

Ross Van Den Eeckhaute-Jones

P151366

by Ross VanDenEeckhaute-Jones

FILE	2865_ROSS_VANDENEECKHAUTE-JONES_ROSS_VAN_DEN_EECKHAUTE-JONES_P151366_100686_224486022.DOCX (8.15M)		
TIME SUBMITTED	13-FEB-2018 10:05PM (UTC+0000)	WORD COUNT	34 14
SUBMISSION ID	83155917	CHARACTER COUNT	18654

Name: Ross Van Den Eeckhaute-Jones

Student P Number: P151366

Module Title: Network Services

Module Code: M5X01050

Assignment Title: Network Services


Module Lecturer: Dr. Tim Bashford

Date: 13/02/2018


Table of contents

Section	Page Number
Worksheet 1	3-21
Worksheet 2	22-40
Worksheet 3	41-53
Worksheet 4	54-59
Worksheet 5	60-68


Worksheet 1




Lab configuration




Adding network policy and access services




Adding routing to remote access role




Swansea configuration




Attempt to ping network 0 from Swansea




Routing and remote access service configuration




Custom configuration




Lan routing




Service ready to use




Pinging 11 network from lab




Internet connection




Ping from Cardiff to lab




Ping from Cardiff to Swansea




New routing protocol




Ripv2



New interface



2 new rip interfaces



Persistent route to Cardiff network 0

Task 5

Benefits of a type 1 hypervisor

- Applications communicate with hardware with less delay than a type 2 due to the exclusion of the client operating system and more bare bones structure.
- Ease of recovery, restoring a machine due to there being no client OS to reinstall as it is all on the vm software
- Can run multiple versions of vm software at once

The role of a router in a network is to forward packets of data between networks. Data packets are typically sent from router to router until it reaches the router in the network of the intended recipient.

Common routing protocols

RIP – RIP is a distance vector routing protocol which uses hop count to work out the suitable path across a network without creating routing loops. It avoids creating routing loops by “poisoning” routes so that it knows are invalid.


EIGRP – Enhanced interior gateway routing. Designed by Cisco, it is a distance vector routing protocol intended to automate routing decisions for the network. EIGRP shares routing paths to other routers in an autonomous system

OSPF – Open shortest path first uses link state routing to constantly update the logical topology and allow the router to work out the shortest path using the Dijkstra algorithm.


A routing table holds the information gathered by the router about devices on the network. The kind of information stored in a routing table includes network addresses, distance in hop counts to devices and costs of packets being sent. This information is stored in a routers RAM.

A subnet mask is used to identify a single subnet of a network, and it is the protocol that allows the world to use private IPs in a business or home setting. Used in a routing table, it can distinguish from 2 devices using the same IP but are part of different subnets, depending on the subnet mask.


Worksheet 2




Arp-a




Add web server role




ftp services




New user




ftp site test




No ssl port 21




User setup authentication




Message banners




ftp file/folder




Accessing the Swansea FTP from Cardiff




Ftp transfer complete




Uploaded file from Cardiff to Swansea ftp




File confirmed uploaded




Selecting 0 network to apply filtering



Filter settings



Filter settings



Ping fail due to filters


Q1 ARP, address resolution protocol. ARP converts IP addresses into the corresponding MAC addresses. Both the MAC address and IP address are required to send a frame packet between networks. ARP information is stored in an ARP table. The table contains IP and MAC addresses of devices that are on the network. ARP entries are usually entered dynamically, although they can also be entered manually.

Q2 Packet filtering is a system that determines what kind of traffic is allowed or blocked on a network. Unused ports can be blocked to reduce chances of intruders. Ports can also be prioritised to allow faster connections. Filtering can be set up to allow or deny traffic from certain destinations. Workplaces can block unused ports to stop employees accessing restricted websites.


Q3 Contradictory filters can occur when the input filter of one rule permits inbound packets, but an outbound filter denies the same type of packet to be sent. This results in the packet not being sent.

Q4 Along with an FTP server, a server can be configured as a DHCP server or DNS server to either resolve IP addresses for devices currently on the network, or control IP domain names/addresses to simplify internet address usage.


Worksheet 3




APIPA address failure to find dhcp server




Adding dhcp server tools



Installing DHCP server tools




DHCP scope







Activating scope




Lab assigned 192.168.0.3 using DHCP



Failure to ping due to no ip configuration



Creating reservation from address leases



Capturing physical address of 0 network from windows 7 client: 00-15-5D-86-03-20

Q1. The role of a DHCP server in a network is to automatically provide a host with an IP address along with subnet mask and default gateway info.

Q2. When a NIC is configured with DHCP but cannot communicate with the DHCP server, the computer may be given an APIPA address which spans from 169.254.0.1 to 169.254.255.254, with a subnet mask of 255.255.0.0

Q3.

DHCP client broadcasts a discover packet across the network


DHCP server receives discover packet and broadcasts an offer packet

DHCP client receives offer packet and broadcasts a request packet


DHCP server receives request packet and broadcasts an acknowledgement packet

Q4. Allows DHCP servers from different subnets to assign host info. Allows the use of centralised DHCP servers. Provides layer of redundancy to the DHCP servers.


Worksheet 4




Adding active directory domain services to Cardiff server




Adding nwtraders.msft new forest on Cardiff server



Selecting functional level windows server 2012 R2



Adding Swansea server to nwtraders.msft domain



Adding windows 7 server to nstraders.msft domain

Q1 To resolve names to IP addresses, a client can use the DNS console. The DNS console allows the user to select a domain name and assign it to an IP address.

Q2 A recursive query is when a computer sends a request (in the form of a website address e.g www.google.com) and in return, the DNS server replies with the IP address, or in cases where the domain name is invalid, a "domain not found" error message is displayed.


An Iterative query is when a computer sends a request for a best answer to the query without the aid of any further DNS servers. An iterative query is typically made from one DNS server to another.

Q3 Root hints are records that are stored on a DNS server from previous requests. They are used to resolve names that are not available to the authoritative DNS domains. This way, if a DNS server receives a request that it cannot resolve, it can check the root hints file to see if it has the correct answer to the request.

Q4

1. Server1.subcontA.nwtraders.msft

2.



Q5 DNS zones are governed by zone files. Zone files contain all the information required for a zone to access DNS records. DNS zones can be either Primary DNS, Secondary DNS or a Stub. Primary DNS is a read/write copy of a DNS database where records can be added/removed or edited. A secondary DNS is a read only copy of a DNS database which serves as a backup to the primary DNS. A stub is a small directory of regularly used records.

Worksheet 5

Evaluation of Clustering strategies

Introduction

This appraisal will evaluate the available network clustering strategies that are suitable for hosting a 3-tier web development system. Each individual clustering strategy will be analysed and evaluated along with the accompanying tools and technologies that are utilized by the clustering strategy.

After each available Clustering strategy has been evaluated, the most suitable strategy will be selected to be implemented for the 3-tier web development system.

Suitable Clustering Strategies

1. High performance computing

1.1 Overview

High performance computing involves the use of multiple powerful computers (Referred to as “nodes”) to process large computational problems. A HPC cluster consists of a minimum of 2 nodes, however it is possible for a HPC cluster to contain thousands of nodes. The use of HPC clusters is to allow processing on a massive scale to be allocated among each node in the cluster to allow the process to be completed much faster and more efficiently than on a single computer. [1]

1.2 Purpose

HPC is regularly used in industries which require processing of simulations/operations on a massive scale. HPC is used regularly for weather forecasting where multiple simulations of weather behavior need to be calculated from data recorded from many different cities/countries, with the more simulations resulting in more accurate predictions. [1]

HPC are also strongly used for simulating testing of designs, prototypes, chemical reactions, automation and other similar examples as a means of reducing the time and cost of carrying these tests out using real materials/prototypes. [1]

1.3 Hardware

Multiple powerful computers, with a cluster making use of the processing power of multiple powerful CPUs to process calculations. As of December 2013, the fastest HPCs have been recorded processing operations at around 300,000 times faster than a standard desktop PC. [1]

1.4 Costs

The cost of creating an HPC cluster depends on how big of a cluster your project requires. An HPC cluster is a network of multiple high-performance supercomputers/servers, which will need to be housed in a suitable environment (taking into account cooling, risk factors etc.). [4] The HPC cluster will have to be installed by a specialist team. Cost of installation can add up considerably when you are dealing with potentially 10s, 100s or 1000s of machines. These machines will usually be insured and included in a service plan to ensure uptime and fault resolution. Along with all the costs of hardware required to implement an HPC cluster, many experienced engineers will be required to maintain and control these machines daily, usually salaried workers. [4]

An HPC cluster the size of 100 nodes can cost over \$1.3m for the hardware alone, with monthly running costs including electricity and management staff wages can cost around \$70,000, exceeding \$100,000 when you include support staff. [4]

If a certain business project requires the use of a HPC cluster to speed up the process of some work, for example prototype testing or scientific simulations, it is possible to rent part of an HPC cluster for the time it takes to complete any given task. This allows even a small company to make use of the processing power of an HPC cluster for a fraction of the cost of a full HPC cluster [4]

1.5 Scalability

Depending on the intended task, a HPC cluster can be composed of a minimum of 2 nodes, increasing into the thousands should the situation require such a large number of nodes. [1]

1.6 Reliability

Reliability of HPC clusters comes down to the individual node failing, or a failure in the network architecture connecting the nodes together which allows communication and distribution of workload between nodes. [1]

1.7 Advantages and Disadvantages

Advantages [1,2]

- Allows businesses to reach levels of processing power deemed unachievable in certain business setups
- Extremely reliable due to ease of resolution of most common problems related to the use of a single computer
- Can process massive workloads in a much shorter time than the average desktop PC
- Hugely scalable due to network capabilities and connecting multiple HPCs together
- Increases safety due to multiple simulations of virtual environment scenarios before performing live demonstrations
- Deals with parallel operations extremely well

Disadvantages [1,2,3]

- Network overload if work is not allocated efficiently
- Software licensing can get complicated
- Poorly designed HPC clusters can cause a single point of failure of multiple machines
- Security concerns with multiple businesses using the resources of a single HPC (As seen in the recent Meltdown and Spectre vulnerabilities in Intel CPUs) [3]
- Primarily deals with significantly parallel operations.

2. Network Load Balancing

2.1 Overview

Network load balancing works to distribute traffic across up to 32 nodes in a cluster through the use of the TCP/IP protocol. NLB can dynamically add or remove hosts to deal with fluctuating traffic. Servers that are part of an NLB cluster are referred to as "hosts", each host runs its own copy of the server applications. [5] When a host failure occurs, the work load is immediately distributed to the other hosts that are still active. Failed hosts can re-join the cluster seamlessly to regain its workload share and continue processing. [5]

2.2 Purpose

NLB clusters are commonly used for mission critical operations, as well as hosting web servers. While an NLB cluster is only capable of running stateless applications, it is highly scalable and fits the purpose of these applications extremely well. [5]

2.3 Hardware

Network load balancing is a method of dealing with multiple nodes (32 per cluster, multiple clusters) using network drivers. The low overheads of implementing an NLB cluster are advantageous. In a single cluster, all devices must be set to either multicast or unicast. [5]

2.4 Costs

Due to the software elements, there is no need to purchase any extra hardware to be able to implement NLB. Unless the project is being created from the ground up, where typical network hardware would have to be purchased, NLB can simply be activated through the Microsoft server OS. [5,9]

2.5 Scalability

NLB can support up to 32 nodes in a single cluster, including multiple clusters. NLB can seamlessly add or remove hosts to or from the cluster as work load increases or decreases, usually without failure. [5]

2.6 Reliability

NLB clustering allows seamless transfer of work load to other hosts on the network, as well as reintegration of failed hosts or adding new hosts to deal with increased traffic. NLB makes use of network drivers rather than adding powerful expensive hardware, failures are largely dealt with by efficiently adding or removing hosts to manage the work load across the network. [5]

2.7 Advantages and Disadvantages

Advantages [5]

- Less complex than a Failover Cluster
- Provides reliability and scalability to large complicated applications
- High scalability
- High reliability
- Can balance multiple requests across the cluster

Disadvantages [5]

- Restricted to stateless applications
- High performance requires complex setup

3. Failover

3.1 Overview

A failover cluster is designed in such a way that if any piece of hardware suffers a failure, there is always a backup device ready to take over from whatever job the failed device was performing. In cases which warrant it, backups of backup devices may be implemented to ensure increasing levels of redundancy. [6,7]

3.2 Purpose

Used mainly for services that require reliable, stable uptime and minimal downtime. Services such as high traffic trading websites (eBay, Amazon etc) will lose money if services suffer downtime. Hospitals may lose critical information such as timetables or medication/dosages for patients if medical information servers suffer downtime. [6,7]

3.3 Hardware

The hardware required to set up a Failover cluster depends on what the cluster is to be used for, and how many levels of redundancy are required for the task.

At minimum it is required to have at least 2+ storage systems (Storage Area Networks), 2+ switches and 2+ servers. It is optimal and, in some cases, mandatory to have exact same hardware models and software versions, this ensures a smooth and streamlined transfer between a failed device and the backup device. [6,7]

3.4 Costs

The cost of creating a Failover cluster can be as low as the price of a simple 2-layer Failover cluster, which contains 2 servers, 2 switches and 2 SANs. This can cost up to £30,000 when you factor in licensing costs and the level of equipment required. Each extra layer of redundancy will increase the total cost of a Failover cluster. [8]

3.5 Scalability

Failover clusters are largely scalable, due to the cluster scaling directly with the levels of redundancy required of the cluster. If you want another layer of backups, you simply add another set of servers, switches and storage systems alongside the current ones. The layer of backups required depends on the service you are intending to support with the cluster. [6,7]

For example, if the loss of revenue for a company due to 1-hour downtime is £10,000, and an extra layer of redundancy (the cost of another server, switch and storage) is £500, it would be extremely cost effective to add increasing layers of redundancy.

There are many factors that can influence the level of redundancy required for any given system, but a cost analysis must be carried tailored to each situation to find out exactly what level of redundancy may be required. Cost is one of the more crucial factors but space, risk levels, environment, business needs and importance of a stable service are all very important in the design and implementation of a Failover cluster. [6,7]

3.6 Reliability

The reliability of a Failover Cluster related directly to the levels of redundancy created in the cluster. If a failover cluster consists of only 2 servers, all it takes is for both servers to fail for the entire system to fail, or even 1 server if the neighbouring server cannot accept the workload passed to it from the failed server. The level of redundancy required for any project must be tailored to the potential losses incurred should a failure occur.

As well as the reliability of the core setup of the Failover Cluster, the layers of redundancy can extend further due to the structure of the SAN. The SAN can be configured in RAID 1, which uses mirrored hard disks so that any failure can be recovered from the mirroring disk, should the mirrored drive be in working order.

3.7 Advantages and disadvantages

Advantages [6,7]

- Reduce downtime by transferring workload from failed devices to other devices on the network
- Can potentially have multiple levels of redundancy reducing the chances of a complete failure.
- Can support multiple apps at once, while providing stability to all applications on the server

Disadvantages [6,7]

- Costs can mount up when adding extra layers of redundancy to a Failover Cluster
- Can take a lot of time to set up a large cluster
- Requires multiple specialized staff to maintain

Conclusion

Due to the nature of a 3-tier web development system, it would be best suited to make use of a failover cluster incorporated with Network load balancing technologies to accommodate multiple users on a website. The 3 layers of a tier 3 architecture are presentation layer, logic tier, and data tier.

The presentation tier is important as it is the customer-facing component of the website. If a customer is having problems with a slow or failing website front end, they may decide to take their custom elsewhere. So, it is crucial that maximum uptime is achieved in order to satisfy customer needs. This is where the layers of redundancy of a failover cluster will ensure adequate levels of reliability should any failures occur.

The logic tier can be responsible for processing customer transactions. Similar to the presentation tier, this service is important to give the customer a smooth, problem free checkout on an e-commerce website. If there were to be any problems with the ecommerce portion of a website, it could lead to customers being charged more than they expected and can cause problems for both the customer and the business. A failover cluster will reduce the likelihood of such problems happening, while the NLB technology will allow fast processing of customer orders in the event of high traffic during sales and holiday shopping etc.

The third tier in the 3-tier architecture is the Data layer. Without stable, reliable connection to the business data, there is no guarantee that website information is accurate or true. Customers could be seeing incorrect prices or information about assorted products, leading to problems such as false advertising, which can affect a business reputation, as well as potentially suffering fines for false advertising or losing money through accidental sales at incorrect prices.

An HPC cluster would be inappropriate for a web development system as there isn't a big enough workload to warrant spending money on such high-performance hardware.

References

- [1] "COLUMN DATA CENTERS Typical HPC Applications."
- [2] "Pros and Cons of HPC Virtualization - insideHPC." [Online]. Available: <https://insidehpc.com/2014/12/pros-cons-hpc-virtualization/>. [Accessed: 08-Feb-2018].
- [3] "The CPU catastrophe will hit hardest in the cloud - The Verge." [Online]. Available: <https://www.theverge.com/2018/1/4/16850120/meltdown-spectre-vulnerability-cloud-aws-google-cpu>. [Accessed: 08-Feb-2018].
- [4] "The Cloud Versus HPC Cluster Cost Conundrum." [Online]. Available: <https://www.nextplatform.com/2015/06/03/the-hpc-cloud-versus-cluster-cost-conundrum/>. [Accessed: 09-Feb-2018].
- [5] "Network Load Balancing | Microsoft Docs." [Online]. Available: <https://docs.microsoft.com/en-us/windows-server/networking/technologies/network-load-balancing>. [Accessed: 09-Feb-2018].
- [6] "Windows Server Failover Clustering with SQL Server | Microsoft Docs." [Online]. Available: <https://docs.microsoft.com/en-us/sql/sql-server/failover-clusters/windows/windows-server-failover-clustering-wsfc-with-sql-server>. [Accessed: 09-Feb-2018].
- [7] "Failover Clustering | Microsoft Docs." [Online]. Available: <https://docs.microsoft.com/en-us/windows-server/failover-clustering/failover-clustering-overview>. [Accessed: 10-Feb-2018].
- [8] "Cost and License considerations between Always On Availability Groups and Always On Basic Availability Groups | StarWind Blog." [Online]. Available: <https://www.starwindsoftware.com/blog/cost-and-license-considerations-between-always-on-availability-groups-and-always-on-basic-availability-groups>. [Accessed: 10-Feb-2018].
- [9] "How To Configure Network Load Balancing Parameters in Windows Server 2003." [Online]. Available: <https://support.microsoft.com/en-us/help/323437/how-to-configure-network-load-balancing-parameters-in-windows-server-2>. [Accessed: 11-Feb-2018].

GRADEMARK REPORT

FINAL GRADE

GENERAL COMMENTS

Instructor

47 /100

PAGE 1

PAGE 2

PAGE 3

PAGE 4

PAGE 5

PAGE 6

PAGE 7

PAGE 8

PAGE 9

PAGE 10

PAGE 11

PAGE 12

PAGE 13

PAGE 14

PAGE 15

PAGE 16

PAGE 17

PAGE 18

PAGE 19

PAGE 20

PAGE 21

PAGE 22

PAGE 23

PAGE 24

PAGE 25

PAGE 26

PAGE 27

PAGE 28

PAGE 29

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

GRADING FORM: SMU GENERIC_GJ

ROSS VANDENEECKHAUTE-JONES

47

RESEARCH-INFORMED LITERATURE

 A reasonable number of sources used, though pretty much exclusively from the web, many being non-authoritative. Source 1 is incomplete - needs either a link, or information about where it was published. Some web sources can be acceptable - e.g. you've used TechNet in a fair few places. Microsoft can be considered authoritative, though not necessarily objective on their technologies - fine to cite, though caution must be taken when talking about anything potentially subjective.

4

KNOWLEDGE AND UNDERSTANDING

 Some knowledge and understanding evident, though the comments throughout the report around screenshots are far too brief. While there doesn't need to be an enormous amount of description per screenshot, it would've been nice to see some overarching descriptions of groups of screenshots, to explain the whys of what's being done. The clustering section, however, is quite good, as are the written tasks in each worksheet.

10

ANALYSIS

 Some evidence of synthesis of material, with a few concluding remarks in the clustering section in particular.

9

PRACTICAL APPLICATION AND DEP.



Worksheets not quite complete, though most of the way there. The work has seemingly been completed to a reasonable standard, though again, this mark is somewhat limited by the lack of attached description.

18

SKILLS FOR PROFESSIONAL PRAC.



You write reasonably well, and this is a fairly well-presented document overall.

6