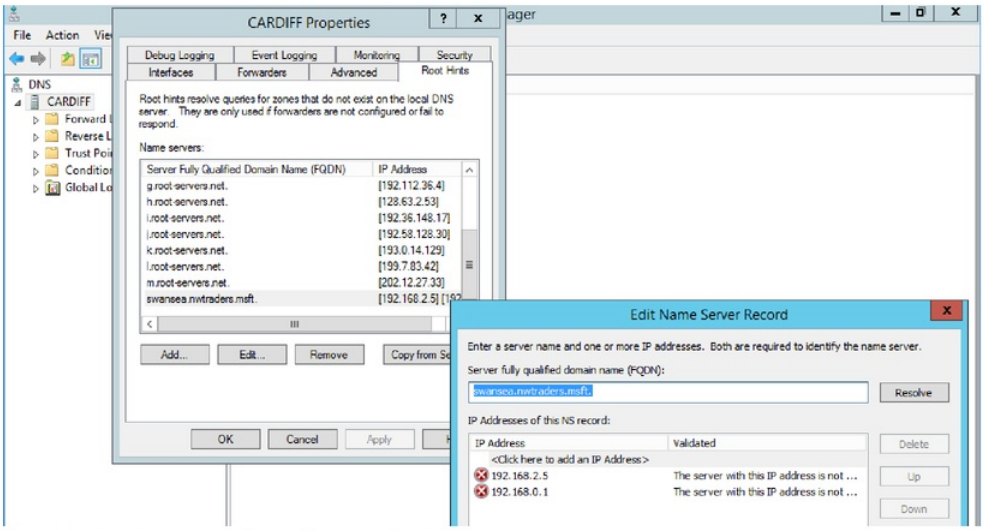


Fig. 92. Configuring a Root Hint on Cardiff

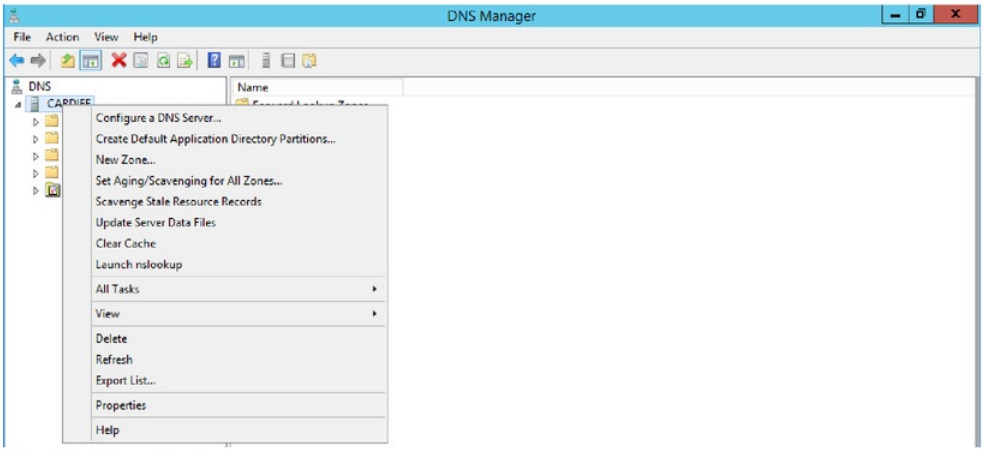


Fig. 93. Clear Cache

Fig. 90. shows the installation of the DNS server on the Swansea server (the DNS Server was installed on Cardiff when the AD Domain Services were installed). A forwarder is configured on Swansea pointing to the Cardiff server as seen in Fig. 91. A forwarder is used to resolve DNS queries that the server cannot resolve (e.g. the server is non-authoritative for the domain requested) [63]. A root hint is configured on Cardiff pointing to Swansea server as seen in Fig. 92 (it fails to validate in this example as Swansea is non-authoritative for the Cardiff zone). A root hint is used to resolve DNS queries for zones that do not exist on the local DNS server [63]. Fig. 93. shows the option to clear the DNS server cache, it is important to be able to control and clear the cache when troubleshooting DNS and connectivity issues (e.g. DNS entries in the cache may become out of date) [64], [65].

IMPORTANT: The DNS Server should be made a client of itself. This is done by setting the interfaces' IP address as the DNS server address in the IPv4 properties of the NICs. This is

important as the DC must look at itself for DNS, without this configuration AD will not function correctly [66]–[68].

4.2. Storing and Maintaining DNS Data



Fig. 94. Creating a new Secondary Forward Lookup Zone on Swansea

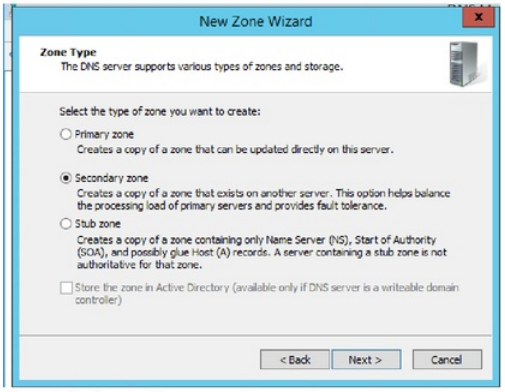


Fig. 95. Zone Type

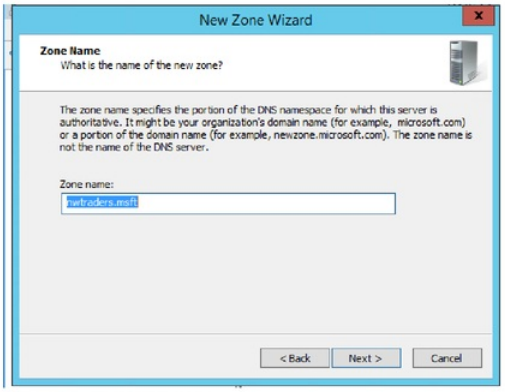


Fig. 96. Zone Name

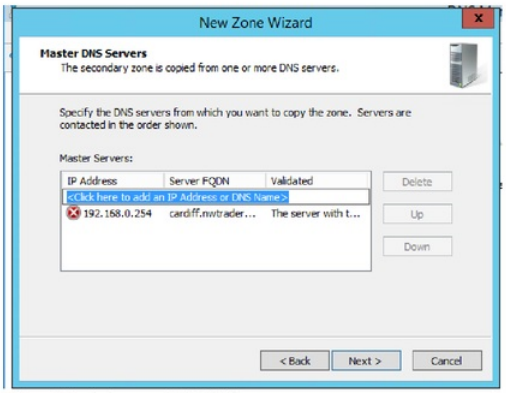


Fig. 97. Master DNS Servers



Fig. 98. Completing Zone creation

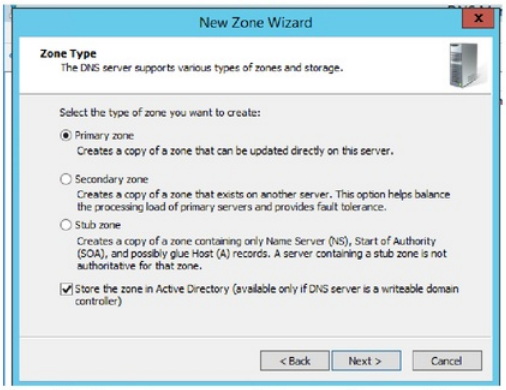


Fig. 99. Creating a new Primary Forward Lookup Zone on Cardiff

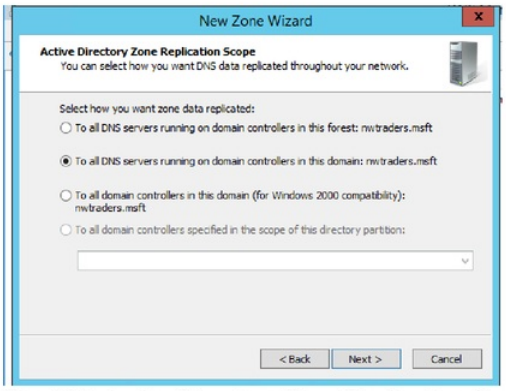


Fig. 100. Active Directory Zone Replication Scope

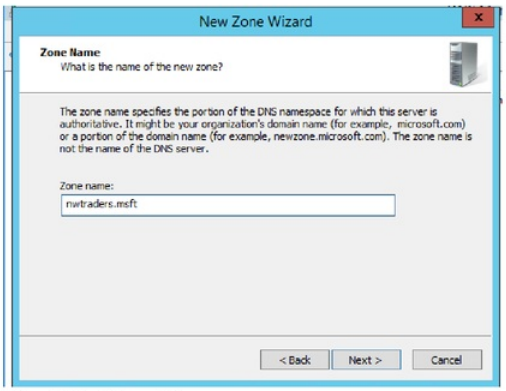


Fig. 101. Zone Name



Fig. 102. Dynamic Update

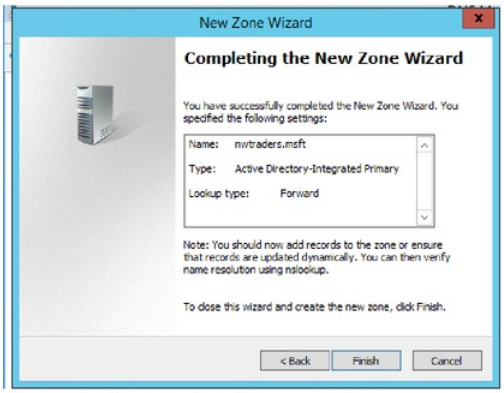


Fig. 103. Completing Zone creation

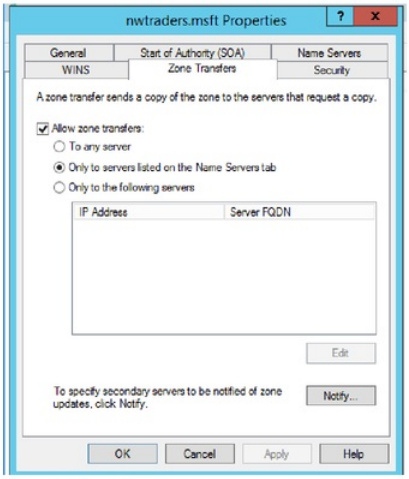


Fig. 104. Allowing zone transfer of nwtraders.msft from Cardiff to Swansea

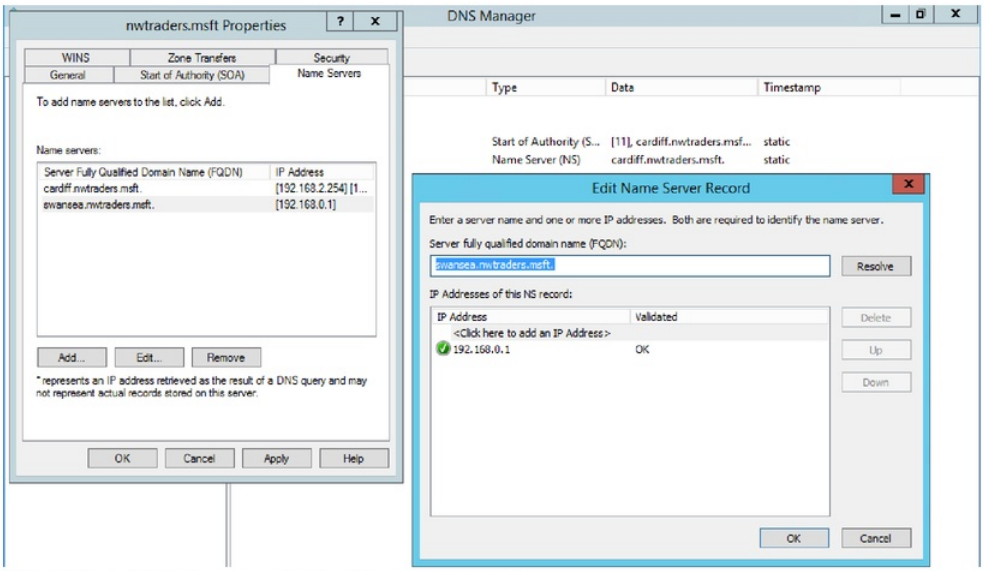


Fig. 105. Adding Swansea to the Name Servers list

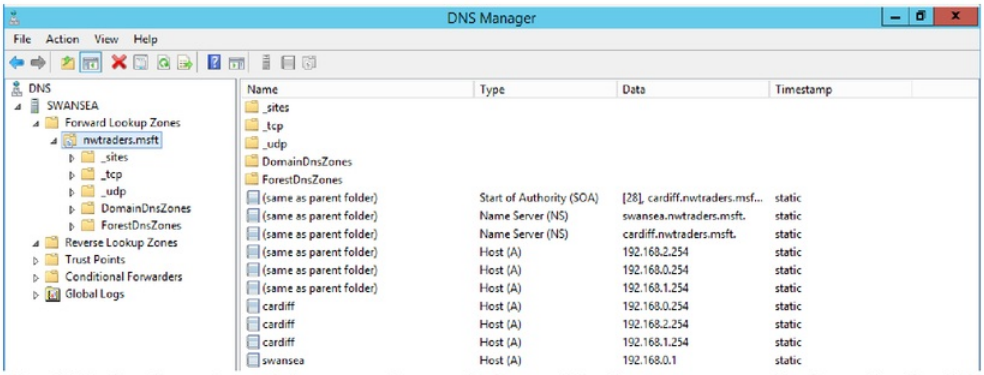


Fig. 106. Confirmation of the nwtraders.msft forward lookup zone transfer from the Cardiff Primary DNS server to Swansea Secondary DNS server

Fig. 94-106. shows a step-by-step guide on how to set up a primary and secondary forward lookup zone (with transfer) for the nwtraders.msft namespace. The forward lookup zone supports the query process that resolves host names to IP addresses [69]. See 4.4.5. DNS Zones for information on Primary and Secondary zones.

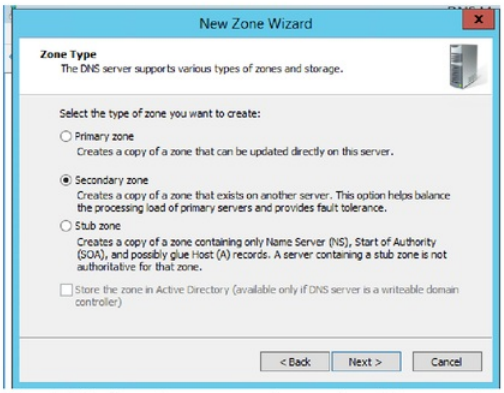


Fig. 107. Creating a new Secondary Reverse Lookup Zone on Swansea

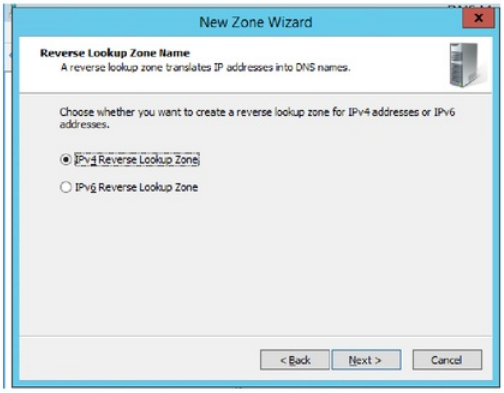


Fig. 108. IPv4 Reverse Lookup Zone

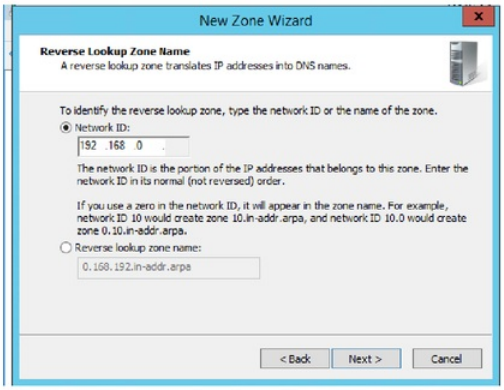


Fig. 109. Reverse Lookup Zone Network ID

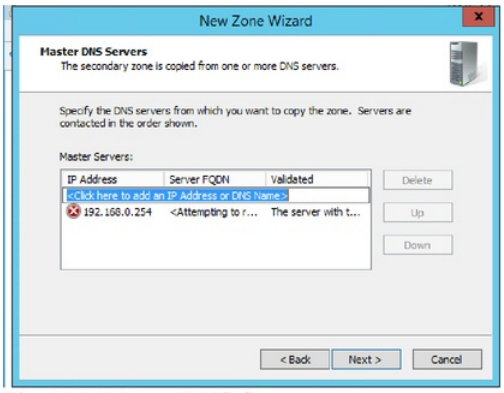


Fig. 110. Master DNS Servers

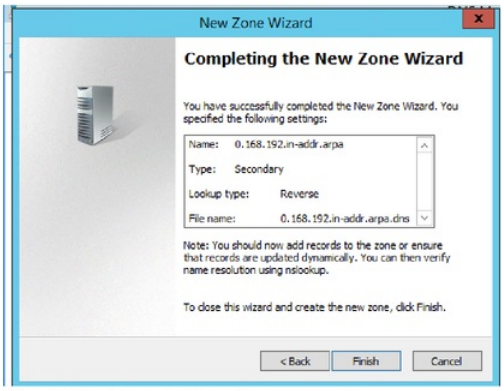


Fig. 111. Completing Zone creation

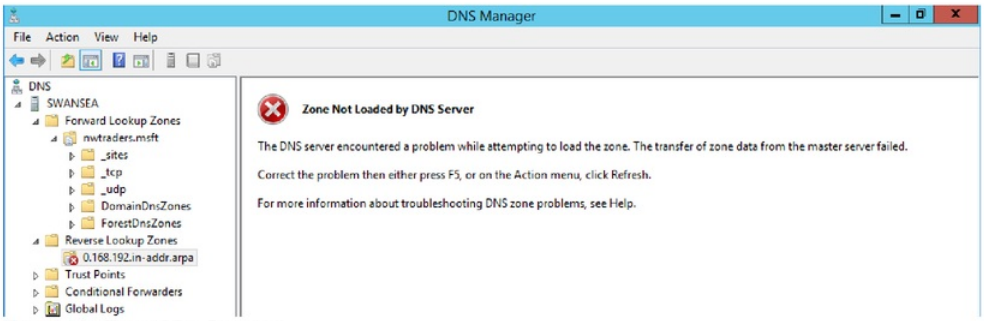


Fig. 112. Zone Not Loaded

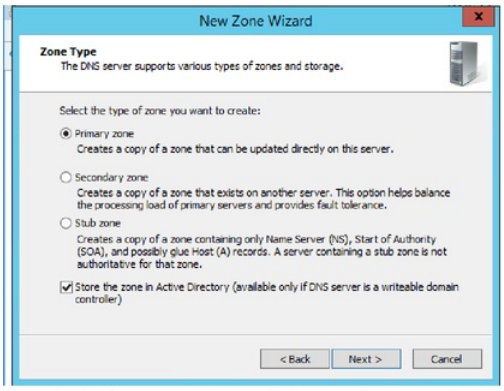


Fig. 113. Creating a new Primary Reverse Lookup Zone on Cardiff

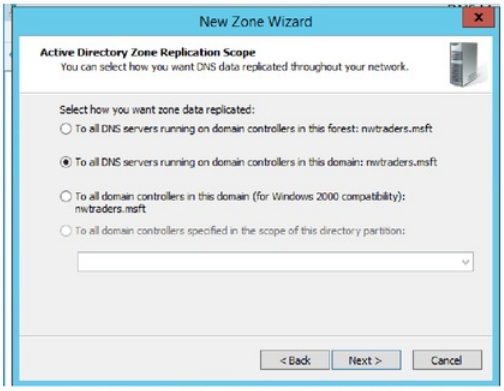


Fig. 114. Active Directory Zone Replication Scope

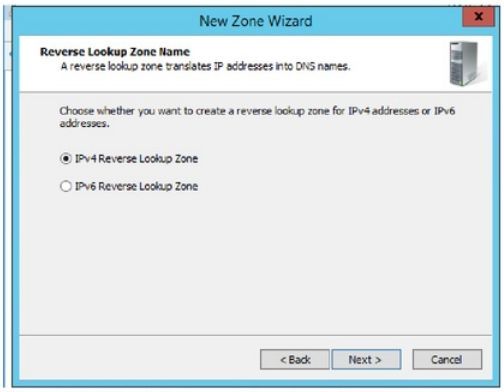


Fig. 115. IPv4 Reverse Lookup Zone



Fig. 116. Reverse Lookup Zone Network ID

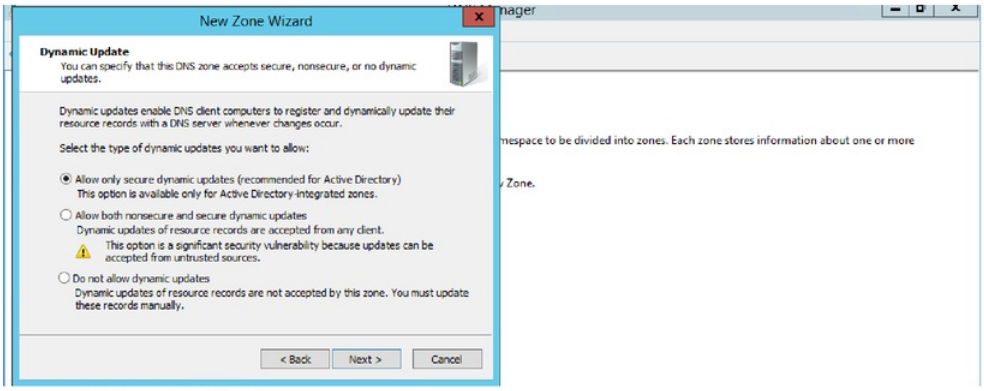


Fig. 117. Dynamic Update

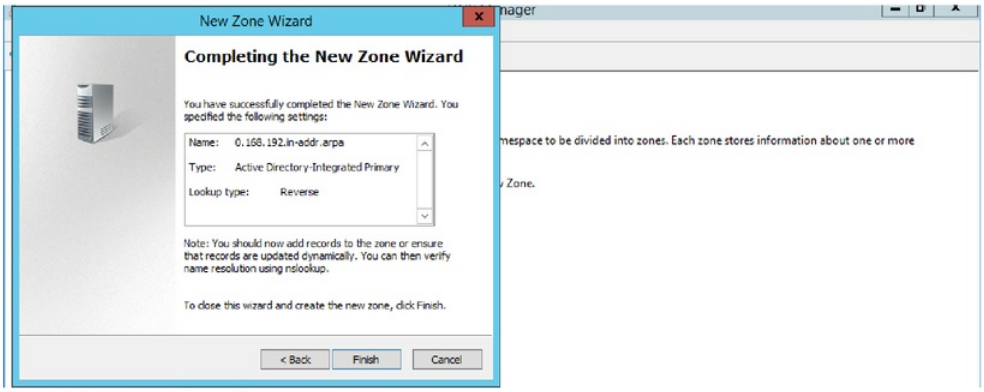


Fig. 118. Completing Zone creation

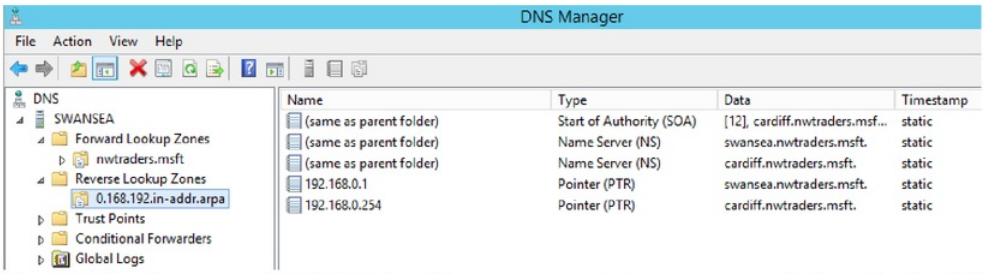


Fig. 119. Confirmation of 0.168.192.in-addr.arpa reverse lookup zone transfer from the Cardiff Primary DNS server to the Swansea Secondary DNS server

Fig. 107-119. shows a step-by-step guide on how to set up a primary and secondary reverse lookup zone for the 192.168.0.0/24 network. A reverse lookup zone contains pointer resource records that map IP addresses to FQDNs [70].

Note: Zone transfer will also need to be configured.

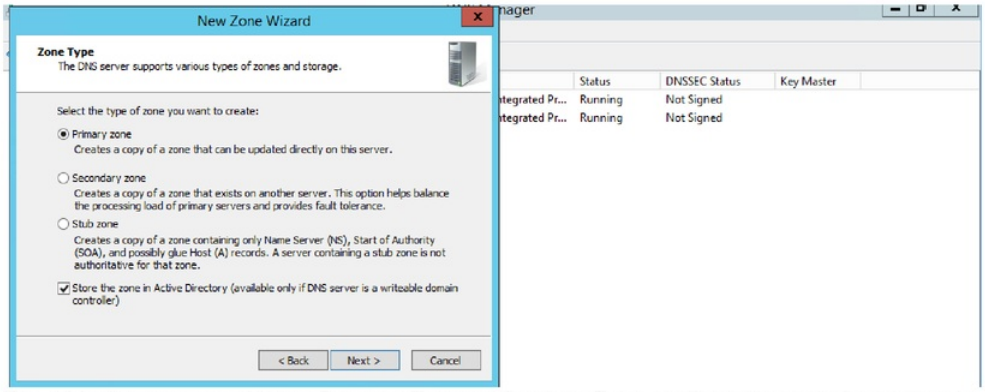


Fig. 120. Creating an addition Primary Forward Lookup Zone on Cardiff to only look after the swanseaserver.nwtraders.msft domain (an example domain, therefore, contains no hosts)

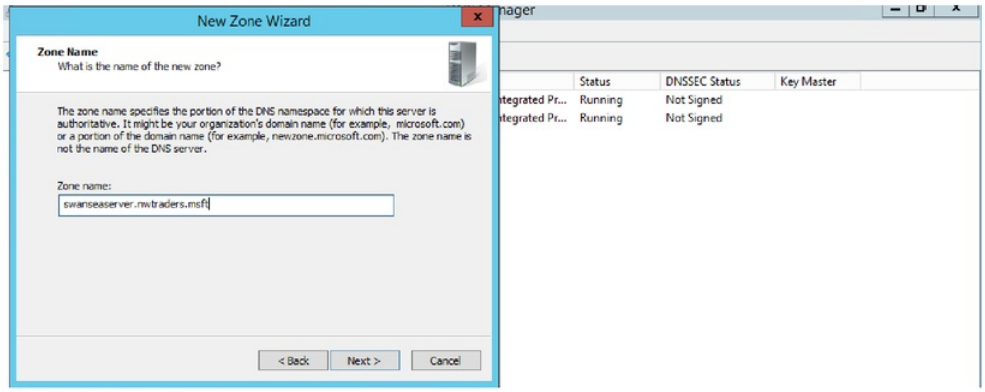


Fig. 121. Zone Name

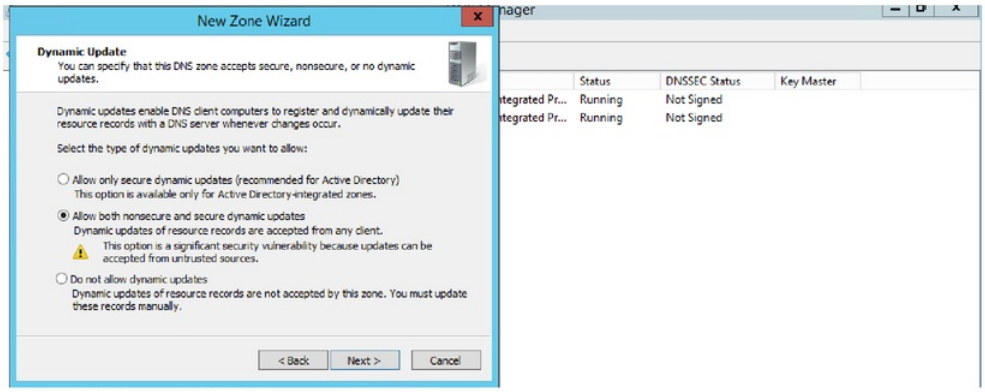


Fig. 122. Dynamic Updates

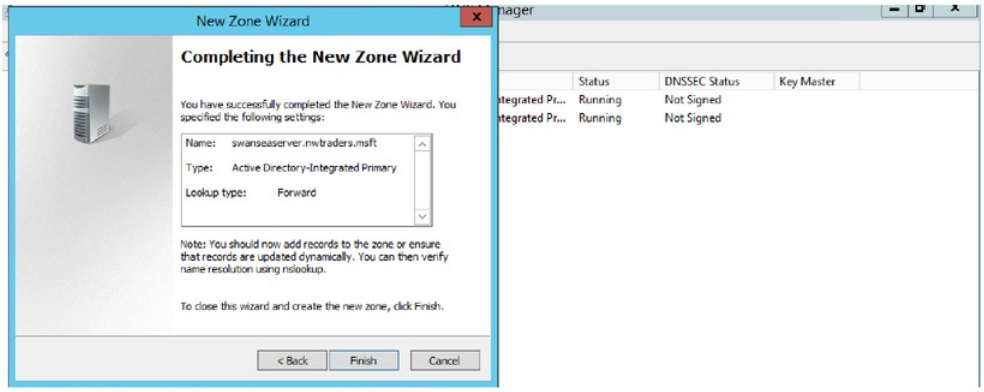


Fig. 123. Completing Zone creation

Fig. 120-123. shows a step-by-step guide on how to set up an additional primary lookup zone for the swanseaserver.nwtraders.msft namespace. The reason for creating additional zones is to add redundancy and load balancing to a DNS server (see 4.4.5. DNS Zones for further information). In this example, a zone has been created to exclusively look after the swanseaserver.nwtraders.msft namespace. This means that if the main nwtraders.msft zone fails for some reason, hosts in the swanseaserver.nwtraders.msft domain will still be able to resolve DNS queries within the domain.

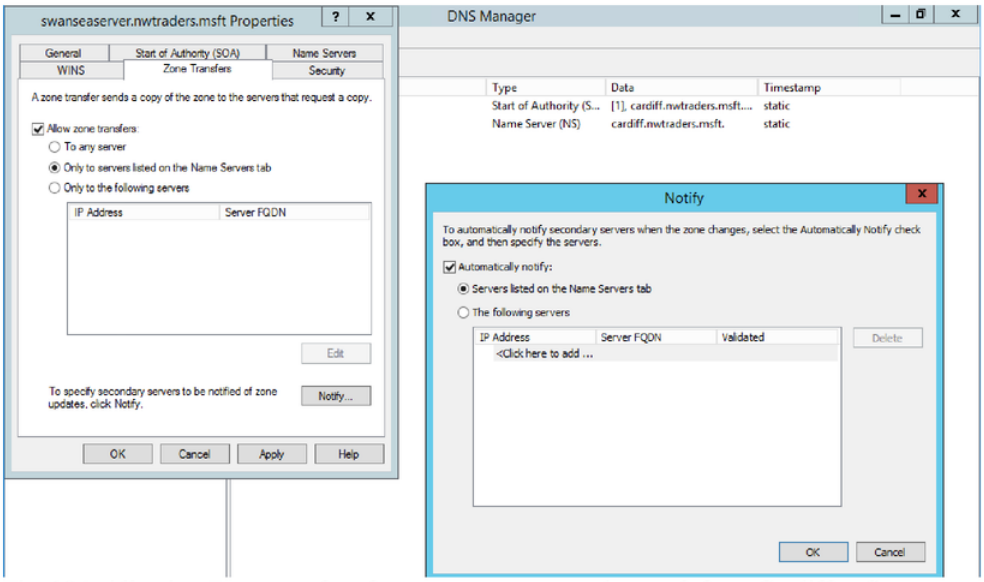


Fig. 124. Allowing Zone transfer of swanseaserver.nwtraders.msft from Cardiff to Swansea

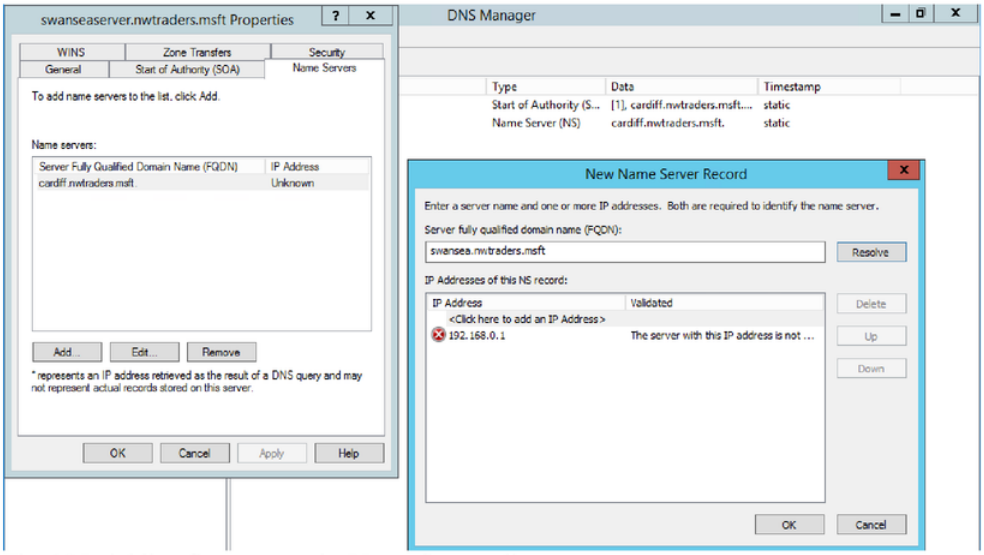


Fig. 125. Adding Swansea to the Name Servers list

Fig. 124-125. show the process of implementing fault tolerance via DNS zone transfer. This allows for all data in a zone to be copied from the primary zone to the secondary zone [71]. Any updates and changes to the primary zone will be copied to the secondary zone [72]. This means mean that the secondary zone is an up to date copy of the primary zone and can resolve queries if the primary zone was to fail [73]. DNS Notify is a process where the primary server notifies the secondary server about changes to a zone. This allows the secondary server to stay up to date and obtain the latest records [74].

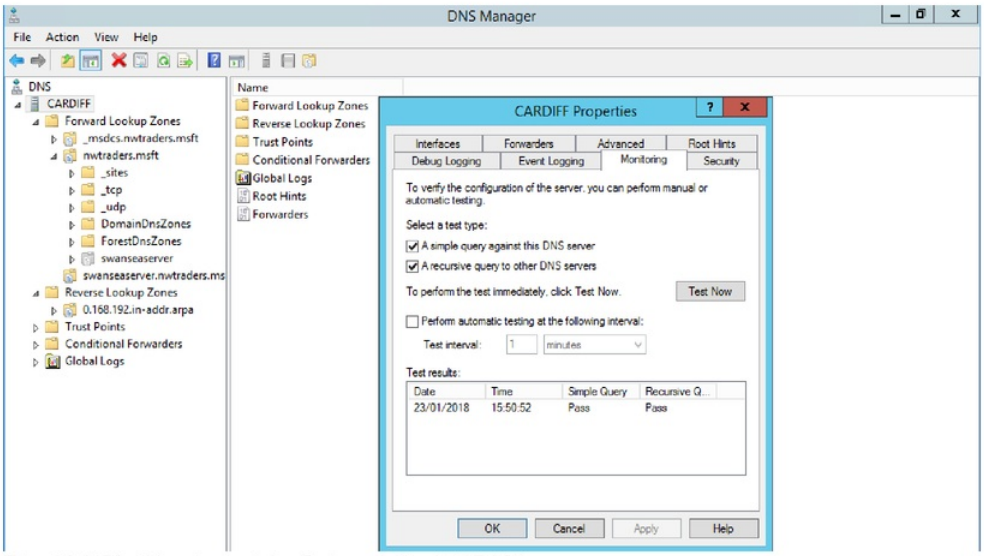


Fig. 126. Verification of the Primary Cardiff DNS server

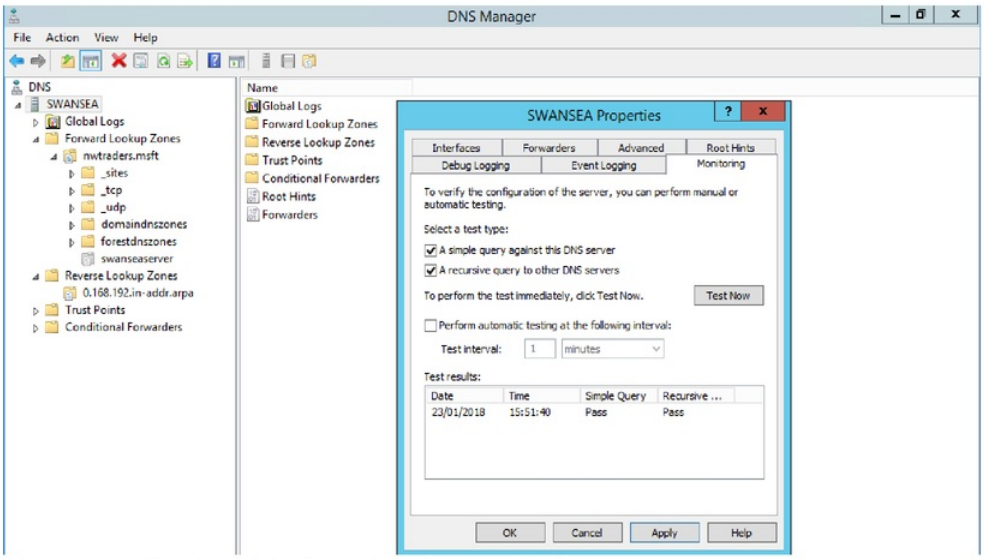


Fig. 127. Verification of the Secondary Swansea DNS server

Fig. 126-127. shows the success of simple and recursive queries on the primary Cardiff DNS server and secondary Swansea DNS server. This verifies the configurations of the DNS servers.

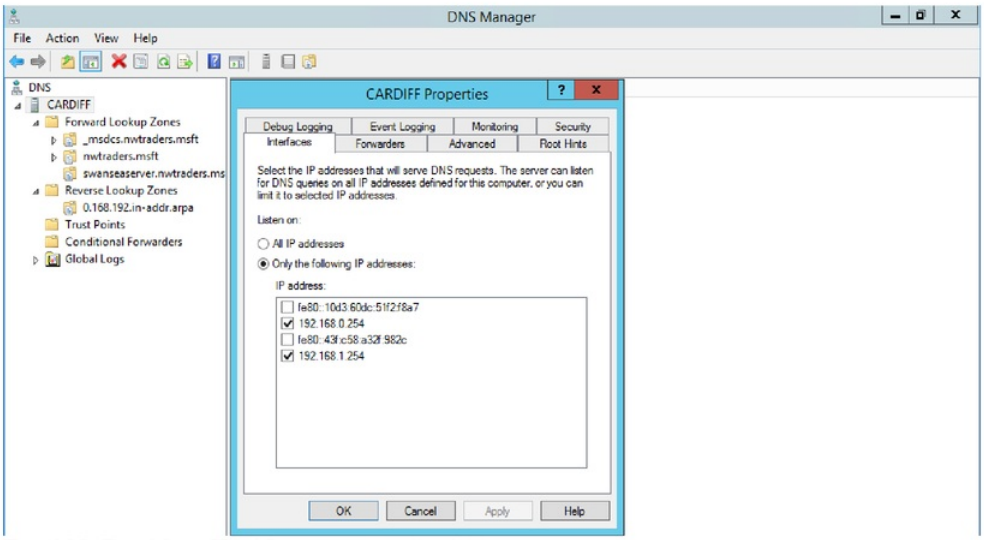


Fig. 128. Disabling IPv6 Interfaces on Cardiff

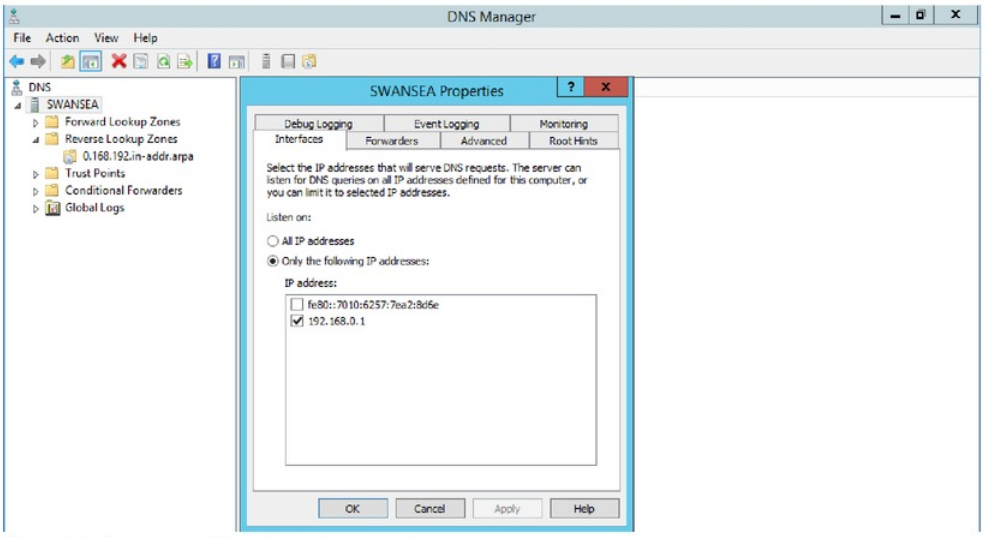


Fig. 129. Disabling IPv6 Interfaces on Swansea

Fig. 128-129. shows a fix to the unknown default DNS server error in nslookup by only listening for IPv4 DNS queries.

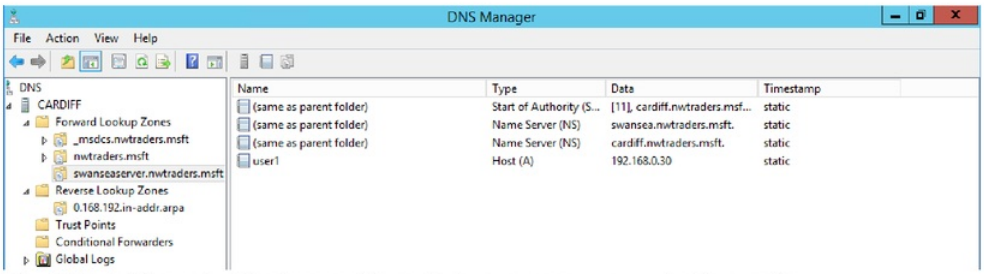


Fig. 130. Adding a test host record (user1) to swanseaserver.nwtraders.msft

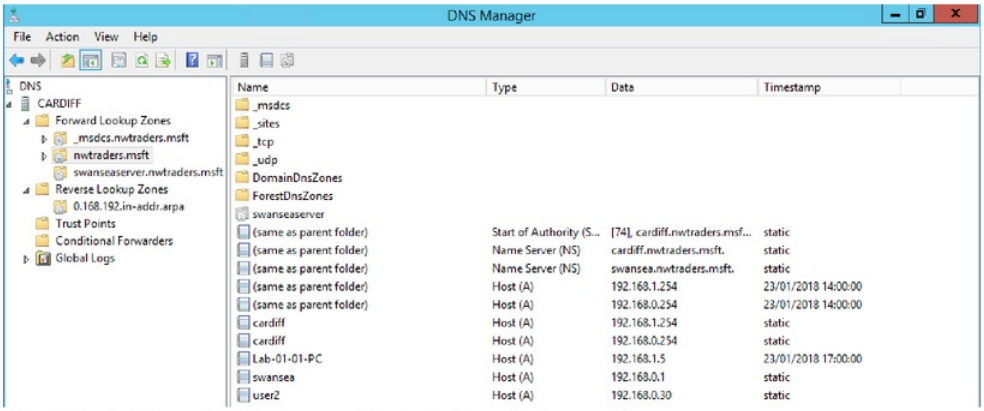


Fig. 131. Adding a test host record (user2) to nwtraders.msft

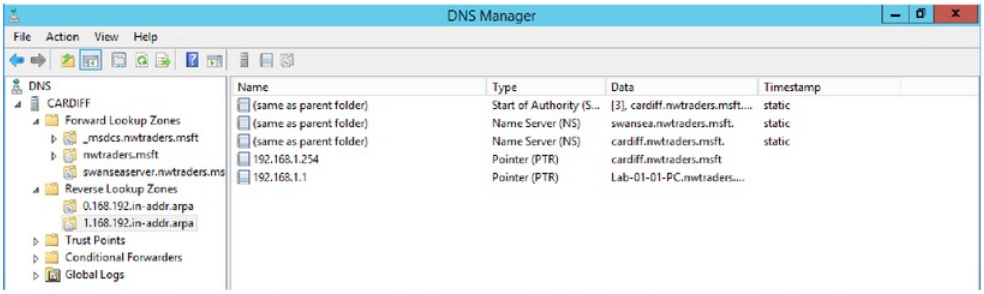


Fig. 132. Addition of another reverse lookup zone for the 192.168.1.0/24 network

Fig. 130-131. shows the addition of test host records that will be used for nslookup tests and DNS application test. Fig. 132. shows the reverse lookup zone for the 192.168.1.0/24 network. This is needed to resolve IP addresses to FQDNs in the 192.168.1.0/24 network.

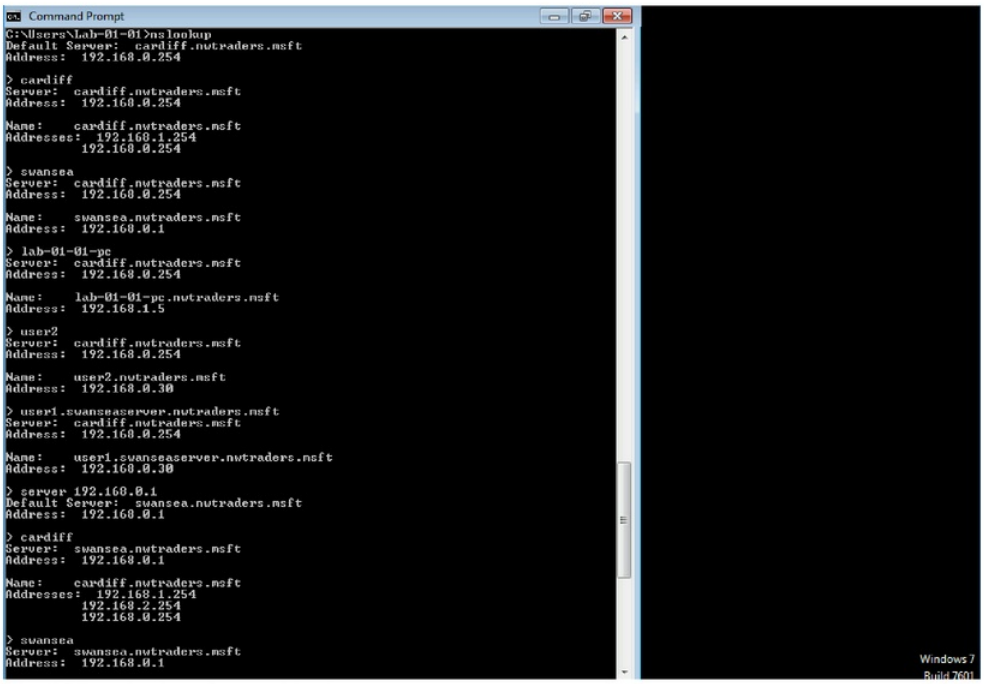


Fig. 133. nslookup forward lookup tests

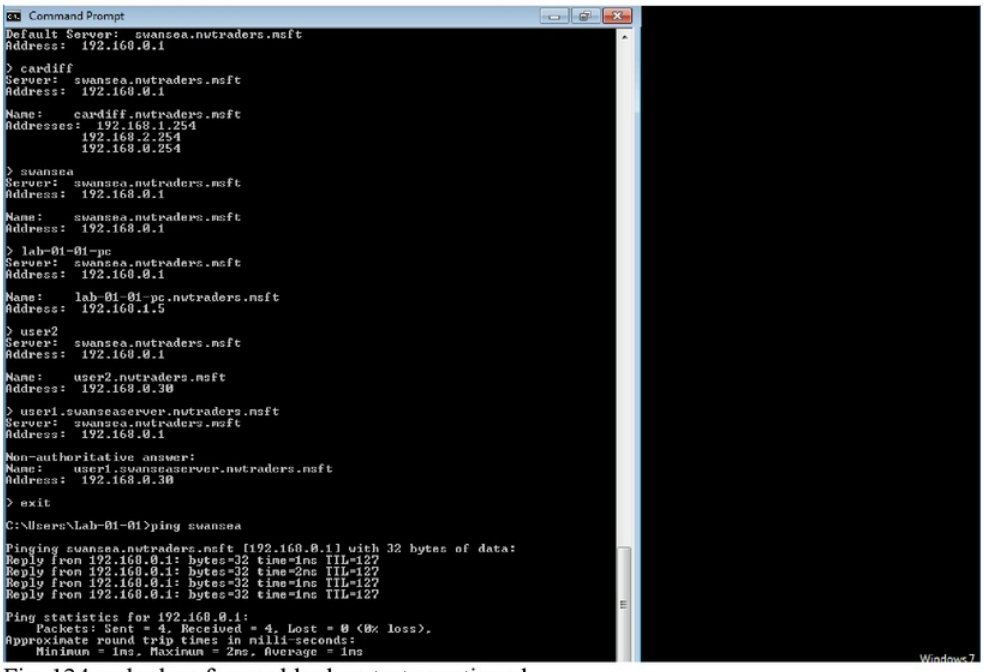


Fig. 134. nslookup forward lookup tests continued



Fig. 135. nslookup reverse lookup tests

Fig. 133-134. shows extensive nslookup tests on the Lab machine. The tests confirm the successful resolution of host names to IP addresses using both the Cardiff and Swansea DNS Server. Fig. 135. shows the successful resolution of IP addresses to FQDNs.

4.3. DNS Application

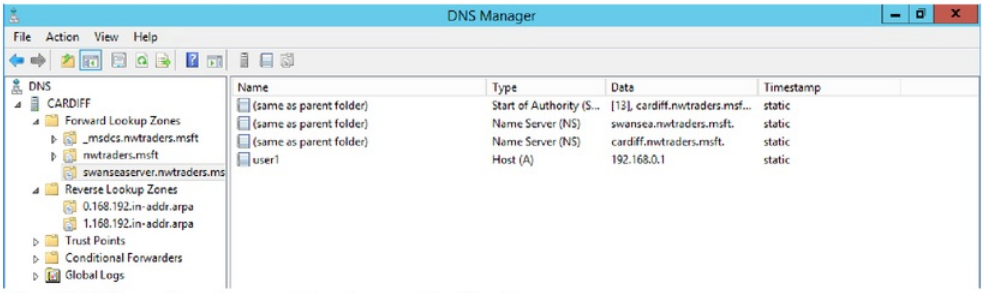


Fig. 136. Changing the user1 host record's IP address

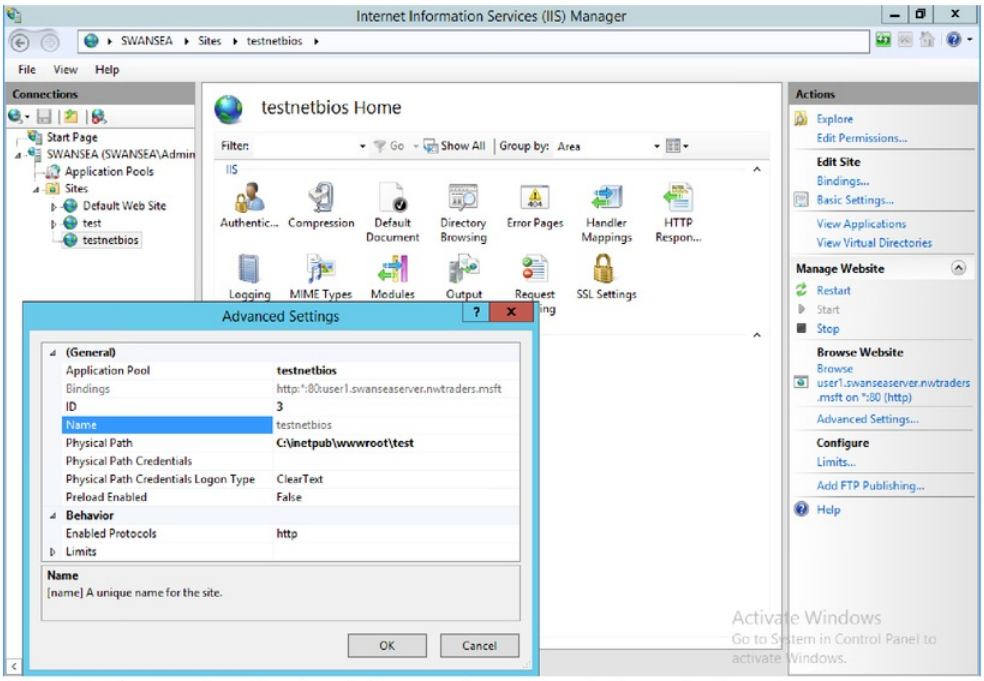


Fig. 137. The creation of a virtual directory for a test website on the Swansea server

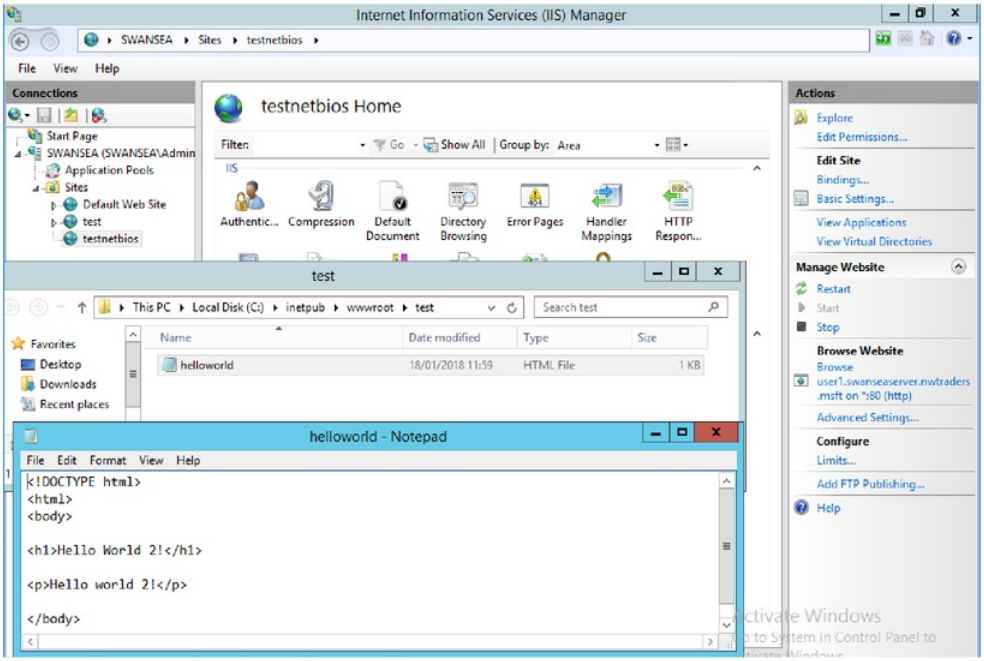


Fig. 138. Creation of a test Hypertext Markup Language (HTML) web page in the 'C:\Inetpub\wwwroot\test' directory specified by the virtual directory

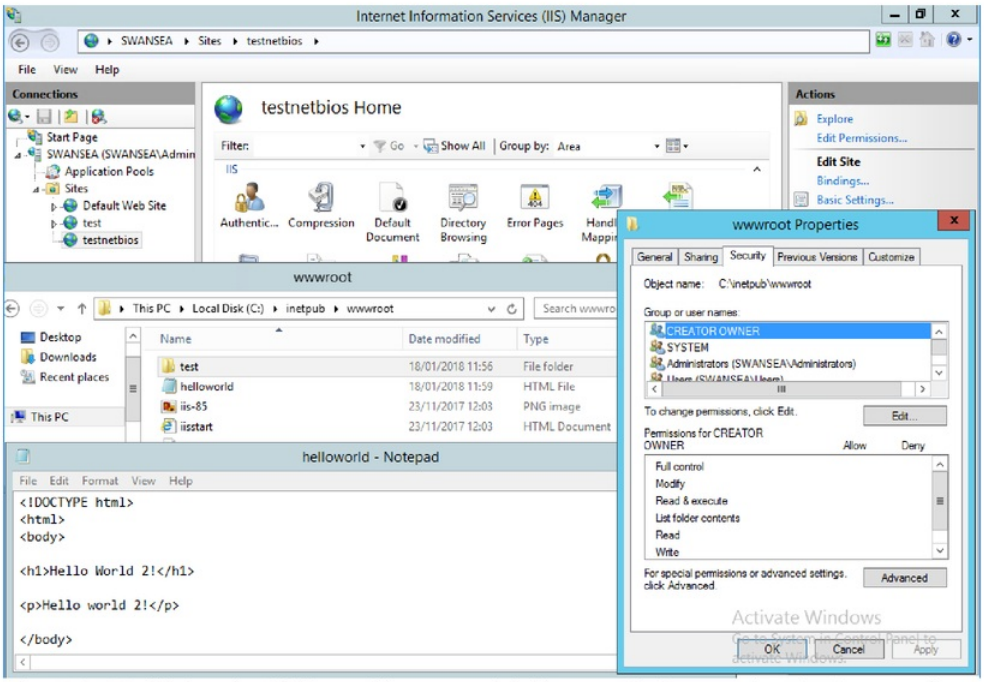


Fig. 139. Modifying the 'C:\Inetpub\wwwroot' folder permissions to give the Internet Guest Account modify permissions

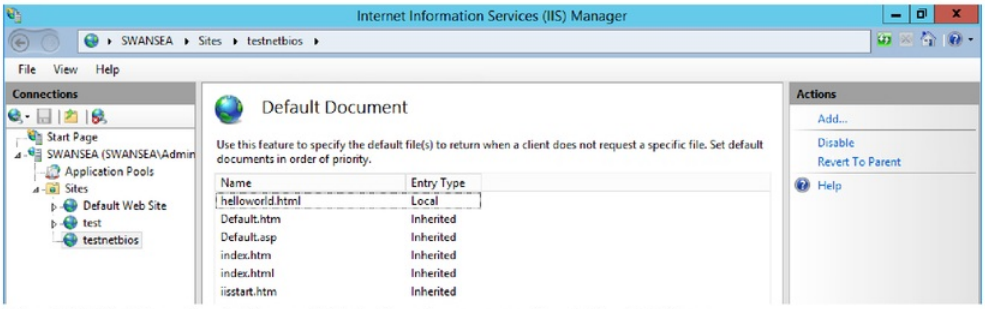


Fig. 140. Setting the helloworld.html web page as the default IIS page

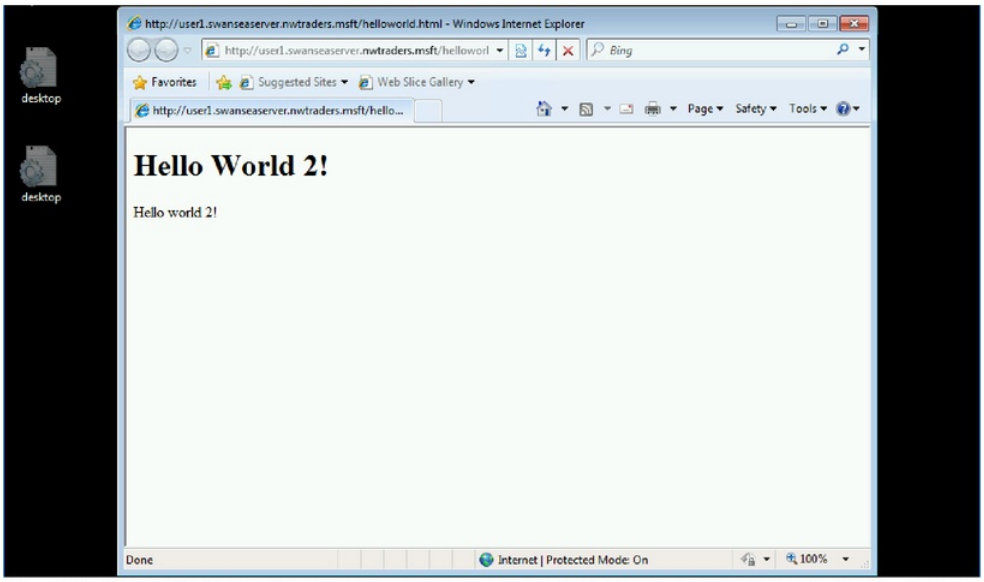


Fig. 141. The website on Swansea server can successfully be accessed from the Lab using the address 'http://user1.swanseaserver.nwtraders.msft/helloworld.html'.

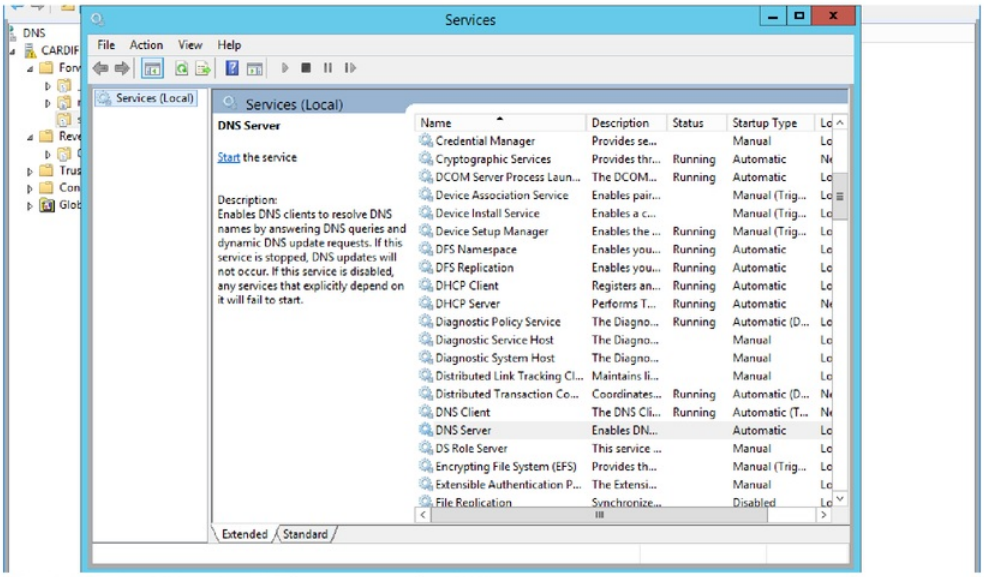


Fig. 142. Disabling the DNS server on Cardiff

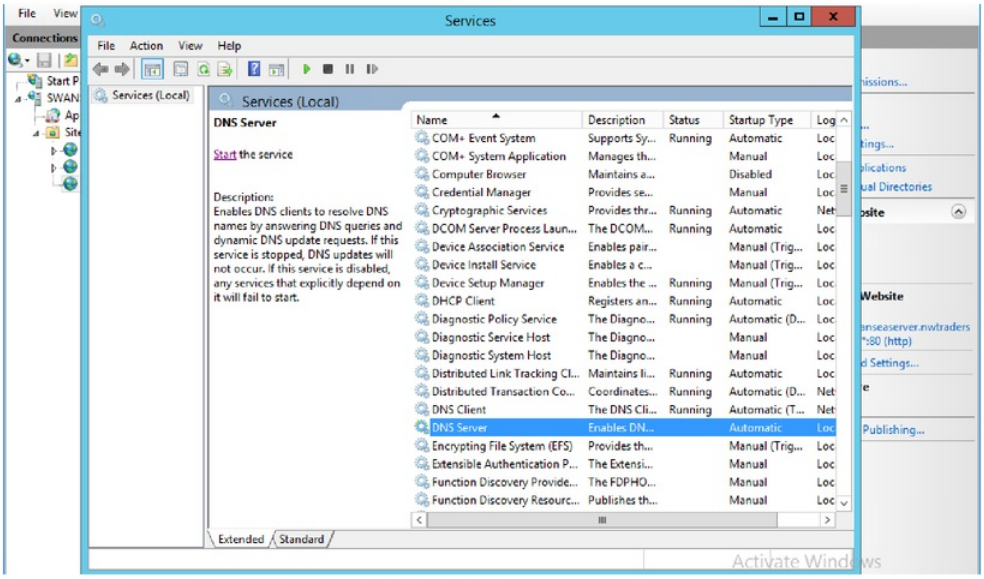


Fig. 143. Disabling the DNS server on Swansea

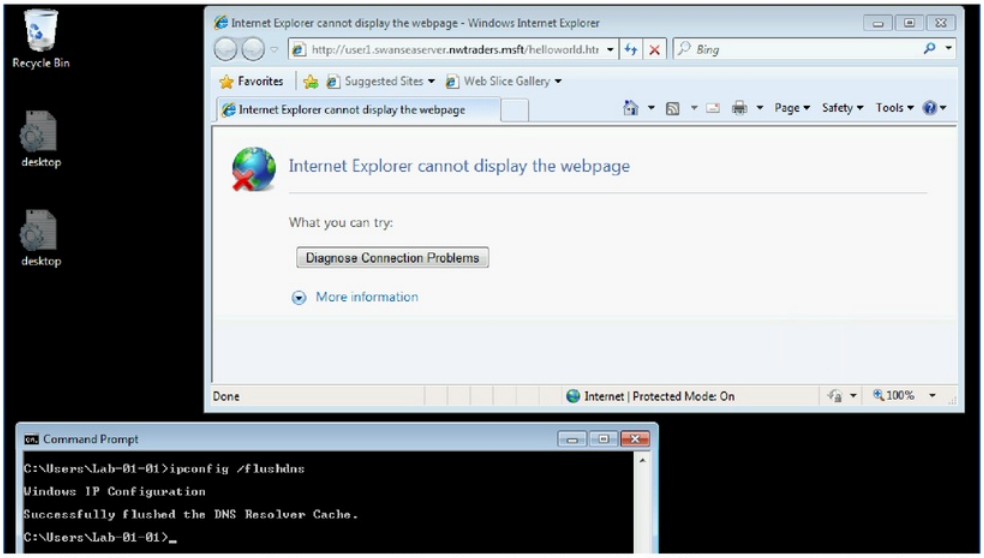


Fig. 144. The name resolution now fails after disabling the DNS servers, clearing the DNS cache and closing and reopening Internet Explorer

Fig. 136-144. demonstrates a practical test of the DNS server using the Web server service. Fig. 137. shows the directory where the website content will be stored. The site is also bound to the user1.swanseaserver.nwtraders.msft host file which points to the Swansea server's IP address (192.168.0.1). The reason for this is to bind the website to a test host record that has no NetBIOS services. If this was not done it would be hard to prove that the DNS server is working as the NetBIOS mechanism would just resolve the name instead of the DNS server. Fig. 141. highlights the DNS server successfully resolving user1.swanseaserver.nwtraders.msft to the IP address of Swansea (192.168.0.1) by displaying the web page.

4.4. DNS Observational Tasks

4.4.1. Resolving a Name to IP address

Host name to IP address resolution is a necessity on a network using the TCP/IP protocol to connect to resources [75].
Microsoft Windows clients will follow a number of steps when attempting to resolve a name to an IP address [75]:

1. The client will check to see if the name queried is its own.
2. The client searches a local Hosts file (located at 'C:\Windows\System32\drivers\etc' on Windows 10/8/7). This file contains a list of mappings of IP addresses to host names on the local computer [76].
3. The client queries Domain Name System (DNS) servers.
4. Lastly, the NetBIOS name resolution sequence is used as a last resort.

4.4.2. Recursive Query and Iterative Query processes

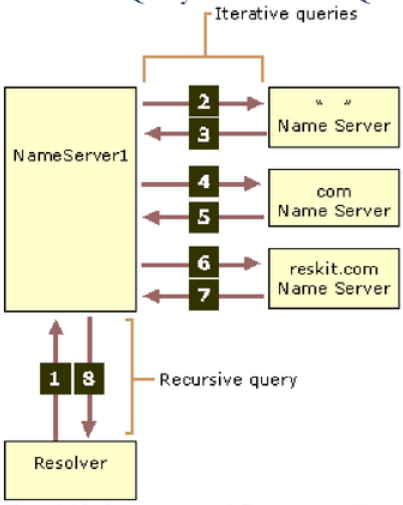


Fig. 145. Iterative and Recursive Queries [77]

A recursive name query is normally made from a DNS client to a DNS server. The DNS client will send a recursive name query to a DNS server requesting a resource record. The DNS server will either respond with the requested record or an error message stating the record or domain name does not exist. If the DNS server does not have the requested information, it will query other servers to get the information. If the information cannot be found the name query fails. Recursive name queries can also be made by DNS servers configured to use a forwarder to pass unresolved name queries to another DNS server. [77]

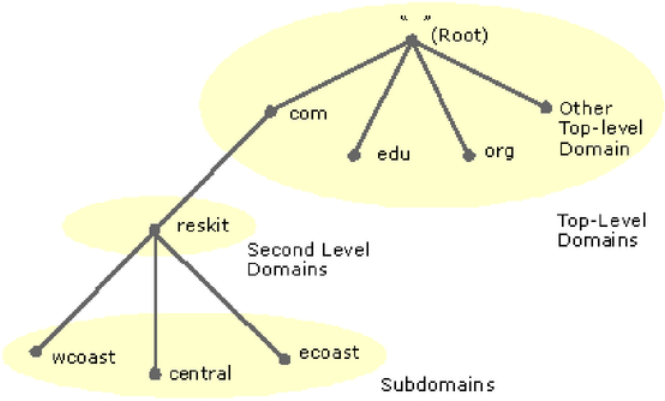


Fig. 146. Domain Name System hierarchy [78]

An iterative name query is when a DNS client allows a DNS server to return the best response it can, based on its cache or zone data. When a queried DNS server does not have a match for the requested name it can make a referral to another DNS server authoritative for a lower level of the domain namespace [77]. For example, in Fig. 145. the client contacts NameServer1 with a recursive query for noam.reskit.com. NameServer1 does not know the answer based on its cache and zones so contacts the root name server of the Internet [77]. The root name server is part of the DNS hierarchical namespace (Fig. 146) and publishes a root zone file to other DNS servers and clients on the Internet. The root zone file contains information about where authoritative servers for the DNS top-level domains are located e.g. .com, .edu, .org [79]. The root name server does not know the answer so refers NameServer1 to a server authoritative for the .com domain. The referrals continue down through the various domain levels until a server knows the answer to the iterative query. In the example, the server authoritative for the secondary domain reskit.com is queried for noam.reskit.com. It knows the answer and responds to NameServer1 with the requested IP address. NameServer1 can now respond to the client with the IP address for noam.reskit.com [77], [78].

4.4.3. Root Hints

Root hints resolve queries for zones that do not exist on the local DNS server [63]. This process uses iterative name queries (see section 4.4.2.) [63]. Root hints use a root hints file which contains host information needed to resolve names outside of the authoritative DNS domains. It contains the names and addresses of either the root DNS servers on the Internet or the DNS servers that are authoritative for the root of the private network [80].

4.4.4. Question

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:

DNS uses a hierarchical namespace, using a Parent and Child Domain structure. This allows for scalability and reuse of names (e.g. the host name *Server1* can be reused as long as it is unique within the Domain). *Server1* identifies the host name, *subcontA* identifies the Child Domain (subdomain) that *Server1* lives within, and *nwtraders* identifies the Parent Domain that *subcontA* is part of. *Nwtraders.msft* is the root domain (highest hierarchical level of a site) with a dot separating the second-level and top-level domains. [81]–[83].

4.4.5. DNS Zones

A DNS Zone is a portion of DNS namespace that has been delegated to other servers [81], [84]. DNS Zones can have various purposes, this includes improved manageability and data distribution of a Domain namespace with many subdomains. The same zone could also be stored on multiple servers to provide load balancing and DNS fault tolerance [81], [85].

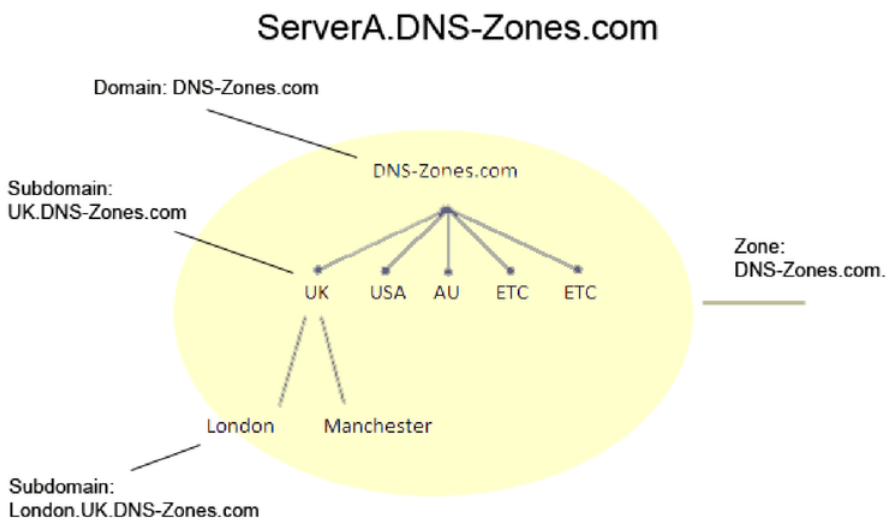


Fig. 147. An example of a single zone responsible for all subdomains in a large enterprise [84]

Fig. 147. shows an example where the domain and zone DNS-Zones.com is hosted on a DNS server called ServerA.DNS-Zones.com. There are likely to be many more subdomains for multiple cities for each country. In this example, as there is only one zone ServerA.DNS-Zones.com would have to store every DNS record of this worldwide organisation. This would be impractical for administrators to manage and likely put enormous strain on a single DNS server [84].

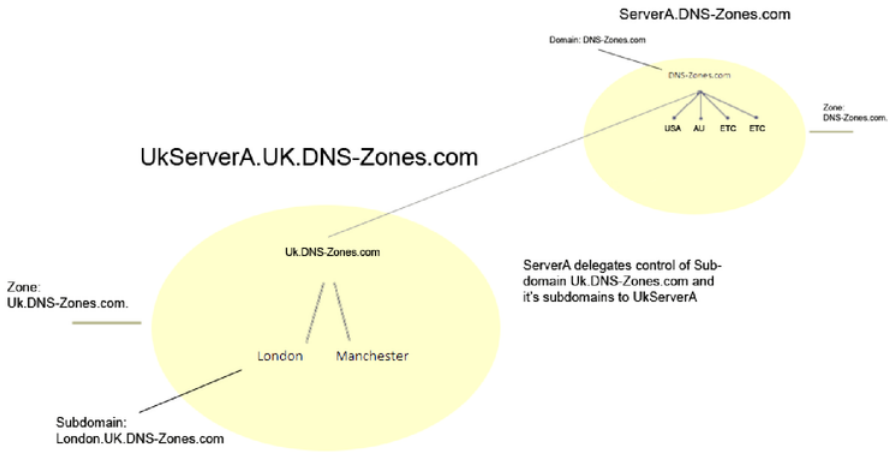


Fig. 148. An example of using multiple zones [84]

Fig. 148. shows how the domain namespace can be split up into different zones. ServerA still hosts the zone for DNS-Zones.com but now UkServerA.UK.DNS-Zones.com hosts the newly created zone called Uk.DNS-Zones.com. This means that ServerA will not contain any records for the UK.DNS-Zones.com domain and UkServerA will not contain any records for the DNS-Zones.com domain. This is because ServerA and UkServerA are only authoritative for their respective zones [85]. ServerA will just be configured to direct all queries for the UK.DNS-Zones.com domain to UkServerA. UkServerA will also be configured to direct all queries for the DNS-Zones.com domain to ServerA [84]. This means the number of records stored on a DNS server can be controlled.
The use of zones to split up a namespace will reduce bandwidth and unnecessary queries [84] (e.g. a query for the UK subdomain will no longer need to query a server in the US. This is because it will query a local UK server hosting a zone responsible for the UK subdomain). Permissions can also be set on zones to allow the delegation of control to different administrators [84].

It is possible to have multiple zones representing the same part of DNS namespace e.g. to provide fault tolerance. There are three zone types [81], [85]:

- **Primary** – This zone is where all updates are made to records that belong to the zone.
- **Secondary** – A read-only copy of the primary zone.
- **Stub** – A read-only copy of the primary zone that only contains records that identify the DNS servers that are authoritative for a DNS domain name.

To successfully implement a fault tolerance DNS server, zone replication and transfer are required [86]. This is achieved by the master DNS server copying its zone file to a secondary DNS server [81]. Master DNS servers are often primary but can also be secondary [81].

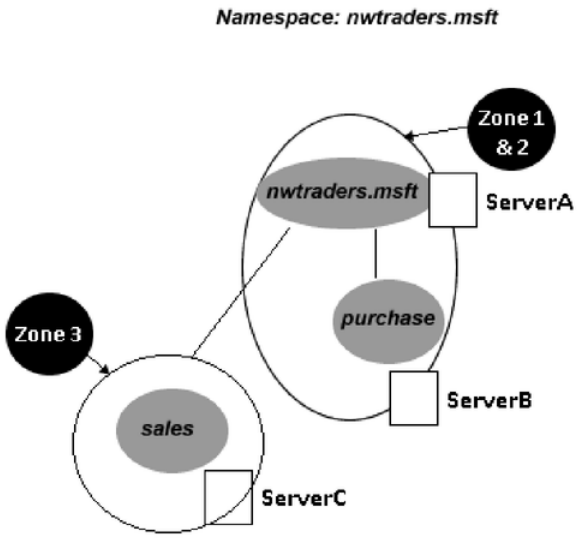


Fig. 149. An example of using multiple zones with fault tolerance [87]

Fig. 149. shows the use of DNS zones to provide fault tolerance. ServerA is hosting the domain nwtraders.msft and Zone 1 (which nwtraders.msft is part of) as a Primary zone. ServerB is also part of the nwtraders.msft namespace and is hosting Zone 2 as a Secondary zone. ServerC is hosting the sales.nwtraders.msft domain and Zone 3 as a Primary zone. ServerA is also configured to replicate its zone file and transfer it to ServerB whenever resource record changes occur. This means that if ServerA fails or becomes unavailable ServerB will have an up-to-date zone file to resolve any queries for the nwtraders.msft domain. If ServerA and ServerB both fail, queries for the nwtraders.msft domain will fail. However, ServerC will still be able to resolve queries for the sales.nwtraders.msft domain as Zone 3 is only responsible for the sales.nwtraders.msft namespace.

It is important to remember that a DNS zone should not be associated with a domain [84], [85]. This is because a DNS zone can contain multiple domains even with multiple zones on the same server [84], [85]. The purpose of a zone is to delegate control and administrative rights of different parts of the namespace (which includes moving resource records of domains to different servers) [81], [84]. DNS zones can also be considered as security boundaries [84].
Factors that require awareness when considering DNZ zones include, the size of a domain namespace and how many subdomains a zone should be responsible for. The need for availability and level of fault tolerance required, as well as security considerations such as zone permissions and administration rights.

Section 5: Evaluation of clustering strategies for hosting a 3-tier web development system

5.1. Introduction

A cluster is a group of servers and other resources that are connected via hardware, networks and software to operate as if they were a single system [88]. There are many reasons to cluster servers that will benefit even small businesses. These reasons include high availability, load balancing, parallel processing, systems management and scalability [88]. High availability is a very important aspect to consider when designing and implementing systems that require minimal downtime. This is because the loss of access to data or a service due to (but not limited to) computer failure and storage failure can have severe consequences to a business. Loss of productivity, revenue, damaged customer relationships, bad publicity and lawsuits are all possible consequences that downtime can cause [89]. Highly available systems use fault tolerance mechanisms to minimise the impact of failures of components and services. Fault tolerance is achieved by reducing single points of failure through the implementation of redundancy [90].
This appraisal will focus on the high availability that failover clustering can provide to a web application architecture, as well as a brief look at application clustering. Cluster topologies, tools and technologies will be investigated along with various suggestions and a proposed clustering strategy.

5.2. Web Architecture

A Web Server is a special type of file server that retrieves files from a server's hard drive, formats them for the web browser and sends the formatted static web content out via a network [91]. Static web content is content that is expected to be the same by everyone that visits the server (HTML pages, files, images, video etc.) [91], [92]. The Web Server standards were intended for the publishing of static documents on the Internet. However, there is a limited capability for accessing dynamic content [91].

Web application servers extend the capabilities of web servers for dynamic content by supporting a high volume of users and highly interactive web application [91]. The application server provides business logic which generates dynamic content for a client. A typical website that provides both static and dynamic content will run web servers for the static content and application servers to generate content dynamically.

5.2.1. Physical Tiers and Distributed Deployment of Applications

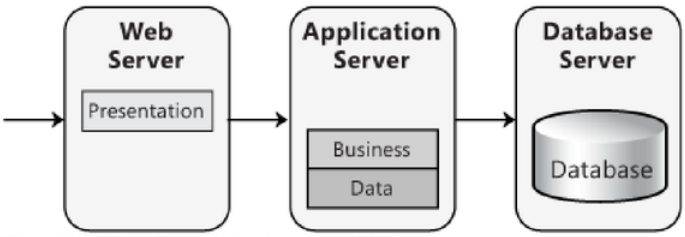


Fig. 150. Distributed deployment [93]

Web servers and application servers can be brought together in the form of a distributed deployment [93]. This is an approach where the layers of an application reside on separate physical tiers as seen in Fig. 150. Multi-tier architectures allow for each tier to be optimised

for its operational requirements and system resource usage [93]. Distributed deployments allow for easier scaling up of each physical tier when processing demands increase [93]. Security can also be enhanced as multiple tiers allow for stricter security policies [93].

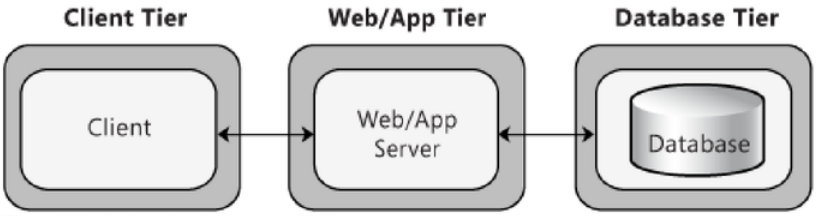


Fig. 151. 3-tier deployment [93]

A 3-tier deployment is a very common approach for most Web applications and services [93]. This is where the application software and database are located on separate servers as seen in Fig. 151. [93]. The client that interacts with the system is also one of the tiers. Note a 3-tier deployment means that the web and application tiers are hosted on one physical layer of hardware and software. The database tier is hosted on another physical layer of hardware and software. This means that three logical tiers (web, application and database) are implemented on two physical tiers [94].

5.2.2. Logical Web Application Tiers

The web/application tier of the 3-tier distributed deployment model can be further broken down into web application tiers [95]. These tiers correspond to the logical services that an application provides.

- Web Tier – This tier provides static content (e.g. simple HTML pages) to the client. The web tier is normally the first contact point (front-end) between a client and a web application [95], [96].
- Presentation Tier – The presentation tier provides dynamic content (e.g. Java servlets, JavaServer Pages or ASP.NET files) to the client [95]–[97].
- Object Tier – The object tier provides business logic to web applications in the form of Java objects or ASP.Net objects (e.g. application beans) [95], [96], [98], [99].

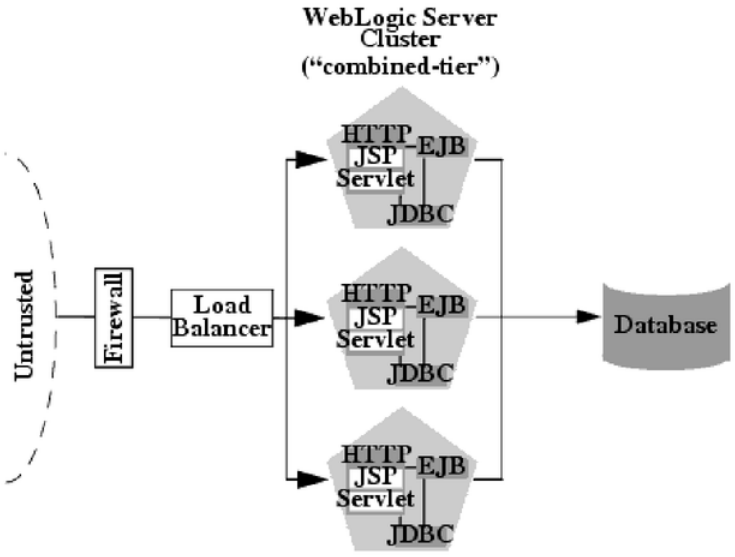


Fig. 152. Basic Combined Tier Architecture [95]

As a 3-tier distributed deployment is being used in this scenario with the three web application tiers deployed on a single server (physical layer), this is called a combined tier architecture (Fig. 152) [95]. The benefits of a combined tier architecture include ease of administration, flexible load balancing, robust security and optimal performance [95]. The web application tiers can be split across multiple physical layers (servers) in a multi-tier architecture [95], [96]. Such a scenario would relate to a 4 or 5-tier distributed deployment [93], [96], [100] which is beyond the scope of this appraisal.

5.3. Application Clustering

The combined tier architecture can take advantage of application clustering [101]. Application clustering allows for an application to take advantage of parallel computing and is used to enhance scalability and high availability (especially with Java applications) [102], [103]. Application clustering can be implemented in two main ways, either state replication or a hub-and-spoke architecture [102]. State replication is where multiple application service instances are running on synchronised cluster nodes to provide fail-over [102], [104]. Hub-and-Spoke is where one application service instance is running on a master cluster node with multiple child nodes invoking the master node service providing improved performance [102]. Application clustering is not always needed but can improve underperforming systems [102]. The benefits of application clustering include load balancing which is the distribution of processing and allows for vertical or horizontal scalability [102], [105], [106]. For example, multiple clients (that are requesting the same application service) can be handled almost simultaneously by the clustered nodes. Fail-over is also a benefit as if an application service fails or is interrupted another instance of the application can take over [101], [102]. Note if the master cluster node fails in a hub-and-spoke architecture all child nodes will fail. However, in a state replicate architecture all nodes are redundant [102]. This fail-over function can be the role of load balancers [102]. As application clustering refers to enabling parallel capabilities of an application on one or more servers. Such capabilities can be extended by taking advantage of application clustering across multiple physical servers, in a strategy called application server clustering [102], [107]. This allows for the utilisation of additional resources (e.g. CPU) while

maintaining the simplicity of a single application server instance [101]. Note application clustering is often deployed with multiple cluster nodes running on VMs or Java Virtual Machines (JVMs) on a physical server, as seen in Fig. 152. This allows a single physical server with multiple CPUs to be fully utilised to its required potential [108]. Clustering of JVMs also allows for more control over Java heap size [108]. This is important to reduce Java Garbage Collection times (an automatic process of freeing up memory resources) [108]–[111].

5.4. Failover Clustering

Application clustering is a worthy consideration, especially when deploying a Java-based web application [105], [107]. However, it should be remembered that multiple application virtual cluster nodes on a single physical server will not protect against hardware failure. This is the role of a failover cluster, to make the physical server(s) highly available [112]. Note an application server cluster (not to be confused with application clustering) is specific to web architecture whereas a failover cluster is a more general term [101], [102], [112].

5.4.1. Cluster Topologies

There are different types of cluster topologies that can be implemented to provide different levels of redundancy in the case of a hardware or software failure [113], [114]:

- **Active-passive** – In this configuration, there is an addition node brought online to replace the role of one specific node that has failed. Every server in a cluster has its own passive redundant server. The failed server will then become the redundant server after repair [115]–[117]. The advantages of this configuration are that it is simple to configure and provides the highest level of redundancy [115]. The disadvantages are hardware costs and a very large overhead as half of all server in a cluster will be offline most of the time [118].
- **Active-active** – This configuration consists of an addition node that is running a specific application/service while also providing redundancy for another specific node. When a failure occurs, the remaining server will take over the role of the failed service in addition to the application that is already running on the server. The redundant server should have enough processing power to take over the role of the failed server, especially if the application is different to what is already running on the redundant server [115]. In active-active, both nodes in the pair will be configured to provide redundancy to each other [117]. The advantages of this configuration are very low overhead as all server in a cluster will be used all the time. It provides a very high level of redundancy [115], [118]. The disadvantages include possible increased hardware cost as every server will need to be in a pair [115]. Another consideration is that as all servers are being used all the time there will be an increased risk of multiple failures as the servers will likely be experiencing similar loads and wear.
- **N+1** – In this configuration a single addition node is brought online to replace the role of a node that has failed [114]. The addition node can be considered a roaming-spare server that is ready to take over the role of any server that fails in the cluster. When the failed server is repaired it will now become the new roaming-spare [118]. This is possible due to the introduction of Storage Area Networks (SANs) as every server can have access to any servers' storage (something not possible in N-to-1 as there is no SAN) [115]. The advantages of this configuration are flexibility as the roaming-space can assume the role of any server in the cluster and reduced hardware costs as only one server will be acting as a space [119]. The disadvantages are limited redundancy as

N+1 cannot provide protection against multiple simultaneous server failures. There will also be some overhead as one server will be offline most of the time [118].

- **N+M** – This configuration is similar to N+1 but instead of a single roaming-space there are multiple. This is often used in a single cluster that is managing many services [114]. The advantages of this configuration are improved redundancy. The disadvantages are increased hardware costs and increased overhead.
- **N-to-1** – This configuration is where a node is a dedicated spare which is cabled to the storage of each server in the cluster [115]. The server will become active to replace the failed server temporarily. When the failed server is repaired the application must be failed back from the acting temporary. The temporary server will now resume its duty as the standby dedicated spare. The same server will always be the spare once the failed server is repaired and put back online [118], [120]. The advantages of this configuration are reduced hardware costs. The disadvantages are limited redundancy, more work for an administrator to fail applications back and some overhead [118].
- **N-to-N** – This configuration is a combination of active-active and N+M [114]. Multiple nodes will be running multiple application/service instances with standby capacity to host instances from a failed node. When a server fails the services or instances that were running on that server will be distributed across the remaining servers. This prevents a single server from becoming overloaded. It is important to have extra capacity on all servers [114], [115]. This configuration provides the same advantages as active-active with even lower overhead and hardware costs as no dedicated redundant pairs need to be formed. The disadvantages are careful planning will be required to make sure all applications are compatible on all servers and that the required extra capacity has been calculated [121]. Multiple failures should also be considered due to similar loads and wear on the servers.

5.4.2. Failover Cluster and Network Load Balancing

A failover cluster is designed for applications that have large, frequently updated data states or in-memory states that are long-running [122]. These include databases, Internet of Things (IoT) and messaging applications which are stateful applications [122], [123].
The general principle of a Network Load Balancer (NLB) is to distribute incoming network traffic across multiple servers and can be either a hardware or software solution [124], [125]. NLBs are designed for stateless applications which do not have long-running in-memory states [122]. This is where every client request is treated as an independent process allowing for requests to be independently load-balanced [122]. NLBs support Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) based services/applications and are often used for front-end web servers, firewalls and proxy servers [122].

Most modern websites and web applications provide some level of dynamic or highly interactive content (making them stateful) [91], [123]. Developers are no longer limited to static content that appears the same to every visitor (HTTP is a stateless protocol) [91], [92]. The rise of stateful web applications makes a failover cluster suited for hosting the presentation and object tiers in a web deployment; providing high availability [122], [126]. NLBs primarily provide scalability which also helps increase the availability of web services [122]. It can also support stateful applications using a combination of various techniques (e.g. client affinity, persistence, session sticky and state replication) [127]–[130]. However, NLBs lacks some functionality when it comes to handling stateful applications [131]. This is because unlike a failover cluster NLBs lacks 100% transparent failover [131]. There is also a potential risk of

data loss if a server fails that is hosting a client's session state [127], [131]. NLBs often cannot detect if a service on a server has crashed meaning a user could be directed to a server that cannot offer the service [132], [133]. For the following reasons and the complexity involved for a network administrator, a failover cluster is an attractive approach for stateful session dependent web applications [126], [134]. However, it should be remembered that hardware or software load balancers are still a vital component to a web architecture and should still be used for scalability (just not as the primary method of failover) [124], [125]. NLBs are ideally suited for front-end web servers which is the web tier of the 3-tier distributed deployment model [122]. Microsoft's Internet Information Services (IIS) is a web server, not an application server [135]. However, IIS can be added as a service to the application server role to host internal and external websites [136], [137].
It should be noted the emergence of advanced Application Delivery Controllers (ADC) that far exceed the capabilities of basic hardware and software network load balancers [138], [139]. ADCs provide a very attractive alternative to failover clustering and may possible supersede such strategies in web application deployments [138], [139].

5.5. Tools and Technologies

There are many technologies and products available for deploying web applications, as well as a huge number of deployment methods and strategies. For example, Windows and Linux could be used as the host OS which could then be hosted on a dedicated server box or a Type 1 Hypervisor such as Hyper-V. Then depending on the programming language platform that the web application is developed a server environment must be chosen (application server) [140]. Java and .Net are popular platforms for building web applications and require different application servers [141], [142]. In the case of Microsoft's .Net framework, Windows Server uses the Application Server Role that can take advantage of various technologies such as ISS and .Net technologies (ASP.NET) [136]. Java uses the Java Platform, Enterprise Edition (Java EE) defined API and features [143], [144]. Java EE is available in commercial application servers such as Oracle WebLogic Server [145] and IBM WebSphere [146]. It is hard for the author to recommend a single set of technologies for a clustered web infrastructure due to the varying requirements, goals and technical expertise in different enterprises.

A few suggestions and considerations include:

- If an organisation's server infrastructure is already based on a set up using Microsoft Hyper-V (or another virtualisation solution). It would seem sensible to integrate the web application into the Hyper-V environment. The failover cluster would be implemented using Hyper-V [112]. The application server installed on one or more VMs would handle application (server) clustering of the VMs [147], [148]. Windows Server 2016 have also introduced a new feature to load balance VMs across nodes based on CPU and memory utilisation [149].
- If an organisation's server infrastructure uses dedicated hardware servers only and/or wishes to implement dedicated hardware for the web system. Such a scenario would allow for increased flexibility as the failover clustering and application clustering can be handled natively by the application server software [150], [151].
- An organisation should consider the expertise and experience of network administrators and technical staff with specific technologies and software.
- Application server software support for application clustering with multiple instances or child nodes. Note application clustering seems far more transparent with Java using

Oracle WebLogic and IBM WebSphere [105], [107], [151]–[153]. The concept of application clustering is a term not used by Microsoft thus finding a similar solution is difficult. A potential solution is using the Microsoft Message Passing Interface Standard (MS-MPI) [154]. A multiple machine virtual High-Performance Computing (HPC) cluster could be built which takes advantage of MPI using the Microsoft HPC pack [155], [156].

5.6. Discussion

Based on the presented research the author has developed the following clustering strategy for a web application:

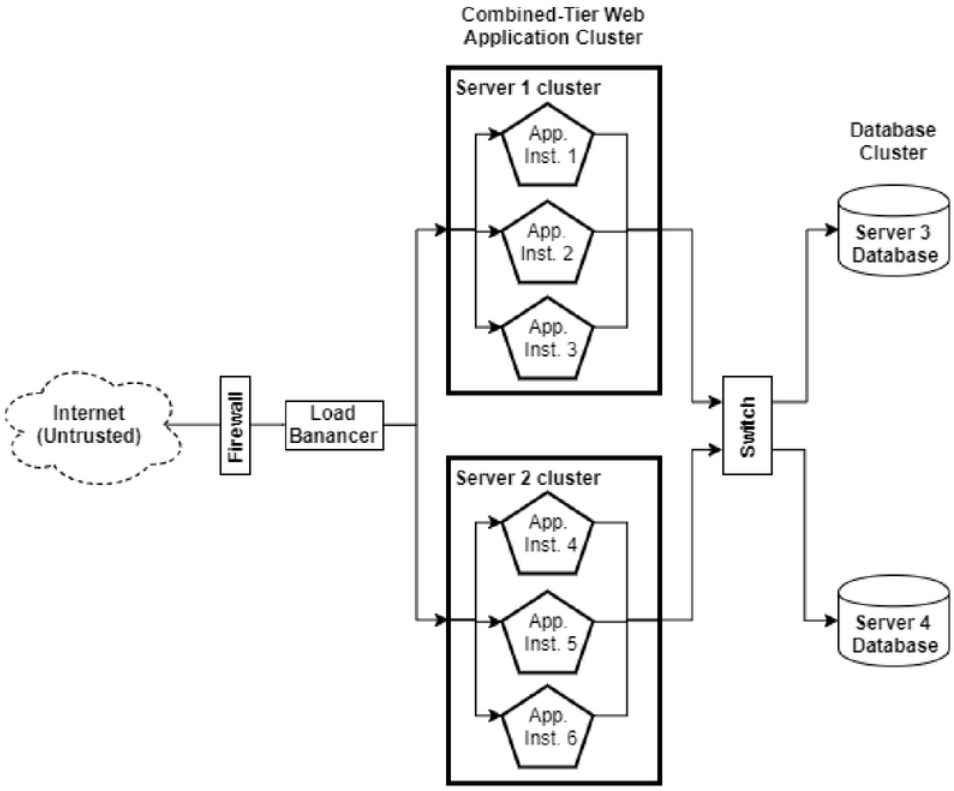


Fig. 153. A generic cluster topology for a web application

Fig. 153. shows a generic vendor neutral N-to-N cluster topology for a web application using a 3-tier distributed deployment architecture. Server 1 and Server 2 are clustered in an N-to-N cluster configuration (can be considered active-active as there are only two servers in the topology). The two servers then have six clustered application instances spread across them to provide application failover, utilise the servers to their full potential and potentially provide parallel computing capabilities. Server 3 and 4 are a separate active-active, multiple-instance database cluster (separate database instances on each server) [157]–[159]. This provides high availability to the back-end database used by the web application. Note database failover clustering has not been mentioned previously but the concept of a failover cluster remains the same no matter the service, although there are unique considerations depending on the database setup and software [45], [159]–[162]. For example, a different software environment will be

used such as SQL server 2017 [163] or Oracle Database 12c Enterprise Edition [164] with Oracle Real Application Clusters (RAC) [162], [165].

The N-to-N configuration was chosen for the application servers as it provides a high level of redundancy with very low overhead. The topology is also very scalable without affecting the level of redundancy as the application instances would be equally spread across remaining servers in the cluster when a failure occurs. All servers would have undergone load calibration to assess average and peak resource usage [166]. It is assumed in this example that the application servers and database servers have an average CPU utilisation of 50% with a peak of around 65%. This means that if a server fails the remaining server can absorb the failed server's workload with no performance penalty most of the time [167]. There would only be some performance issues if a server failed during peak utilisation, which can easily be fixed by adding more servers.

There are a number of further considerations and improvements that could be made to the topology presented in Fig. 153. For example, NIC teaming on the application and database servers for increased bandwidth and further fault tolerance [168], [169]. The implications of clustering on network and application security should be investigated [170], [171]. For example, if the databases support both internal data and externally available web applications an addition firewall should be considered between the application servers and databases [95].

5.7. Conclusion

In conclusion, clustering has many uses and is critically important for achieving the high availability of systems. Failover clustering provides high availability and scalability to servers while application clustering provides high availability, scalability and improved resource efficiency at the application level. There are multiple cluster topologies providing different levels of redundancy and levels of expertise to implement. The N-to-N topology is one of the most attractive options for the enterprise, providing a high level of redundancy with very low overhead.
Web applications require application servers that can run on dedicated server boxes or a Type 1 Hypervisor to provide an environment for hosting software frameworks used to develop web applications (e.g. Java and .Net). Failover clusters primarily provide high availability while NLBs provides scalability. The author's suggested cluster topology provides a feasible clustering strategy for a scalable web application, taking advantage of both failover and application clustering.

Future recommendations include research into the deployment of a specific application server on either a Windows or Linux server. Vendor specific features and technologies can then be explored in detail to provide a more comprehensive clustering strategy for web applications. A more detailed investigation of application clustering, HPCCs and cluster security should also be explored.

References

[1] B. Possey, "System Center Virtual Machine Manager for Beginners (Part 1)," *techgenix.com*, 2014. [Online]. Available: http://techgenix.com/system-center-virtual-machine-manager-beginners-part1/. [Accessed: 20-Dec-2017].

[2] S. Crawford, "What is an IP address?," *HowStuffWorks*. [Online]. Available: https://computer.howstuffworks.com/internet/basics/question549.htm. [Accessed: 18-

Dec-2017].

[3] GROK Knowledge Base, "Definition: IP Address vs. MAC Address," *networking.grok.lsu.edu*, 2016. [Online]. Available: https://networking.grok.lsu.edu/article.aspx?articleid=14849. [Accessed: 18-Dec-2017].

[4] TechTerms, "Intranet Definition," *techterms.com*, 2015. [Online]. Available: https://techterms.com/definition/intranet. [Accessed: 19-Dec-2017].

[5] S. Antoniou, "How to Configure Routing Information Protocol: RIPv2," *pluralsight.com*, 2013. [Online]. Available: https://www.pluralsight.com/blog/tutorials/cisco-how-to-configure-rip-2. [Accessed: 19-Dec-2017].

[6] C. Schluting, "Networking 101: Understanding RIP Routing," *enterprisenetworkingplanet.com*, 2006. [Online]. Available: http://www.enterprisenetworkingplanet.com/netsp/article.php/3609151/Networking-101-Understanding-RIP-Routing.htm. [Accessed: 19-Dec-2017].

[7] A. Bruno and J. Kim, "CCDA Self-Study: RIP, IGRP, and EIGRP Characteristics and Design," *ciscopress.com*, 2003. [Online]. Available: http://www.ciscopress.com/articles/article.asp?p=102174&seqNum=4. [Accessed: 19-Dec-2017].

[8] Scottish Qualifications Authority, "Rip V2," *sqa.org.uk*, 2010. [Online]. Available: https://www.sqa.org.uk/e-learning/NetInf104CD/page_40.htm. [Accessed: 19-Dec-2017].

[9] Microsoft Technet, "IP Routing Table," *technet.microsoft.com*. [Online]. Available: https://technet.microsoft.com/en-us/library/cc958823.aspx. [Accessed: 19-Dec-2017].

[10] B. Mitchell, "What Purpose Does the 127.0.0.1 IP Address Serve?," *lifewire.com*, 2017. [Online]. Available: https://www.lifewire.com/network-computer-special-ip-address-818385. [Accessed: 19-Dec-2017].

[11] Cisco Networking Academy, "Cisco Networking Academy's Introduction to Routing Dynamically," *www.ciscopress.com*, 2014. [Online]. Available: http://www.ciscopress.com/articles/article.asp?p=2180210&seqNum=7. [Accessed: 20-Dec-2017].

[12] VapourApps, "What is Hypervisor and what types of hypervisors are there?," *vapour-apps.com*, 2016. [Online]. Available: http://vapour-apps.com/what-is-hypervisor/. [Accessed: 19-Dec-2017].

[13] B. P. Tholeti, "Hypervisors, virtualization, and the cloud: Learn about hypervisors, system virtualization, and how it works in a cloud environment," *www.ibm.com*, 2011. [Online]. Available: https://www.ibm.com/developerworks/cloud/library/cl-hypervisorcompare/. [Accessed: 19-Dec-2017].

[14] D. Ruest, "Virtualization hypervisor comparison: Type 1 vs. Type 2 hypervisors,"

http://searchservervirtualization.techtarget.com. [Online]. Available: http://searchservervirtualization.techtarget.com/tip/Virtualization-hypervisor-comparison-Type-1-vs-Type-2-hypervisors. [Accessed: 20-Dec-2017].

[15] wing, "Types of Virtualization Technology-Advantages vs. Disadvantages," *securitywing.com*, 2014. [Online]. Available: https://securitywing.com/types-virtualization-technology/. [Accessed: 20-Dec-2017].

[16] D. Prasad, "Comparison Type 1 vs Type 2 Hypervisor," *golinuxhub.com*, 2014. [Online]. Available: http://www.golinuxhub.com/2014/07/comparison-type-1-vs-type-2-hypervisor.html. [Accessed: 20-Dec-2017].

[17] E. Siebert, "Choosing vSphere vs. Hyper-V vs. XenServer," *http://searchitchannel.techtarget.com/*, 2016. [Online]. Available: http://searchitchannel.techtarget.com/feature/Choosing-vSphere-vs-Hyper-V-vs-XenServer. [Accessed: 20-Dec-2017].

[18] A. Sharif, "ESXi vs. ESX: A comparison of features," *blogs.vmware.com*, 2009. [Online]. Available: https://blogs.vmware.com/vsphere/2009/06/esxi-vs-esx-a-comparison-of-features.html. [Accessed: 20-Dec-2017].

[19] Cisco Networking Academy, "Routing and Switching Essentials," *netacad.com*, 2017. [Online]. Available: http://netacad.com. [Accessed: 21-Dec-2017].

[20] E. Devaney, "What Are the Different Types of Routers?," *itstillworks.com*. [Online]. Available: https://itstillworks.com/different-types-routers-8594987.html. [Accessed: 20-Dec-2017].

[21] Computer Hope, "Router," *www.computerhope.com*, 2017. [Online]. Available: https://www.computerhope.com/jargon/r/router.htm. [Accessed: 20-Dec-2017].

[22] ciscorouterswitch, "Routers, To See Different Types of Routers," *ciscorouterswitch.over-blog.com*, 2011. [Online]. Available: http://ciscorouterswitch.over-blog.com/article-routers-to-see-different-types-of-routers-77257064.html. [Accessed: 20-Dec-2017].

[23] Cisco, "Overview of Cisco Interface Cards for Cisco Access Routers," *www.cisco.com*, 2007. [Online]. Available: https://www.cisco.com/c/en/us/td/docs/routers/access/interfaces/ic/hardware/installation/guide/oview_ic.html. [Accessed: 20-Dec-2017].

[24] sdx central, "What's a Virtual Router (vRouter)?," *www.sdxcentral.com*. [Online]. Available: https://www.sdxcentral.com/nfv/definitions/whats-a-virtual-router-vrouter/. [Accessed: 20-Dec-2017].

[25] D.A.R.Y.L., "The Routing Table," *www.think-like-a-computer.com*, 2011. [Online]. Available: http://www.think-like-a-computer.com/2011/08/24/the-routing-table/. [Accessed: 20-Dec-2017].

[26] M. Martin, "The importance of the subnet mask,"

http://searchnetworking.techtarget.com. [Online]. Available: http://searchnetworking.techtarget.com/tip/The-importance-of-the-subnet-mask. [Accessed: 20-Dec-2017].

[27] Network Sorcery, "Multicast addresses," *RFC Sourcebook*, 2012. [Online]. Available: http://www.networksorcery.com/enp/protocol/ip/multicast.htm. [Accessed: 21-Dec-2017].

[28] Learn Networking, "How a Broadcast Address Works," *learn-networking.com*, 2008. [Online]. Available: http://learn-networking.com/network-design/how-a-broadcast-address-works. [Accessed: 21-Dec-2017].

[29] Meridian Outpost, "Well-Known TCP/IP Port Numbers," *www.meridianoutpost.com*. [Online]. Available: http://www.meridianoutpost.com/resources/articles/well-known-tcpip-ports.php. [Accessed: 21-Dec-2017].

[30] S. Gibson, "Port 21," *Gibson Research Corporation*. [Online]. Available: https://www.grc.com/port_21.htm. [Accessed: 21-Dec-2017].

[31] Cisco, "Internet Control Message Protocol (ICMP) Filtering Configuration on the 300 Series Managed Switches," *sbkb.cisco.com*. [Online]. Available: https://sbkb.cisco.com/CiscoSB/GetArticle.aspx?docid=1f54ff141e964770bd35c6d746bca7e5_ICMP_Filtering_Configuration_on_Cisco_300_Managed_Switch.xml&pid=2&converted=0. [Accessed: 21-Dec-2017].

[32] Microsoft Technet, "IPv4 Physical Address Resolution," *technet.microsoft.com*, 2009. [Online]. Available: https://technet.microsoft.com/en-us/library/dd392257(v=ws.10).aspx. [Accessed: 22-Dec-2017].

[33] Cisco Networking Academy, "CCNA Routing and Switching: Introduction to Networks," *netacad.com*, 2017. [Online]. Available: http://netacad.com. [Accessed: 22-Dec-2017].

[34] DistributedNetworks, "Purpose of Address Resolution Protocol," *www.distributednetworks.com*. [Online]. Available: https://www.distributednetworks.com/linux-network-administration/module3/address-resolution-protocol.php. [Accessed: 22-Dec-2017].

[35] Microsoft Technet, "IP Packet Filtering," *technet.microsoft.com*. [Online]. Available: https://technet.microsoft.com/en-gb/library/cc957881.aspx. [Accessed: 22-Dec-2017].

[36] Techopedia, "What is Packet Filtering?," *www.techopedia.com*. [Online]. Available: https://www.techopedia.com/definition/4038/packet-filtering. [Accessed: 22-Dec-2017].

[37] D. Strom, "The Packet Filter: A Basic Network Security Tool," *Glob. Inf. Assur. Certif. Pap.*, 2002.

[38] checkpoint.com, "Packet Filters," *www.checkpoint.com*. [Online]. Available: https://www.checkpoint.com/smb/help/utm1/8.2/7078.htm. [Accessed: 22-Dec-2017].

[39] Microsoft Technet, "Server Roles and Technologies in Windows Server 2012 R2 and Windows Server 2012," *technet.microsoft.com*, 2014. [Online]. Available: https://technet.microsoft.com/en-us/library/hh831669(v=ws.11).aspx. [Accessed: 22-Dec-2017].

[40] Pluralsight, "Active Directory Administration," *www.pluralsight.com*, 2017. [Online]. Available: https://www.pluralsight.com/paths/active-directory-administration. [Accessed: 22-Dec-2017].

[41] Microsoft Press, "Active Directory," *msdn.microsoft.com*. [Online]. Available: https://msdn.microsoft.com/en-us/library/bb742424.aspx. [Accessed: 22-Dec-2017].

[42] Techopedia, "What is Active Directory (AD)?," *www.techopedia.com*. [Online]. Available: https://www.techopedia.com/definition/25/active-directory. [Accessed: 22-Dec-2017].

[43] T. Jacoberger, "What is Active Directory? Why Does My Company Need It?," *TCI Technologies Blog*, 2016. [Online]. Available: http://blog.tcitechs.com/blog/active-directory-company-need/. [Accessed: 22-Dec-2017].

[44] Microsoft Technet, "Networking Overview," *technet.microsoft.com*, 2014. [Online]. Available: https://technet.microsoft.com/en-us/library/hh831357(v=ws.11).aspx. [Accessed: 22-Dec-2017].

[45] Microsoft Docs, "Windows Server Failover Clustering with SQL Server," *docs.microsoft.com*, 2017. [Online]. Available: https://docs.microsoft.com/en-us/sql/sql-server/failover-clusters/windows/windows-server-failover-clustering-wsfc-with-sql-server. [Accessed: 22-Dec-2017].

[46] Microsoft Technet, "File and Storage Services Overview," *technet.microsoft.com*, 2016. [Online]. Available: https://technet.microsoft.com/en-us/library/hh831487(v=ws.11).aspx. [Accessed: 22-Dec-2017].

[47] Microsoft Technet, "Remote Desktop Services Overview," *technet.microsoft.com*, 2014. [Online]. Available: https://technet.microsoft.com/en-us/library/hh831447(v=ws.11).aspx. [Accessed: 22-Dec-2017].

[48] Microsoft Technet, "DHCP: The server should be authorized," *technet.microsoft.com*. [Online]. Available: https://technet.microsoft.com/es-es/library/ee941131(v=ws.10).aspx. [Accessed: 03-Feb-2018].

[49] D. Huggins and E. Tittel, "Creating and Managing Scopes, Superscopes, and Multicast Scopes," in *DHCP for Windows 2000 Network Infrastructure Exam #70-216*, 1st ed., Upper Saddle River, United States: Pearson Education (US), 2003, p. 416.

[50] Dell, "Windows Server: DHCP Reservations vs. Exclusions," *Dell Knowledge Base*, 2014. [Online]. Available: https://www.dell.com/support/article/uk/en/ukbsdt1/sln290628/windows-server-dhcp-reservations-vs-exclusions?lang=en. [Accessed: 03-Feb-2018].

[51] B. Steele, "Why Use Static Addresses for Printers?," *Chron*. [Online]. Available: http://smallbusiness.chron.com/use-static-addresses-printers-57587.html. [Accessed: 03-Feb-2018].

[52] Microsoft Technet, "What Is DHCP?: Dynamic Host Configuration Protocol (DHCP)," *technet.microsoft.com*, 2003. [Online]. Available: https://technet.microsoft.com/en-gb/library/cc781008(v=ws.10).aspx. [Accessed: 23-Dec-2017].

[53] Microsoft Technet, "Troubleshooting DHCP Clients," *technet.microsoft.com*. [Online]. Available: https://technet.microsoft.com/en-us/library/cc958902.aspx. [Accessed: 23-Dec-2017].

[54] Microsoft Technet, "Clustering DHCP Servers," *technet.microsoft.com*. [Online]. Available: https://technet.microsoft.com/en-us/library/cc958897.aspx. [Accessed: 23-Dec-2017].

[55] Microsoft Technet, "How DHCP Technology Works," *technet.microsoft.com*, 2003. [Online]. Available: https://technet.microsoft.com/en-gb/library/cc780760(v=ws.10).aspx. [Accessed: 23-Dec-2017].

[56] Microsoft Technet, "Configure the IPv4 DHCP Relay Agent," *technet.microsoft.com*, 2009. [Online]. Available: https://technet.microsoft.com/en-us/library/dd469685(v=ws.11).aspx. [Accessed: 23-Dec-2017].

[57] R. Droms, "Dynamic Host Configuration Protocol," *Netw. Work. Gr.*, 1997.

[58] Microsoft Technet, "Design Option 2: BOOTP/DHCP Forwarding on Routers," *technet.microsoft.com*, 2008. [Online]. Available: https://technet.microsoft.com/en-us/library/dd296698(v=ws.10).aspx. [Accessed: 23-Dec-2017].

[59] The Network Encyclopedia, "DHCP relay agent," *www.thenetworkencyclopedia.com*. [Online]. Available: http://www.thenetworkencyclopedia.com/entry/dhcp-relay-agent/. [Accessed: 23-Dec-2017].

[60] Microsoft Technet, "What Are Domains and Forests?," *technet.microsoft.com*, 2014. [Online]. Available: https://technet.microsoft.com/en-us/library/cc759073(v=ws.10).aspx#w2k3tr_logic_what_ovkc. [Accessed: 04-Feb-2018].

[61] Pluralsight, G. Grudzinskas, T. Miller, and G. Shields, "Active Directory Administration," 2017. [Online]. Available: https://www.pluralsight.com/paths/active-directory-administration. [Accessed: 22-Dec-2017].

[62] University of York, "Active Directory," *IT Services*. [Online]. Available: https://www.york.ac.uk/it-services/services/active-directory/. [Accessed: 04-Feb-2018].

[63] Dell, "DNS Forwarders vs. Root Hints in Windows Server," *Dell Knowledge Base*, 2015. [Online]. Available: https://www.dell.com/support/article/uk/en/ukbsdt1/sln156952/dns-forwarders-vs-root-hints-in-windows-server?lang=en. [Accessed: 13-Jan-2018].

[64] L. Notenboom, "How does flushing a DNS cache help resolve some issues, and while you're at it what's DNS?," *Ask Leo!*, 2008. [Online]. Available: https://askleo.com/how_does_flushing_a_dns_cache_help_resolve_some_issues_and_while_youre_at_it_whats_dns/. [Accessed: 04-Feb-2018].

[65] S. Laverty, "Is There a Need to Clean the DNS Cache?," *It Still Works*. [Online]. Available: https://itstillworks.com/there-need-clean-dns-cache-25636.html. [Accessed: 04-Feb-2018].

[66] M. Tulloch, "Active Directory Insights (Part 1): Configuring DNS on domain controllers," *TechGenix*, 2015. [Online]. Available: http://techgenix.com/active-directory-insights-part1/. [Accessed: 04-Feb-2018].

[67] A. Fekay (MVP), "Recommended DNS TCP/IP configuration for Domain Controllers," *social.technet.microsoft.com*, 2011. [Online]. Available: https://social.technet.microsoft.com/Forums/en-US/045b0965-2b75-42b0-bafc-45d2d839c0d5/recommended-dns-tcpip-configuration-for-domain-controllers?forum=winserverNIS. [Accessed: 04-Feb-2018].

[68] B. Mayo, "Server Based Network Guide," *speedguide.net*, 2005. [Online]. Available: https://www.speedguide.net/articles/server-based-network-guide-1660. [Accessed: 04-Feb-2018].

[69] Microsoft Docs, "Add a Forward Lookup Zone," *docs.microsoft.com*, 2012. [Online]. Available: https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2008-R2-and-2008/cc771566(v=ws.10). [Accessed: 04-Feb-2018].

[70] Microsoft Technet, "Adding a Reverse Lookup Zone," *technet.microsoft.com*. [Online]. Available: https://technet.microsoft.com/en-us/library/cc961414.aspx. [Accessed: 04-Feb-2018].

[71] Microsoft Technet, "Full Zone Transfer," *technet.microsoft.com*. [Online]. Available: https://technet.microsoft.com/en-us/library/cc958966.aspx. [Accessed: 04-Feb-2018].

[72] Microsoft Technet, "Incremental Transfer," *technet.microsoft.com*. [Online]. Available: https://technet.microsoft.com/en-us/library/cc958973.aspx. [Accessed: 04-Feb-2018].

[73] JH Software, "Definition - Zone Transfer," *Simple DNS Plus*. [Online]. Available: http://simpledns.com/help/definition-zone-transfer. [Accessed: 04-Feb-2018].

[74] Microsoft Technet, "DNS Notify," *technet.microsoft.com*. [Online]. Available: https://technet.microsoft.com/en-us/library/cc958986.aspx. [Accessed: 04-Feb-2018].

[75] Microsoft Support, "Microsoft TCP/IP Host Name Resolution Order," *support.microsoft.com*. [Online]. Available: https://support.microsoft.com/en-us/help/172218/microsoft-tcp-ip-host-name-resolution-order. [Accessed: 24-Dec-2017].

[76] A. Khanse, "Hosts File in Windows 10/8/7," *www.thewindowsclub.com*, 2013. [Online]. Available: http://www.thewindowsclub.com/hosts-file-in-windows. [Accessed: 24-

Dec-2017].

[77] Microsoft Technet, "Recursive and Iterative Queries," *technet.microsoft.com*. [Online]. Available: https://technet.microsoft.com/en-us/library/cc961401.aspx. [Accessed: 24-Dec-2017].

[78] Microsoft Technet, "Host Name Resolution," *technet.microsoft.com*. [Online]. Available: https://technet.microsoft.com/en-us/library/cc958812.aspx. [Accessed: 24-Dec-2017].

[79] D. Karrenberg, "DNS Root Name Servers Frequently Asked Questions," *www.isoc.org*, 2008. [Online]. Available: https://www.isoc.org/briefings/020/. [Accessed: 24-Dec-2017].

[80] Microsoft Technet, "Root Hints Files," *technet.microsoft.com*. [Online]. Available: https://technet.microsoft.com/en-us/library/cc958982.aspx. [Accessed: 13-Jan-2018].

[81] Microsoft Technet, "DNS Architecture," *technet.microsoft.com*. [Online]. Available: https://technet.microsoft.com/en-us/library/dd197427(v=ws.10).aspx. [Accessed: 13-Jan-2018].

[82] Microsoft Support, "DNS Namespace Planning," *support.microsoft.com*, 2017. [Online]. Available: https://support.microsoft.com/en-us/help/254680/dns-namespace-planning. [Accessed: 13-Jan-2018].

[83] Raven, "What is a Root Domain?," *Raven Marketing Glossary*. [Online]. Available: https://raventools.com/marketing-glossary/root-domain/. [Accessed: 13-Jan-2018].

[84] D.A.R.Y.L, "DNS Zones Explained," *Think Like A Computer*, 2011. [Online]. Available: https://www.think-like-a-computer.com/2011/06/11/dns-zones-explained/. [Accessed: 13-Jan-2018].

[85] Microsoft Technet, "Zones," *technet.microsoft.com*. [Online]. Available: https://technet.microsoft.com/en-us/library/cc958980.aspx. [Accessed: 13-Jan-2018].

[86] Microsoft Technet, "Understanding Zones," *technet.microsoft.com*. [Online]. Available: https://technet.microsoft.com/en-us/library/cc725590(v=ws.11).aspx. [Accessed: 13-Jan-2018].

[87] T. Bashford and L. Cobley, "B.Sc. Year 2 Network Services and Web Technology. WorkSheet 4: Name Resolution using the Domain Name Service (DNS)," Swansea, 2018.

[88] IBM developerWorks, "Clustering: A basic 101 tutorial," *www.ibm.com*, 2002. [Online]. Available: https://www.ibm.com/developerworks/aix/tutorials/clustering/clustering.html. [Accessed: 26-Dec-2017].

[89] Oracle, "Overview of High Availability," *docs.oracle.com*. [Online]. Available: https://docs.oracle.com/cd/B19306_01/server.102/b14210/overview.htm. [Accessed:

26-Dec-2017].

[90] Microsoft Technet, "High Availability and Clustering Technologies," *technet.microsoft.com*. [Online]. Available: https://technet.microsoft.com/en-gb/library/cc535053.aspx. [Accessed: 26-Dec-2017].

[91] C. Bishop, "What is a Web Application Server?," *Resultant Systems*, 2003. [Online]. Available: http://www.resultantsys.com/index.php/general/what-is-a-web-application-server/. [Accessed: 26-Dec-2017].

[92] NGINX, "What Is a Web Server vs. an Application Server," *www.nginx.com*. [Online]. Available: https://www.nginx.com/resources/glossary/application-server-vs-web-server/. [Accessed: 26-Dec-2017].

[93] Microsoft Developer Network, "Chapter 19: Physical Tiers and Deployment," *msdn.microsoft.com*. [Online]. Available: https://msdn.microsoft.com/en-gb/library/ee658120.aspx. [Accessed: 26-Dec-2017].

[94] P. R. Reed, "Logical versus Physical Tiers," *informIT*, 2002. [Online]. Available: http://www.informit.com/articles/article.aspx?p=26018&seqNum=4. [Accessed: 26-Dec-2017].

[95] Oracle, "Cluster Architectures," *docs.oracle.com*, 2016. [Online]. Available: https://docs.oracle.com/cd/E13222_01/wls/docs103/cluster/planning.html#wp1115757. [Accessed: 26-Dec-2017].

[96] R. Tyagi, "Enterprise application clustering architecture," *TechRepublic*, 2002. [Online]. Available: https://www.techrepublic.com/article/enterprise-application-clustering-architecture/. [Accessed: 26-Dec-2017].

[97] MicroStrategy, "Presentation Layer," *lw.microstrategy.com*. [Online]. Available: https://lw.microstrategy.com/MSDZ/MSDL/GARelease_Current/docs/projects/WebSDK/Content/topics/webarch/Presentation_Layer.htm. [Accessed: 27-Dec-2017].

[98] MicroStrategy, "Application Objects Layer," *lw.microstrategy.com*. [Online]. Available: https://lw.microstrategy.com/MSDZ/MSDL/GARelease_Current/docs/projects/WebSDK/Content/topics/webarch/Application_Objects_Layer.htm. [Accessed: 27-Dec-2017].

[99] Microsoft Developer Network, "ObjectDataSource Web Server Control Overview," *msdn.microsoft.com*. [Online]. Available: https://msdn.microsoft.com/en-us/library/9a4kyhcx.aspx. [Accessed: 27-Dec-2017].

[100] V. Thakur, "5-Tier Architecture," *ASP.NET 3.5 Application Architecture and Design*, 2008. [Online]. Available: https://www.packtpub.com/mapt/book/web_development/9781847195500/4/ch04lvl1sec03/5-tier-architecture. [Accessed: 27-Dec-2017].

[101] Oracle, "Application Server Clustering," *docs.oracle.com*, 2002. [Online]. Available: https://docs.oracle.com/cd/A97329_03/core.902/a92171/cluster.htm#1010630.

[Accessed: 27-Dec-2017].

[102] A. Brown, "Application Clustering For Scalability and High Availability," *DZone Performance*, 2009. [Online]. Available: https://dzone.com/articles/application-clustering. [Accessed: 26-Dec-2017].

[103] S. Popa, "Architectures for J2EE Application Clustering in BEA and IBM Application Servers," *T-110.551 Internetworking Semin.*, 2004.

[104] B. Goetz, "State replication in the Web tier," *IBM developerWorks*, 2004. [Online]. Available: https://www.ibm.com/developerworks/library/j-jtp07294/index.html. [Accessed: 15-Jan-2018].

[105] IBM Knowledge Center, "Application server clustering," *www.ibm.com*. [Online]. Available: https://www.ibm.com/support/knowledgecenter/en/SSZJPZ_11.5.0/com.ibm.swg.im.iis.productization.iisinfsv.ha.install.doc/topics/wsisinst_pln_ha_appserver_clustering.html. [Accessed: 27-Dec-2017].

[106] D. Beaumont, "How to explain vertical and horizontal scaling in the cloud," *IBM Cloud computing news*, 2014. [Online]. Available: https://www.ibm.com/blogs/cloud-computing/2014/04/explain-vertical-horizontal-scaling-cloud/. [Accessed: 27-Dec-2017].

[107] Oracle, "Understanding WebLogic Server Clustering," *docs.oracle.com*. [Online]. Available: https://docs.oracle.com/cd/E11035_01/wls100/cluster/overview.html. [Accessed: 27-Dec-2017].

[108] A. Kang, "J2EE clustering, Part 2," *JavaWorld*, 2001. [Online]. Available: https://www.javaworld.com/article/2075478/core-java/j2ee-clustering--part-2.html. [Accessed: 15-Jan-2018].

[109] S. Lee, "Understanding Java Garbage Collection," *CUBRID blog*, 2017. [Online]. Available: https://www.cubrid.org/blog/understanding-java-garbage-collection. [Accessed: 27-Dec-2017].

[110] Stackify, "What is Java Garbage Collection? Best Practices, Tutorials & More," *Developer Tips, Tricks & Resources*, 2017. [Online]. Available: https://stackify.com/what-is-java-garbage-collection/. [Accessed: 15-Jan-2018].

[111] P. Haridas, "Importance of the Java heap size In WebSphere Application Server," *IBM developerWorks Blog*, 2015. [Online]. Available: https://www.ibm.com/developerworks/community/blogs/aimsupport/entry/java_heap_was?lang=en. [Accessed: 15-Jan-2018].

[112] Microsoft Technet, "Failover Clustering Overview," *technet.microsoft.com*, 2013. [Online]. Available: https://technet.microsoft.com/en-us/library/hh831579(v=ws.11).aspx. [Accessed: 15-Jan-2018].

[113] IBM Knowledge Center, "High availability models," *www.ibm.com*. [Online]. Available:

https://www.ibm.com/support/knowledgecenter/en/SSLKT6_7.6.0.5/com.ibm.mbs.doc/gp_highavail/c_cluster_config_models.html. [Accessed: 26-Dec-2017].

[114] R. Shetty, "High Availability: Planning Redundancy," *IBM developerWorks*, 2014. [Online]. Available: https://www.ibm.com/developerworks/community/blogs/RohitShetty/entry/high_availability_planning_redundancy?lang=en. [Accessed: 26-Dec-2017].

[115] UnixArena, "Veritas cluster concepts," *www.unixarena.com*, 2012. [Online]. Available: https://www.unixarena.com/2012/07/veritas-cluster-concepts.html. [Accessed: 16-Jan-2018].

[116] IBM Knowledge Center, "Multiple clusters," *www.ibm.com*. [Online]. Available: https://www.ibm.com/support/knowledgecenter/SSHRKX_8.0.0/plan/top_multipleservers.html. [Accessed: 26-Dec-2017].

[117] J. C. Villanueva, "Active-Active vs Active-Passive High Availability Cluster," *Jscape*, 2015. [Online]. Available: http://www.jscape.com/blog/active-active-vs-active-passive-high-availability-cluster. [Accessed: 26-Dec-2017].

[118] Eileen, "N+1 is not the same as N-to-1," *vox.veritas.com*, 2009. [Online]. Available: https://vox.veritas.com/t5/Articles/Multinode-clustering-N-1-is-not-the-same-as-N-to-1-and-it/ta-p/805822. [Accessed: 26-Dec-2017].

[119] Veritas Technologies LLC, "About the N + 1 configuration," *About advanced failover configurations*. [Online]. Available: https://sort.veritas.com/public/documents/vie/7.1/aix/productguides/html/vcs_admin/ch02s02s01.htm. [Accessed: 26-Dec-2017].

[120] Veritas Technologies LLC, "About N-to-1 configuration," *Basic failover configurations*. [Online]. Available: https://sort.veritas.com/public/documents/via/7.0/aix/productguides/html/vcs_admin/ch02s01s03.htm. [Accessed: 16-Jan-2018].

[121] Veritas Technologies LLC, "About the N-to-N configuration," *About advanced failover configurations*. [Online]. Available: https://sort.veritas.com/public/documents/vie/7.1/aix/productguides/html/vcs_admin/ch02s02s02.htm. [Accessed: 26-Dec-2017].

[122] Microsoft Technet, "Network Load Balancing Overview," *technet.microsoft.com*. [Online]. Available: https://technet.microsoft.com/en-us/library/cc725946(v=ws.11).aspx. [Accessed: 16-Jan-2018].

[123] A. Stannard, "The Inevitable Rise of the Stateful Web Application," *Petabridge*, 2015. [Online]. Available: https://petabridge.com/blog/stateful-web-applications/. [Accessed: 16-Jan-2018].

[124] Elastic Load Balancing, "What Is a Network Load Balancer?," *docs.aws.amazon.com*. [Online]. Available: https://docs.aws.amazon.com/elasticloadbalancing/latest/network/introduction.html.

[Accessed: 17-Jan-2018].

[125] NGINX Software, "What Is Load Balancing?," *www.nginx.com*. [Online]. Available: https://www.nginx.com/resources/glossary/load-balancing/. [Accessed: 17-Jan-2018].

[126] Microsoft Docs, "Understanding Generic Services and Applications that Can Be Configured in a Failover Cluster," *docs.microsoft.com*, 2009. [Online]. Available: https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2008-R2-and-2008/cc753938(v=ws.11). [Accessed: 17-Jan-2018].

[127] B. Assmann, "load balancing, affinity, persistence, sticky sessions: what you need to know," *HAProxy Technologies*, 2012. [Online]. Available: https://www.haproxy.com/blog/load-balancing-affinity-persistence-sticky-sessions-what-you-need-to-know/. [Accessed: 17-Jan-2018].

[128] Microsoft Technet, "Chapter 5 - Load Balancing," *technet.microsoft.com*. [Online]. Available: https://technet.microsoft.com/en-us/library/bb734910.aspx. [Accessed: 17-Jan-2018].

[129] Microsoft Technet, "Using NLB," *technet.microsoft.com*. [Online]. Available: https://technet.microsoft.com/en-us/library/bb687542.aspx. [Accessed: 17-Jan-2018].

[130] Microsoft Technet, "Preserving Session State in Network Load Balancing Web Server Clusters (IIS 6.0)," *Improving IIS 6.0 Scalability and Availability with Network Load Balancing*. [Online]. Available: https://www.microsoft.com/technet/prodtechnol/WindowsServer2003/Library/IIS/77cb4318-75f8-4310-a05f-3605b5768007.mspx?mfr=true. [Accessed: 01-Jan-2018].

[131] D. Van Hoye (MVP), "NLB for Stateful Applications is possible ?," *social.technet.microsoft.com*, 2011. [Online]. Available: https://social.technet.microsoft.com/Forums/windowsserver/en-US/c124b4b4-fa2a-4764-b99a-5b2228b16430/nlb-for-stateful-applications-is-possible-?forum=winservergen. [Accessed: 17-Jan-2018].

[132] James, "Microsoft Cluster Service vs NLB," *jbcomp.com*, 2014. [Online]. Available: http://jbcomp.com/microsoft-cluster-service-vs-nlb/. [Accessed: 16-Jan-2018].

[133] S. Perriman, "Introducing NLB Application Health Awareness," *Failover Clustering and Network Load Balancing Team Blog*, 2008. [Online]. Available: https://blogs.msdn.microsoft.com/clustering/2008/06/24/introducing-nlb-application-health-awareness/. [Accessed: 20-Jan-2018].

[134] LaunchSchool, "Stateful Web Applications," *launchschool.com*. [Online]. Available: https://launchschool.com/books/http/read/statefulness. [Accessed: 17-Jan-2018].

[135] Microsoft Docs, "Web Server (IIS) Support for Application Server," *docs.microsoft.com*, 2013. [Online]. Available: https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2008-R2-and-2008/cc753830(v=ws.11). [Accessed: 20-Jan-2018].

[136] Microsoft Docs, "Application Server Role," *docs.microsoft.com*, 2012. [Online]. Available: https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2008-R2-and-2008/cc754024(v=ws.10). [Accessed: 19-Jan-2018].

[137] Microsoft Docs, "Application Server: Frequently Asked Questions," *docs.microsoft.com*, 2012. [Online]. Available: https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2008-R2-and-2008/cc725850(v=ws.10). [Accessed: 20-Jan-2018].

[138] F5 Networks Inc., "The Evolution of Application Delivery Controllers," 2017.

[139] L. MacVittie, "Clustering versus load-balancing," *devcentral.f5.com*, 2017. [Online]. Available: https://devcentral.f5.com/articles/clustering-versus-load-balancing. [Accessed: 16-Jan-2018].

[140] B. Mitchell, "Application Servers in Computer Networking," *INTERNET & NETWORK*, 2017. [Online]. Available: https://www.lifewire.com/what-is-an-application-server-818199. [Accessed: 19-Jan-2018].

[141] M. Marinina, "Java vs .NET: Factors to Consider," *codeburst*, 2017. [Online]. Available: https://codeburst.io/java-vs-net-factors-to-consider-cde1d22b06a7. [Accessed: 19-Jan-2018].

[142] Rishabh Software, ".Net vs Java: Comparison Of The Platforms," *www.rishabhsoft.com*, 2017. [Online]. Available: https://www.rishabhsoft.com/blog/dot-net-vs-java. [Accessed: 19-Jan-2018].

[143] M. Rouse, "What is Java Platform, Enterprise Edition (Java EE)," *TheServerSide*, 2017. [Online]. Available: http://www.theserverside.com/definition/J2EE-Java-2-Platform-Enterprise-Edition. [Accessed: 19-Jan-2018].

[144] Oracle, "Java Platform, Enterprise Edition (Java EE)," *Oracle Technology Network*. [Online]. Available: http://www.oracle.com/technetwork/java/javaee/overview/index.html. [Accessed: 19-Jan-2018].

[145] Oracle, "Oracle WebLogic Server," *Oracle Fusion Middleware*. [Online]. Available: https://www.oracle.com/uk/middleware/weblogic/index.html. [Accessed: 19-Jan-2018].

[146] IBM, "IBM WebSphere software," *IBM Software*. [Online]. Available: https://www-01.ibm.com/software/uk/websphere/. [Accessed: 19-Jan-2018].

[147] L. Dignan, "Oracle databases head to Microsoft's Hyper-V, Azure," *ZDNet*, 2013. [Online]. Available: http://www.zdnet.com/article/oracle-databases-head-to-microsofts-hyper-v-azure/. [Accessed: 20-Jan-2018].

[148] IBM, "WebSphere Application Server Hypervisor Edition overview," *IBM Knowledge Center*, 2017. [Online]. Available: https://www.ibm.com/support/knowledgecenter/en/SSCKBL_8.5.5/com.ibm.websphere.virtualimages.doc/ae/cins_virt.html. [Accessed: 20-Jan-2018].

[149] S. Bhattacharya, "Failover Cluster VM Load Balancing in Windows Server 2016," *Failover Clustering and Network Load Balancing Team Blog*, 2016. [Online]. Available: https://blogs.msdn.microsoft.com/clustering/2016/04/29/failover-cluster-node-fairness-in-windows-server-2016/. [Accessed: 19-Jan-2018].

[150] Oracle, "Setting up WebLogic Clusters," *docs.oracle.com*, 2018. [Online]. Available: https://docs.oracle.com/cd/E13222_01/wls/docs81/cluster/setup.html. [Accessed: 20-Jan-2018].

[151] B. Roehm *et al.*, "WebSphere Application Server Network Deployment V6: High Availability Solutions," 2005.

[152] Oracle, "Failover and Replication in a Cluster," *docs.oracle.com*. [Online]. Available: https://docs.oracle.com/cd/E11035_01/wls100/cluster/failover.html. [Accessed: 20-Jan-2018].

[153] IBM, "WebSphere Application Server clustering for high availability and scalability," *IBM Knowledge Center*. [Online]. Available: https://www.ibm.com/support/knowledgecenter/en/SSZJPZ_8.5.0/com.ibm.swg.im.iis.productization.iisinfsv.ha.install.doc/topics/wsisinst_ha_wasclusteringtop.html. [Accessed: 20-Jan-2018].

[154] Microsoft Developer Network, "Microsoft MPI," *msdn.microsoft.com*. [Online]. Available: https://msdn.microsoft.com/en-us/library/bb524831(v=vs.85).aspx. [Accessed: 19-Jan-2018].

[155] iBrarian, "Virtual HPC Cluster Deployment on Hyper-v," 2008.

[156] Windows HPC, "Overview of Microsoft HPC Pack," *technet.microsoft.com*, 2014. [Online]. Available: https://technet.microsoft.com/en-us/library/ff919321(v=ws.11).aspx. [Accessed: 20-Jan-2018].

[157] K. Little, "An Introduction to SQL Server Clusters with Diagrams," *Brent Ozar*, 2012. [Online]. Available: https://www.brentozar.com/archive/2012/02/introduction-sql-server-clusters/. [Accessed: 21-Jan-2018].

[158] H. Cotter, "Active/Active clusters in SQL Server," *TechTarget SearchSQLServer*, 2006. [Online]. Available: http://searchsqlserver.techtarget.com/tip/Active-Active-clusters-in-SQL-Server. [Accessed: 21-Jan-2018].

[159] B. Consulting, "Failover Cluster Architecture - Oracle Database Service in HA Cluster," *http://www.dba-oracle.com*. [Online]. Available: http://www.dba-oracle.com/real_application_clusters_rac_grid/fo_cluster.htm. [Accessed: 21-Jan-2018].

[160] Microsoft Docs, "High Availability Solutions (SQL Server)," *docs.microsoft.com*, 2016. [Online]. Available: https://docs.microsoft.com/en-us/sql/sql-server/failover-clusters/high-availability-solutions-sql-server. [Accessed: 20-Jan-2018].

[161] Microsoft Developer Network, "Clustering the SQL Servers," *msdn.microsoft.com*,

2005. [Online]. Available: https://msdn.microsoft.com/en-us/library/ee784936(v=cs.20).aspx. [Accessed: 21-Jan-2018].

[162] Oracle, "Using WebLogic Server with Oracle RAC," *docs.oracle.com*, 2016. [Online]. Available: https://docs.oracle.com/cd/E13222_01/wls/docs103/jdbc_admin/oracle_rac.html#wp1080030. [Accessed: 20-Jan-2018].

[163] Microsoft, "SQL Server 2017," *www.microsoft.com*, 2017. [Online]. Available: https://www.microsoft.com/en-gb/sql-server/sql-server-2017. [Accessed: 20-Jan-2018].

[164] Oracle, "Database 12c Enterprise Edition," *Database*. [Online]. Available: https://www.oracle.com/uk/database/enterprise-edition/index.html. [Accessed: 20-Jan-2018].

[165] Oracle, "Comparing Oracle Real Application Cluster to Failover Clustering for Oracle Database."

[166] D. Boutin, "Enough is enough! Or is it? Best practices for load server calibration," *SOASTA*, 2015. [Online]. Available: https://www.soasta.com/blog/load-performance-testing-best-practices/. [Accessed: 21-Jan-2018].

[167] IBM, "Determining how many servers to include in a cluster," *IBM Knowledge Center*, 2009. [Online]. Available: https://www.ibm.com/support/knowledgecenter/en/SSKTMJ_8.5.3/com.ibm.help.domino.admin85.doc/H_DETERMINING_HOW_MANY_SERVERS_TO_INCLUDE_IN_A_CLUSTER_OVER.html. [Accessed: 21-Jan-2018].

[168] Dell, "Windows Server: How Many Network Cards are Recommended for a Hyper-V Cluster Node using Cluster Shared Voumes?," *Dell Knowledge Base*, 2014. [Online]. Available: http://www.dell.com/support/article/uk/en/ukbsdt1/sln155767/windows-server-how-many-network-cards-are-recommended-for-a-hyper-v-cluster-node-using-cluster-shared-voumes-?lang=en. [Accessed: 20-Jan-2018].

[169] K. Mauldin, "5 reasons to use NIC Teaming with Windows Server 2012," *Pluralsight*, 2013. [Online]. Available: https://www.pluralsight.com/blog/it-ops/nic-teaming-windows-server-2012. [Accessed: 21-Jan-2018].

[170] J. D. Meier, A. Mackman, M. Dunner, S. Vasireddy, R. Escamilla, and A. Murukan, "Securing Your Database Server," *Improving Web Application Security: Threats and Countermeasures*, 2006. [Online]. Available: https://msdn.microsoft.com/en-us/library/ff648664.aspx. [Accessed: 21-Jan-2018].

[171] J. D. Meier, A. Mackman, M. Dunner, S. Vasireddy, R. Escamilla, and A. Murukan, "Securing Your Application Server," *Improving Web Application Security: Threats and Countermeasures*, 2006. [Online]. Available: https://msdn.microsoft.com/en-us/library/ff648657.aspx. [Accessed: 21-Jan-2018].

NOS Worksheets 1-5

GRADEMARK REPORT

FINAL GRADE

90/100

GENERAL COMMENTS

Instructor

PAGE 1

PAGE 2

PAGE 3

PAGE 4

Comment 1

Software is a mass-noun, so can't be modified by an indefinite article.

PAGE 5

PAGE 6

PAGE 7

PAGE 8

PAGE 9

PAGE 10

PAGE 11

PAGE 12

PAGE 13

PAGE 14

PAGE 15

Comment 2

While this is quite correct, it's effectively opinion - make sure you state explicitly who says it, and cite it inline rather than above.

PAGE 16

Comment 3

Chunks of this are too close to the source for comfort. Aim to paraphrase robustly by stepping away from the source material before writing your argument. Quoting would also be acceptable, but far less preferable.

Comment 4

Too close to source.

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

PAGE 30

PAGE 31

PAGE 32

PAGE 33

PAGE 34

PAGE 35

PAGE 36

PAGE 37

PAGE 38

PAGE 39

PAGE 40

PAGE 41

PAGE 42

PAGE 43

PAGE 44

PAGE 45

PAGE 46

PAGE 47

PAGE 48

PAGE 49

PAGE 50

PAGE 51

PAGE 52

PAGE 53

PAGE 54

PAGE 55

PAGE 56

PAGE 57

PAGE 58

PAGE 59

PAGE 60

PAGE 61

PAGE 62

PAGE 63

PAGE 64

PAGE 65

PAGE 66

PAGE 67

PAGE 68

PAGE 69

PAGE 70

PAGE 71

PAGE 72

PAGE 73

PAGE 74

PAGE 75

Comment 5

Again, this is generally too close to the source material.

PAGE 76

PAGE 77

PAGE 78

PAGE 79

PAGE 80

PAGE 81

PAGE 82

PAGE 83

PAGE 84

PAGE 85

PAGE 86

Comment 6

Hyphenate

PAGE 87

PAGE 88

PAGE 89

PAGE 90

PAGE 91

PAGE 92

PAGE 93

PAGE 94

PAGE 95

PAGE 96

PAGE 97

PAGE 98

PAGE 99

PAGE 100

GRADING FORM: SMU GENERIC_GJ

GEORGE SEARLE	**90**

RESEARCH-INFORMED LITERATURE

8.5

An incredible number of sources used - I've seen fewer in a PhD thesis. The reference style is correctly applied and consistent. There is, however, a very 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. So my suggestion for the next assignment is to cut down on the number of sources (not that the huge number you have is a bad thing), and try to aim for a higher portion of peer-reviewed material.

KNOWLEDGE AND UNDERSTANDING



Some excellent knowledge and understanding evident. You've obviously researched this work to an extreme degree, and have evidently read around the topics involved in minute detail. You've dropped marks in this section, however, because in a fair few places your work is far too close to the source material you're intending to paraphrase. Paraphrased work should never effectively take a portion of another source and mix in a few synonyms and joining words. The purpose of paraphrasing is to combine a number of sources (and evidently you have used a massive number of sources) into a single, concentrated argument using your own voice, displaying your own understanding of the material. My suggestion is to try to step away from the original material before writing your narrative. Read everything you can around a topic, walk away for a while (to get the phrasing of the last source you read out of your head), then write your version of the content, and cite everything you read.

14

ANALYSIS



Excellent analyses and conclusions drawn. I like your use of explicit conclusions throughout the report. Very good synthesis of materials, discounting the issues mentioned above.

18

PRACTICAL APPLICATION AND DEP.



Fantastic work on the practical implementation of the work. A huge amount of work has clearly gone into the worksheets. I don't give perfect marks easily, but I can't find meaningful fault with your work. Extremely well done, George.

40

SKILLS FOR PROFESSIONAL PRAC.



You have a very good written style - a few very minor grammatical errors (admittedly, being picky). A nicely-presented report. Great work.

9.5