

L.Collins - DWA Report

by Lauren Paige Collins

FILE	12964_LAUREN_PAIGE_COLLINS_L.COLLINS_-_DWA_REPORT_181263_2038618011.DOCX (10.65M)		
TIME SUBMITTED	01-MAY-2019 03:51PM (UTC+0100)	WORD COUNT	5842
SUBMISSION ID	105990553	CHARACTER COUNT	32885

MSc Applied Computing

Distributed Web Applications

Cluster Concepts

Dr. Tim Bashford

Lauren Paige Collins – P134209

Contents

Table of Figures	3
1.0 Introduction.....	4
2.0 The Build Process for a Failover Cluster	5
2.1 The Storage Area Network.....	5
2.2 The Windows Operating System.....	11
2.3 Networking and Network Layout	18
2.3.1 Network Layout of the Storage Area Network.....	18
2.3.2 Cardiff (Domain Controller) Network.....	21
2.3.3 Swansea (Node 1) Network.....	21
2.3.4 Bridgend (Node 2) Network	22
2.3.5 Lab 01-01 (Client) Network.....	22
2.4 The Failover Cluster Feature.....	23
3.0 Concepts of a Failover Cluster	30
3.1 Rationale behind Failover Clustering	30
3.1.1 Rules of a Failover Cluster	31
3.2 The Quorum Configuration.....	31
3.3 Cluster Modes	32
3.4 Resource Groups	33
3.5 The Heartbeat Network Operation	34
4.0 The Difference between Failover and Network Load Balancing Clusters	35
4.1 Failover Clusters	35
4.2 Network Load Balancing Clusters	36
4.3 Stateful and Stateless Applications	36
5.0 Evaluation of the Implemented Cluster Architecture for Operation and Performance.....	37
6.0 Deployment of Software to such a Cluster in a Cloud-Based Environment	39
6.0.1 IaaS.....	39
6.0.2 PaaS	39
6.0.3 SaaS	39
7.0 References	41
8.0 Bibliography.....	44

Table of Figures

Figure 1 Failover Clustering [22]	30
Figure 2 Resource Group for Database	33
Figure 3 Resource Group for Windows System [9]	33
Figure 4 Failover Diagram [23]	35
Figure 5 Network Load Balancing Diagram [24]	36
Figure 6 IaaS, PaaS & SaaS Diagram [25]	40

1.0 Introduction

Throughout this report, the process for a failover clusters is going to be evaluated from a technical perspective. This will include the storage area network, the operating system and the failover cluster feature. Within the evaluation, the steps that were taken to configure the cluster will be detailed, including appropriate print screens displaying the necessary steps.

Alongside this, independent research will be conducted describing the foundation concepts of a failover cluster. Included in this will be basis and rules of a failover cluster, defining a Quorum and its relevance within this topic, cluster modes and the heartbeat network.


The difference between failover and network load balancing clusters will be discussed, including a detailed evaluation of the cluster architecture for both operation and performance. Finally, the cloud based environment software will be discussed in relation to clusters.

2.0 The Build Process for a Failover Cluster


To begin the build process of the failover cluster, a virtual storage area network was required. Virtual Machine Manager was used throughout this process.

2.1 The Storage Area Network


The Virtual Machines that were required, needed to be created in Virtual Machine Manager.




To start this process, the user chose the option in the top tool bar named "Create Virtual Machine". This option then displayed the below window enabling the user to click on the option, "Create the new virtual machine with a blank virtual disk".




The next stage in this process, was to modify the hardware and settings for the virtual machine.




Within the compatibility option, the user ensured that the Hyper-V option was selected.




Next was the General section. Four processors were required at this stage which was inputted into the settings by the user.




Alongside the processor modification, the memory of the machine was then set to 4GB.




Once the General settings had been applied, the user was then required to move the Advanced section to click on Availability. The option “Make this virtual machine highly available” was then selected.




Once the hardware had been configured to the appropriate settings, it was then required to select the destination the machine would be deployed to. The “Deploy the virtual machine to a private cloud” was selected.




Selecting the cloud was the next stage. As a result of the user being part of Group 2, the user ensured that the correct group was available to deploy the machine.




Within properties, the automatic actions of the machine were selected stating for the machine to not turn on automatically and ensure that if the server was stop operating, that it saved correctly.



Finally, the summary page is presented providing a breakdown of the machine that has been created and its destination. This print screens show the process of creating Cardiff. Swansea was created following the same process that has been explained for Cardiff machine.




The user then clicked create, creating the machine for Cardiff. In the below print screen, the two virtual machines can be viewed within the SAN cloud.




2.2 The Windows Operating System

Once the servers had been created correctly, the user was then required to install the Windows operating system on to the two machines.


The first step was to log on to the server via the user's virtual machine, using the user name SOAC\S134209 and password.




Once logged in, the Hyper-V screen was presented stating that the user was required to select a key on the keyboard to initiate the install for the operating system.




Below is the next window that was displayed to the user. This is where the user was required to state what language the operating system was to be installed in, the time and currency format and the keyboard input method. It was ensured that English was chosen along with the United Kingdom out of the provided options.




Once the user had clicked next, the below screen was then presented for the user to begin installing Windows Server 2016.




The next step was to insert the product activation key.




Once the product key had been inserted, the user then needed to choose which operating system was to be installed on the machine. Windows Server 2016 (Desktop Experience) was selected from the list.




The user was then required to accept the license terms of installing the product.




Next was the type of installation. The user chose Custom: Install windows only (advanced), as instructed.




Finally, the next two print screens show the user selecting where the Windows Operating system was to be installed.




The primary drive was chosen as it had the largest amount of memory available to install the operating system.




The installation process of the Windows Operating system than began.




Once installed, the below print screen was displayed showing the operating system loading up.



The Administrator account was then created by giving the account a username and a password.



Windows had then successfully been installed on the machine, allowing the user to log in with the previous account details created.



This process shown was conducted on the virtual machine named, Swansea. The same process was then carried out on the Cardiff virtual machine.

2.3 Networking and Network Layout


Within this process, there has been a variety of different networks established which has allowed the SAN and the machines to operate appropriately. This section is going to recognise and explain these networks.

2.3.1 Network Layout of the Storage Area Network

Virtual Machine Manager has been used as the Storage Area Network (SAN) for the virtual machines. On the virtual machine there were three private networks created, these were:-

- ❖ PrivateCloud2Network1
- ❖ PrivateCloud2Network2
- ❖ PrivateCloud2Network3


On each virtual machine, two network adapters were established within the hardware configuration properties, which both connected to one of the three networks. In the print screen below, it can be viewed which machines were connected to which network.




Name	Type	Logical Network	VM network	Resource
E1 - Bridge2	VM	PrivateCloud2Network2	PrivateCloud2Network2	
E1 - Bridge2	vNIC	PrivateCloud2Network1	PrivateCloud2Network1	
E1 - CardB1	VM	PrivateCloud2Network1	PrivateCloud2Network1	
E1 - CardB1	vNIC	PrivateCloud2Network2	PrivateCloud2Network2	
E1 - Lab 01-01	VM	PrivateCloud2Network1	PrivateCloud2Network1	
E1 - Lab 01-01	vNIC	PrivateCloud2Network2	PrivateCloud2Network2	
E1 - Lab 01-01	VM	PrivateCloud2Network2	PrivateCloud2Network2	
E1 - Lab 01-01	vNIC	PrivateCloud2Network1	PrivateCloud2Network1	
E1 - Searcase	VM	PrivateCloud2Network2	PrivateCloud2Network2	
E1 - Searcase	vNIC	PrivateCloud2Network1	PrivateCloud2Network1	

Allocating these networks through the network adapters allowed the machines to connect and communicate over the provided network and with one another. Using these private cloud networks also ensured that the users of the created machines could not directly access the internet which ensured no errors would occur.

Originally, when the hardware was configured on each of the servers, it was ensured that both Cardiff and Swansea had one network adapter. Once the servers had successfully loaded up, the network connections were then modified. The below print screen was taken on the Swansea server.




Right clicking on the network allowed the user to open the Ethernet properties to access the Internet Protocol Version 4 (TCP/IPv4). Doing this enabled the user to input the IP address, the subnet mask and the default gateway that was required. The same process was carried out on the Cardiff network. Once completed the Windows firewall was then turned off on the servers, allowing the servers to communicate.



Once this connection had been established, the Ethernet connection was then named the 0 Network on both Cardiff and Swansea. After this stage the second network adapter was then added to the two servers, both set up to connect to the VM network PrivateCloud2Network2, due to the first network adapter being connected through VM network PrivateCloud2Network1.

The second Ethernet connection on the network connections for both Cardiff and Swansea was then named 11 Network. The same previous process was then conducted through the network properties, setting the network IP address for Swansea to 192.168.11.2 and Cardiff to 192.168.11.1.




The connection was then tested through the same Ping communication in the Command Prompt to ensure that the servers could communicate on the 11 network also.

Further on in the process, when the Lab 01-01 server was added, the network connection was set up with the same process that was conducted on the Swansea and Cardiff machines. Lab 01-01 was set up to connect through Cardiff. Swansea and Lab 01-01 could only communicate through the Cardiff Network. Bridgend was then created as a clone from the Swansea Server. All of the network connection details for the two network connections, on each of the four machines can be viewed in sections 2.3.2 – 2.3.5.

The two network connections that were set up, one was set to be an application network and the other to be a cluster network. “Application clustering typically refers to a strategy of using software to control multiple servers. Clustered servers can help to provide fault-tolerant systems and provide quicker responses and more capable data management for large networks” [1]. The application network is set up to run clustering jobs, which is why a cluster was built further along in the process. This then allows for massive parallel applications to be run, which are the jobs that require number crunching. While the other is in place for cluster controlling, flagging, monitoring the heartbeat network and ensure that the necessary nodes are running correctly.


2.3.2 Cardiff (Domain Controller) Network

In the below print screen, the network connection details for the Cardiff machine can be viewed.




2.3.3 Swansea (Node 1) Network

In the below print screen, the network connection details for the Swansea machine can be viewed.




2.3.4 Bridgend (Node 2) Network

In the below print screen, the network connection details for the Bridgend machine can be viewed. Bridgend was created by cloning the Swansea network.




2.3.5 Lab 01-01 (Client) Network

In the below print screen, the network connection details for the Lab 01-01 machine can be viewed.




2.4 The Failover Cluster Feature


Before the cluster could be created, the Failover cluster feature was required to be installed on the Swansea and Bridgend Machines. Within the server manager and dashboard, the user began by clicking on “Add Roles and Features”.




This then allowed for the features of the failover cluster to be installed on to the two machines.




The below print screen shows the failover cluster feature being installed on Swansea and Bridgend.




Once installed, the user was then required to go to start menu within the machine and click on Administrative Tools. Within this setting, the user then needed to choose Failover Cluster Manager.




By launching Failover Cluster Manager, the operation for the validation of a cluster could then be conducted on the machine. It was only necessary to do this on the Swansea machine. First step seen below is validating the configuration.




During this stage the cluster needed to be named and linked to the servers.



Once named, the user was then able to move on to the testing section. "Run all tests" was selected for this cluster.




Once all the previous had been completed, the validation configuration was then confirmed ready for testing.




The left screenshot shows the 'Failover Cluster Manager' console. The 'Validate a Configuration Wizard' is running, and the 'Summary' tab is selected. The 'Summary' tab shows a list of validation tests and their results. The 'Summary' tab is selected, showing a list of validation tests and their results. The 'Summary' tab is selected, showing a list of validation tests and their results.

The right screenshot shows the 'Failover Cluster Validation Report' window. The report displays the results categorized by group, including 'Network', 'Storage', 'System Configuration', and 'System Configuration'. The report shows the results of the validation tests, including the 'Network' group, 'Storage' group, 'System Configuration' group, and 'System Configuration' group.


As a result of the validation tests being completed successfully, the user was then able to begin the process of creating the cluster.



The same process and steps from the validation, were then taken to start the process of creating the cluster.



Once all the necessary steps and tests had been completed, the cluster was then successfully created on the Swansea machine.



3.0 Concepts of a Failover Cluster

Clustering has become a vital part within the running of tasks on a computer system. It allows for a range of tasks to take place at one time while running the tasks on a multiple machines.

Although technology has come a long way throughout the years, there are still aspects of computers that are not 100% reliable as of yet. As a result of this, and to ensure that if any servers or nodes were to fail, there is strategy in place to aid in resolving this, which is referred to as a failover cluster. "A failover cluster is a group of independent computers that work together to increase the availability and scalability of clustered roles". [2]

Throughout this section, Failover Clustering is going to be analysed in detail, examining the basis, the rules, the Quorum Configuration, Cluster Modes and Resource Groups and also analysing the Heartbeat Network.

3.1 Rationale behind Failover Clustering

The service of failover clustering is provided by Windows Microsoft Server. Within this service, the clustered servers, which are often referred to as nodes, are all linked through appropriate cables and software programs. Having the system arranged in this way, ensures that if a node was to fail, the alternative nodes provide the service requested. The process of this occurring is then known as a failover.

Once this occurs, the service is then monitored to check if it is working correctly. If this is not the case and there is an error, the service will be restarted or moved to another node within the network to resolve the issue. Using this on a system ensures that the user has minimum disruptions.




Figure 1 Failover Clustering [22]

3.1.1 Rules of a Failover Cluster

There are a number of rules which are required to be followed when putting a failover cluster in place. A number of nodes are required, but there needs to be at least two. The hardware required should be identical to ensure the alteration and process runs correctly. There can be some flexibility with this, but it is ideal to be consistent with the hardware choices.

The nodes will need to be within the domain, but ensure that those nodes are not listed as domain controllers. The network adapters connect in groups and if there are multiple adapters, the process of grouping together takes place to avoid any faults occurring and overall develop the network.

Network switches are also required for this process. Two are compulsory but more could be added, and they are typically connected via Ethernet cables. Finally, to ensure successful completion of the failover clusters, storage area network trays are essential as they also need to be duplicated. A storage area would be required for this, e.g. RAID 10 and Disk witness e.g. RAID 1.

3.2 The Quorum Configuration

Windows server failover clustering gives a high accessibility for a variety of capacities. These are only deemed accessible though if the nodes that host the resources are available; but a cluster does require over half the nodes to be running, which is when a quorum is in place. Windows also has a flexible interface for a quorum model, as the configuration can be easily modified within Configure Cluster Quorum Wizard.

Within failover clustering, the quorum configuration defines the amount of failures that the system is able to sustain while the computers remain online. If a number of failures do occur, this could result in the cluster completely shutting down. Often the failures result in nodes no longer operating or no longer able to communicate with one another.

How a quorum works is, is that it is constructed on a voting process. The voting essentials is for it to be part of an active cluster, if it is not, it may not start or run appropriately. If a quorum was not in place, this could result in issues arising in the network and especially if there is e.g. a two-node cluster which are not able to converse. "If there is no Quorum, what prevents both nodes to operate independently and take disks ownership on each side? This situation is called Split-Brain. Quorum exists to avoid Split-Brain and prevents corruption on disks". [3]

There are three quorum configuration selections within Configure Cluster Quorum Wizard. These are as follows [4]: -

- ❖ **Use Typical Settings** – This is when the cluster allocates a vote to the nodes automatically and handle the node votes. If appropriate for the cluster and there is already a cluster storage ready, the cluster is able to choose a disk witness. This particular option is suggested as it decides on the quorum and witness configuration best suited to that particular cluster.
- ❖ **Add or change the quorum witness** – This allows the user to add, modify or delete a witness resource. Also organise files to be shared or a disk witness. The cluster allocates a vote to the nodes automatically and handle the node votes.
- ❖ **Advance quorum configuration and witness selection** – This particular option should only be selected if the obligations for modifying the quorum are application specific or site specific. This allows the user to edit the quorum witness, add, delete and to decide whether the cluster handles votes. The votes are allocated by default and managed.

3.3 Cluster Modes

The properties of cluster can depend on two main factors, if there is more than one network adapter and the amount of traffic originating from outside the cluster. If more than one adapter is present, transfer between cluster modes can be enabled through another network. There are two modes available for clusters, these are: -

❖ Unicast

"This type of information transfer is useful when there is a participation of single sender and single recipient. So, in short you can term it as a one-to-one transmission" [5]. This particular mode is the common choice for transferring data. This particular method is often used in latest versions, as a result of unicast not requiring the cross-network configuration, unlike multicast. As a result of this, potential network errors can then be avoided.

❖ Multicast

"In multicasting, one/more senders and one/more recipients participate in data transfer traffic," [5] often described as a one to many communication. Within the multicast mode, all the servers that are available communicate with servers that are located in the cluster. Throughout the communication process, heartbeats are sent to the servers. However, this can result in "a heavy load on the application's multicast buffer, so if the buffer is full, new multicast messages cannot be written to the buffer and the application is not notified when messages are dropped". [6] To use this method, certain hardware configuration is required, alongside support for packets.

3.4 Resource Groups

Cluster resources are held together within a resource group. “A resource group is a unit of failover. Resources that are related or dependent on each other are grouped using resource groups” [7]. When all resources assemble together within a resource group, this means that the resources have come from the same node. If an operation is performed on a resource group, this then results in the same action being performed on all of the resources that are within the resource group. A single resource or a resource group can also be backed up on request and scheduled backups for the resources can easily be arranged.

One of the main purposes of creating a cluster, is so it can be communicated with as a single resource. This then allows for combining all the necessary resources within the cluster to one group.

“A cluster software script or agent executes operations on resources, including starting, stopping, restarting and monitoring at the resource group level. Resource group operations initiate administrative operations for all the resources within the group” [8]. If there is failure from any of the resources, this results in all of resources failing but once it has been resolved, the resource group will then return to the original server where it was previously operating.

“Think of a resource group as defining what you want to protect and when you want to protect it in terms of day and time. Think of a policy as defining how you want to protect it” [9]. Creating a resource group allows for protecting exactly what the user would like to protect. Alongside this, the user is able attach policies and organise the groups backup settings, e.g. daily, hourly etc.

The below two images demonstrate the connection between the resources, resource groups and the policies for both a database and Windows files systems.




Figure 2 Resource Group for Database [9]




Figure 3 Resource Group for Windows System [9]

3.5 The Heartbeat Network Operation

When a network is in place, it is vital to ensure a high availability solution has been implemented to recover the web servers. Within a high availability solution, at least one redundant system is required but the more redundant systems that are available, ensures a higher chance of recovery if the network was to fail.

Another solution which could be implemented for increasing the availability, is adding a way that will fail over automatically if the primary sever within the network does fail. A failover device requires a way of defining whether the primary server has failed and can no longer operate. However, there is a system which can either be a separate program or included within the functionality of the cluster program available, which can inform whether the system is operating or not. This is referred to as the Heartbeat Network. "Its purpose is to continually poll the servers in a cluster configuration to ensure that they're up and responding" [10]. It continually checks on the server, if it has a response, the backup servers are not required to do anything. On the other hand, if a response is not received, the heartbeat will then be configured to arrange the necessary steps to recover the backup server.

"The heartbeat protocol employs the FIFO (First In, First Out) property of communication channels to ensure that all previously sent messages have been received before the events are ordered. This is achieved by having every node send a message in time intervals" [11]. The heartbeat application requires a method where it is able to communicate with the other computers on the network. The messages received from the nodes are the main source of control information, which ensures no messages are overlooked or postponed within the network.

Ideally the heartbeat should be on a private network and used specifically for the required task. If the heartbeat was to be put on a public network, it could be interrupted or obstructed which could easily cause an unwanted failure on the network. Fortunately to avoid this, the heartbeat can operate on both serial and Ethernet connectivity by crossover cable or a private network.

4.0 The Difference between Failover and Network Load Balancing Clusters

There are two methods that can be applied to ensure a service is available to a user, these methods are known as Load Balancing (attained using Load Balancing Clusters) and Fault Tolerance (attained using Failover Clusters).

The main difference between Failover and Network Load Balancing (NLB) Clusters is that a failover cluster is viewed as a more complicated system. "Load Balancing is all about improvising performance (scale) while Failover Clustering is improvising uptimes mitigating system failures" [12]. Load balancing is often found within web based application servers, which only runs stateless applications. Failover clustering is then found within database servers, which runs stateful requests e.g. Hyper V and DHCP, but can run stateless also. However, "failover clusters are a primary method of achieving high availability at the server level" [13]. Both of these methods have been analysed in the next two sections, 4.1 and 4.2.

4.1 Failover Clusters

Failover does include the system going down, or completely failing and as a result of this, the data does need to be present within all servers. This cluster also constructs two or more servers together but ensures that some servers are available for active nodes, while others are retained for the nodes that are inactive. Typically there is a primary server and a secondary server. The primary sever within the network will take on the average load, handling the requests from the user. The secondary then monitors the primary server. If this were to shut down, the secondary server is then available. The aim of the failover is to have high availability and dependability.




Figure 4 Failover Diagram [23]

4.2 Network Load Balancing Clusters

“The Network Load Balancing (NLB) feature distributes traffic across several servers by using the TCP/IP networking protocol” [14]. Within this cluster, two or more servers are constructed to one another which then allows the load to be balanced. When new clients join, the system will automatically redirect the client to an available server within the cluster, guarantying the balance between the loads remains. Both the loads and the clients can be modified and NLB can also control sending the traffic load to a selected server, which would label it the default. This allows for all of the computers within a cluster to be given the same set of IP addresses, but a set of unique of IP addresses are also kept for the hosts. This clusters main goal is to ensure that the scalability of the service or application is improved but as a result there is a limit on usability of the available applications.




Figure 5 Network Load Balancing Diagram [24]

4.3 Stateful and Stateless Applications


“The stateful and stateless models of software application behaviour define how a user’s web browser communicates with a web server” [15]. As stated in 4.0, Failover runs stateful applications and Network Load Balancing runs stateless applications, but there is a significant difference between the two. “Applications in Stateful react by the current state, while Stateless applications act independently with taking into consideration previous/next request” [16].

If a web server was to store data behind the scenes, the server then uses that data to identify a particular client as continuously connected user, this is then referred to as a stateful service. However, with a stateless service, the server will store the data within a database and use that to confirm the user, when a connection needs to be established. Depending on the website or system, the initiator must ensure that the differences between these two are evaluated, including the advantages and disadvantages. Doing this will ensure that the most suitable model can be chosen which would meet the requirements of the created page or system.


5.0 Evaluation of the Implemented Cluster Architecture for Operation and Performance

Once the cluster has been created on the Swansea machine, the user was then able to test both the operation and the performance of the cluster.


In the below print screen, it shows that the cluster is successfully running on the Swansea machine but is currently being run by the nodes on the Bridgend machine.




The Bridgend machine was powered off by the user, causing a failure event within the network. This has been done to test that the application is working as expected.



Once the machine had successfully been powered off, the user then monitored the cluster within Failover Cluster Manager. In the below print screen, the cluster can be viewed showing that the status was pending while ownership was transferred back to the Swansea machine.



The user was then able to check the status of the nodes within the cluster manager. As seen below, the Bridgend machine had gone down but the Swansea machine had successfully taken over. After putting this failure event in place, it can be seen that the cluster is operating correctly overall.



6.0 Deployment of Software to such a Cluster in a Cloud-Based Environment

Using cloud computing services allows a user to access a computing source within a virtualised setting which is characterised as a public connection e.g. the internet. Within cloud computing, there are three fundamental service models. These are:-

- ❖ **IaaS (Infrastructure as a Service)** – The hardware is provided and managed by an external provider.
- ❖ **PaaS (Platform as a Service)** – Alongside the hardware, the operating system is also managed.
- ❖ **SaaS (Software as a Service)** – This would involve an application layer being provided and managed, resulting in the user having no control over the above two layers. [17]

6.0.1 IaaS

Infrastructure as a Service is also referred to as cloud infrastructure service, “are self-service models for accessing, monitoring, and managing remote datacentre infrastructures, such as compute (virtualized or bare metal), storage, networking, and networking services (e.g. firewalls)” [18]. This aids a user greatly as IaaS can be purchased based on the usage levels and the user would only pay for what is used. “Each resource is offered as a separate service component, and you only need to rent a particular one for as long as you need it” [19].

This is where there is a fundamental difference between IaaS, PaaS, and SaaS as the users of IaaS are liable for the management of the programs, data, runtime, the operating system etc.

6.0.2 PaaS

Platform as a Service which is also cloud platform services, “are used for applications, and other development, while providing cloud components to software. What developers gain with PaaS is a framework they can build upon to develop or customize applications” [18]. Using PaaS allows the progress, testing and distribution of applications to be quick, and profitable. Exactly like IaaS, PaaS can manage the operating system, servers, storage etc. “but also middleware, development tools, business intelligence (BI) services, database management systems and more” [20]. The way that PaaS is designed is so that it allows and supports the full web application sequence, which includes structure, testing, set up, handling and modernising.

6.0.3 SaaS

Cloud application services, known as SaaS, represent the major increasing cloud market. Software as a service uses the internet to provide a variety of applications which are maintained by a third party vendor while the interface can be opened and

used via the client. “With SaaS, it’s easy for enterprises to streamline their maintenance and support, because everything can be managed by vendors: applications, runtime, data, middleware, Oses, virtualization, servers, storage and networking” [18]. Most of these applications can be used via a web browser which eliminates the downloading and installation running on individual computers. “Common examples are email, calendaring and office tools (such as Microsoft Office 365). SaaS provides a complete software solution that you purchase on a pay-as-you-go basis from a cloud service provider” [21]. Some organisations which are not often thought of as software providers, are now building SaaS as a an extra source of income, mainly as a competitive advantage.



Figure 6 IaaS, PaaS & SaaS Diagram [25]

Within the cloud infrastructure of IaaS, allows the organisation and administration team to have highest level of control, relating to the hardware and software. As a result of this the organisation will be responsible for the security and maintaining of the technology. IaaS is also the most expensive choice.

While PaaS also allows the administration team to have control over the platform software, while the applications are being build. PaaS overall allows great levels of flexibility and easier operation. However, if an issue was to occur with the hardware or OS (Operating System), this will also extract the software.

SaaS is mainly praised for its easy set up and use. As all required applications in one, no local hosting is required but this is also a weakness. Due to the user having no control over cloud infrastructure, if a software error was to occur with the provider, the user would also be affected by this, influencing the organisation elements that rely on SaaS applications. SaaS is currently labelled as the cheapest option.

All three of these models have both advantages and disadvantages, but the choice between the three will depend on the organisation, the type of work that is conducted and what cloud service would be most suitable. The organisation will need to identify its overall business needs and exactly what aspects the organisation require control of, within the IT stack.

7.0 References

- [1] "Application Clustering," 2019. [Online]. Available: <https://www.techopedia.com/definition/16181/application-clustering>. [Accessed 25 April 2019].
- [2] Microsoft, "Failover Clustering in Windows Server," 2019. [Online]. Available: <https://docs.microsoft.com/en-us/windows-server/failover-clustering/failover-clustering-overview>. [Accessed 04 March 2019].
- [3] R. Serre, "Understand Failover Cluster Quorum," 17 November 2015. [Online]. Available: <https://www.tech-coffee.net/understand-failover-cluster-quorum/>. [Accessed 04 March 2019].
- [4] W. I. P. Center, "Configure and manage quorum," 18 January 2019. [Online]. Available: <https://docs.microsoft.com/en-us/windows-server/failover-clustering/manage-cluster-quorum>. [Accessed 27 March 2019].
- [5] "Computer Network | Difference between Unicast, Broadcast and Multicast," [Online]. Available: <https://www.geeksforgeeks.org/computer-network-difference-unicast-broadcast-multicast/>. [Accessed 02 April 2019].
- [6] Divya, "Comparison Between Cluster Multicast Messaging and Unicast Messaging Mode," 17 December 2010. [Online]. Available: <http://weblogic-wonders.com/weblogic/2010/12/17/comparison-between-cluster-multicast-messaging-and-unicast-messaging-mode/>. [Accessed 17 April 2019].
- [7] "Types of Server Cluster Resources," 2019. [Online]. Available: <http://www.omniseccu.com/windows-2003/clustering/server-cluster-resources-and-resource-groups.php>. [Accessed 15 April 2019].
- [8] M. A. e. al., Oracle9i RAC: Oracle Real Application Clusters Configuration and Internals, Rampant TechPress, 2003.
- [9] NetApp, "Resources, resource groups, and policies," May 2018. [Online]. Available: <http://docs.netapp.com/ocsc-40/index.jsp?topic=%2Fcom.netapp.doc.ocsc-con%2FGUID-617B849D-3F78-4829-BC98-1D1D508A9AFB.html>. [Accessed 15 April 2019].
- [10] C. Brookman, Linux Clustering: Building and Maintaining Linux Clusters, 2003: New Riders.
- [11] S. Nikolettseas, J. D. P. Rolim, "Heartbeat," in *Theoretical Aspects of Distributed Computing in Sensor Networks*, Springer, 2011, p. 304.
- [12] "Load Balancing vs. Failover Clustering," 20 November 2009. [Online]. Available: <https://nirajrules.wordpress.com/2009/11/20/load-balancing-vs-failover-clustering/>. [Accessed 02 April 2019].

- [13] D. Gibson, MCITP Guide to Microsoft Windows Server 2008, Enterprise Administration, Course Technology, 2011.
- [14] W. I. P. Centre, "Network Load Balancing," 13 September 2018. [Online]. Available: <https://docs.microsoft.com/en-us/windows-server/networking/technologies/network-load-balancing>. [Accessed 11 April 2019].
- [15] J. Papiewski, "What is Stateless and Stateful Models of Web Development," 18 December 2015. [Online]. Available: <https://www.atlantic.net/managed-server-hosting/what-is-stateless-stateful-models-web-development/>. [Accessed 15 April 2019].
- [16] "Stateful and Stateless Applications Best Practices and Advantages," 11 December 2018. [Online]. Available: <https://www.xenonstack.com/insights/what-are-stateful-and-stateless-applications/>. [Accessed 15 April 2019].
- [17] "What is IaaS?," 2019. [Online]. Available: <https://www.interoute.com/what-iaas>. [Accessed 23 April 2019].
- [18] "IaaS, PaaS, SaaS (Explained and Compared)," 2019. [Online]. Available: <https://apprenda.com/library/paas/iaas-paas-saas-explained-compared/>. [Accessed 23 April 2019].
- [19] "What is IaaS?," 2019. [Online]. Available: <https://azure.microsoft.com/en-gb/overview/what-is-iaas/>. [Accessed 23 April 2019].
- [20] Microsoft, "What is PaaS?," 2019. [Online]. Available: <https://azure.microsoft.com/en-gb/overview/what-is-paas/>. [Accessed 23 April 2019].
- [21] Microsoft, "What is SaaS?," 2019. [Online]. Available: <https://azure.microsoft.com/en-gb/overview/what-is-saas/>. [Accessed 23 April 2019].
- [22] "Failover Clustering (I)," [Online]. Available: <https://networksandservers.blogspot.com/2011/04/failover-clustering-i.html>. [Accessed 04 March 2019].
- [23] "Implementing Microsoft Network Load Balancing in a Virtualized Environment," 16 May 2012. [Online]. Available: <http://www.electricmonk.org.uk/2012/05/16/implementing-microsoft-network-load-balancing-in-a-virtualized-environment/>. [Accessed 15 April 2019].
- [24] "Load balancing and Failover systems," 07 December 2013. [Online]. Available: <https://www.tonido.com/blog/index.php/2013/12/07/load-balancing-and-failover-systems/#.XMmtcEhKiUI>. [Accessed 15 April 2019].

[25] "What is the difference between IaaS, SaaS, and PaaS," 02 April 2016. [Online]. Available: <https://www.quora.com/What-is-the-difference-between-IaaS-SaaS-and-PaaS>. [Accessed 26 April 2019].

8.0 Bibliography

1. <https://searchwindowsserver.techtarget.com/definition/failover-cluster>
[Accessed 04/03/19]
2. <https://erg.abdn.ac.uk/users/gorry/eg3567/intro-pages/uni-b-mcast.html>
[Accessed 02/04/19]
3. <https://standardwisdom.com/softwarejournal/2009/09/clustering-vs-load-balancing-what-is-the-difference/>
[Accessed 15/04/19]
4. <https://nordicapis.com/defining-stateful-vs-stateless-web-services/>
[Accessed 15/04/19]
5. <https://www.wirehive.com/thoughts/understanding-cloud-services-stateful-vs-stateless/>
[Accessed 15/04/19]
6. <https://www.quora.com/What-is-the-difference-between-iaaS-SaaS-and-PaaS>
[Accessed 26/04/19]
7. https://subscription.packtpub.com/book/virtualization_and_cloud/9781783281633/1/ch01lv1sec11/cloud-and-clustering
[Accessed 26/04/19]
8. <https://searchcloudcomputing.techtarget.com/feature/How-cloud-computing-kills-clustering>
[Accessed 26/04/19]
9. <https://blog.hubspot.com/service/iaas-paas-saas>
[Accessed 26/04/19]
10. <https://www.datamation.com/cloud-computing/iaas-vs-paas-vs-saas-which-should-you-choose.html>
[Accessed 26/04/19]

L.Collins - DWA Report

GRADEMARK REPORT

FINAL GRADE

61 / 100

GENERAL COMMENTS

Instructor

Engagement with Literature Skills (50): 26

Knowledge and Understanding Skills (10): 6

Cognitive and Intellectual Skills (15): 9

Practical Application Skills (15): 13

Transferable Skills for Life and

Professional Practice (10): 7

Your report is too quote-oriented. You should only ever use quotes if preserving the original wording of something is crucial - and even then, only sparingly.

Far too many of your sources are from unauthoritative websites. Quora, for example, is far worse than Wikipedia! Try to rely on journals first and foremost. Have a chat with me about these issues at some point to ensure they aren't repeated in your dissertation.

Practical work is good - very nicely documented.

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