BAŞKENT UNIVERSITY
INSTITUTE OF SCIENCE AND ENGINEERING

SERVER VIRTUALIZATION FOR ENTERPRISES AND AN APPLICATION

ONUR MİLLİ

MASTER OF SCIENCE THESIS
2016

SERVER VIRTUALIZATION FOR ENTERPRISES AND AN APPLICATION

KURUM VE KURULUŞLARDA SUNUCU SANALLAŞTIRMA VE BİR UYGULAMA

ONUR MİLLİ

Thesis Submitted
in Partial Fulfillment of the Requirements
For the Degree of Master of Science
in Department of Computer Engineering
at Başkent University

2016

This thesis, titled: “SERVER VIRTUALIZATION FOR ENTERPRISES AND AN APPLICATION”, has been approved in partial fulfillment of the requirements for the degree of **MASTER OF SCIENCE IN COMPUTER ENGINEERING**, by our jury, on 06/05/2016.

Chairman : Prof. Dr. İbrahim AKMAN

Member (Supervisor) : Prof. Dr. A. Ziya AKTAŞ

Member : Prof. Dr. Hasan OĞUL

APPROVAL

…./05/2016

Prof. Dr. Emin AKATA

Institute of Science and Engineering

BAŞKENT ÜNİVERSİTESİ
FEN BİLİMLERİ ENSTİTÜSÜ
YÜKSEK LİSANS TEZ ÇALIŞMASI ORİJİNALLİK RAPORU

Tarih : 16/05/2016

Öğrencinin Adı, Soyadı	:	Onur MİLLİ
Öğrencinin Numarası	:	21420210
Anabilim Dalı	:	Bilgisayar Mühendisliği
Programı	:	Tezli Yüksek Lisans
Danışmanın Unvanı/Adı, Soyadı	:	Prof. Dr. A. Ziya AKTAŞ
Tez Başlığı	:	Server Virtualization for Enterprises and An Application

Yukarıda başlığı belirtilen Yüksek Lisans tez çalışmamın; Giriş, Ana Bölümler ve Sonuç Bölümünden oluşan, toplam 77 sayfalık kısmına ilişkin, 16/05/2016 tarihinde şahsım/tez danışmanım tarafından Turnitin adlı intihal tespit programından aşağıda belirtilen filtrelemeler uygulanarak alınmış olan orijinallik raporuna göre, tezimin benzerlik oranı % 16'dır.

Uygulanan filtrelemeler:

1. Kaynakça hariç
2. Alıntılar hariç
3. Beş (5) kelimeden daha az örtüşme içeren metin kısımları hariç

"Başkent Üniversitesi Enstitüleri Tez Çalışması Orijinallik Raporu Alınması ve Kullanılması Usul ve Esaslarını" inceledim ve bu uygulama esaslarında belirtilen azami benzerlik oranlarına tez çalışmamın herhangi bir intihal içermediğini; aksinin tespit edileceği muhtemel durumda doğabilecek her türlü hukuki sorumluluğu kabul ettiğimi ve yukarıda vermiş olduğum bilgilerin doğru olduğunu beyan ederim.

Öğrenci

Onay
16/05/2016
Danışman

ACKNOWLEDGEMENTS

I would like to thank Prof. Dr. A. Ziya AKTAŞ not only for his critical remarks and valuable guidance during the whole period of my research but also for his moral support, which helped me a lot during my hard times.

I am also deeply grateful for my family's everlasting support and love. Finally, special thanks to my beloved wife for her boundless patience and support for my long work hours during this study.

ABSTRACT

SERVER VIRTUALIZATION FOR ENTERPRISES AND AN APPLICATION

Onur MİLLİ

Başkent University Institute of Science and Engineering

The Department of Computer Engineering

The popularity of virtualization has increased considerably during the past few decades. One claim that virtualization is not a brand new technology. It is known that the concept of virtualization has its origins in the mainframe days in the late 1960s and early 1970s, when IBM invested a lot of time and effort and, of course money, in developing robust time-sharing solutions. The best way to improve resource utilization, and at the same time to simplify data center management was seen as virtualization during those years. Data centers today use virtualization techniques to make abstraction of the physical hardware, create large aggregated pools of logical resources consisting of CPUs, memory, disks, file storage, applications and networking.

In this thesis, going over available references, a life cycle for server virtualization is proposed. Later, the steps of this life cycle were applied on a fictitious enterprise information system for security reasons. As an extension of the study, cloud computing has been planned and summarized in the Summary and Conclusions chapter.

KEYWORDS: Enterprise Information System, Virtualization, Server Virtualization

Advisor: Prof. A. Ziya AKTAŞ, Başkent University, Department of Computer Engineering.

ÖZ

KURUM VE KURULUŞLARDA SUNUCU SANALLAŞTIRMA VE BİR UYGULAMA

Onur MİLLİ

Başkent Üniversitesi Fen Bilimleri Enstitüsü

Bilgisayar Mühendisliği Anabilim Dalı

Sanallaştırma teknolojilerinin popülerliği son yıllarda önemli ölçüde artmıştır. Sanallaştırma aslında yeni bir teknoloji değildir. Sanallaştırma kavramı 1960'ların sonu ve 1970'lerin başında anabilgisayar (mainframe) olarak bilinen IBM sistemlerinde kaynak paylaşım çözümü ile ortaya çıkmıştır. O yıllarda kaynak kullanımını artırmanın ve aynı zamanda veri merkezinin yönetimini basitleştirmenin en iyi yolu sanallaştırma olarak görülmüştür. Günümüzde ise veri merkezleri, sanallaştırma teknolojisi ile işlemci, bellek, depolama, uygulama ve ağ bileşenlerini fiziksel donanımdan soyutlayarak mantıksal kaynaklardan oluşan büyük bir havuz oluşturabilmektedirler.

Bu tez çalışmasında bir kuruluştaki sunucuların sanallaştırılma uygulaması amaçlanmış ve sanallaştırma için gereken teknik hazırlık çalışmaları anlatılmıştır. Daha sonra, bilgi güvenliği açısından gerçek bir kurum/kuruluş yerine sanal bir kuruluşun bilgi sistemi ele alınıp onun üzerinde, öngörülen sunucu sanallaştırma süreçleri gerçekleştirilmiştir. Tez çalışmasının devamı olarak bulut çözümün de katkısı öngörülmüştür. Bu nedenle bununla ilgili kısa bir ön çalışma da Özet ve Sonuçlar bölümünde Tezin Devamı başlığı altında özetlenmiştir.

ANAHTAR SÖZCÜKLER: Kurumsal Bilgi Sistemi, Sanallaştırma, Sunucu Sanallaştırma

Danışman: Prof. Dr. A. Ziya AKTAŞ, Başkent Üniversitesi, Bilgisayar Mühendisliği Bölümü.

TABLE OF CONTENTS

Page

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS

AD	Active Directory
AMD-NX	AMD No-Execute
AMD-V	AMD Virtualization
BITS	Background Intelligent Transfer Service
CAS	Client Access Server
CAU	Cluster-Aware Updating
CD	Compact Disk
CoE	Center of Excellence
CPU	Central Processing Unit
CSV	Cluster Shared Volume
DBA	Database Administrator
DC	Domain Controller
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name System
DPM	Data Protection Manager
FC	Fiber Channel
FT	Fault Tolerance
FTP	File Transfer Protocol
GPS	Global Positioning System
GPT	GUID Partition Table
HA	High Availability
HTTP	HyperText Transfer Protocol
HTML	HyperText Markup Language
HV	Hyper-V
IAAS	Infrastructure as a Service
IBM	International Business Machines
ICT	Information and Communication Technology
IDC	International Data Corporation
Intel-VT	Intel Virtualization Technology
Intel-XD	Intel Execute Disable
IP	Internet Protocol
IPsec	Internet Protocol Security

iSCSI	Internet Small Computer System
JeOS	Just enough Operating System
LAN	Local Area Network
LM	Live Migration
LUN	Logical Unit Number
MMC	Microsoft Management Console
MVMC	Microsoft Virtual Machine Converter
NFS	Network File System
OS	Operating System
OVA	Open Virtual Appliance/Application
OVF	Open Virtualization Format
P2V	Physical to Virtual
PAAS	Platform as a Service
QA	Quality Assurance
R2	Release 2
RAM	Random Access Memory
ROI	Return on Investment
SAAS	Software as a Service
SAN	Storage Area Network
SAS	Serial Attached SCSI
SLA	Service Level Agreement
SP1	Service Pack 1
SQL	Structured Query Language
TAR	Tape Archive
TCP	Transmission Control Protocol
VHD	Virtual Hard Disk
VLAN	Virtual Local Area Network
VM	Virtual Machine
VMM	Virtual Machine Monitor
VSS	Volume Shadow Copy Service
WAN	Wide Area Network
WSB	Windows Server Backup

1. INTRODUCTION

1.1. Statement of the Problem

During the last 50 years, some vital trends in computing services have resulted. The most notably, internet spanned the last few decades and continues [Portnoy, 2012].

In order to meet the needs of computing application services, nowadays one must use servers. Generally, servers are placed in server rooms or datacenters in enterprises. These server systems must support application services by running every minute of every day, even weeks. Yet, a recent report states that servers regularly have a CPU (Central Processing Unit) use rate of below 10% [Silva, 2013]. It means that servers are not used efficiently. In addition to that, all servers usually have a normal warranty period, after the end of warranty period, hardware maintenance costs increase and it is preferred to replace existing hardware with a new one, instead of maintaining existing server, since it becomes costly to continue maintenance. Thus, after buying a new server and transferring old application to new and more powerful server, the new server system becomes underutilized.

Within the past ten years, need of new applications has increased. Since deploying new applications to existing servers is not usually possible, new server(s) must be purchased. Increased need in the quantity of applications results in the increase need in the quantity of servers. This results an unfortunate increase in installation, operation and maintenance costs.

Server virtualization is one of the best ways of efficient use in server resources and it is a broad subject covering most parts of ICT (Information and Communication Technology). In this thesis, therefore, only the technical steps of an application of server virtualization for a fictitious enterprise information system is examined and implemented.

1.2. Previous Work

Trends in virtualization are continuously dynamic since the beginning. When the server virtualization technology matures, there are many choices available to system administrators and a lot of price saving server virtualization projects are completed. There are many articles, books and other publications about server virtualization.

Bailey [2009] stated that server virtualization technologies are essential for datacenters. According to Bailey's IDC (International Data Corporation) report in 2009, roughly more than %50 of all server workloads are performed in a virtual machine.

Portnoy [2012] gives a broad overview of server virtualization. He introduces the basic concepts of computer virtualization beginning with mainframes and continues with the computing trends that have led to current technologies. He focuses on hypervisors and compares hypervisor vendors in the marketplace. He describes what a virtual machine is composed of and explains how it interacts with the hypervisor that supports its existence, and provides an overview of managing virtual machine resources. He also shows how to convert physical servers into virtual machines and provides a walkthrough of process.

According to Dell and Intel [2010], virtualization is the key element of computing for today and future. One may state that virtualization performs better utilization in many areas such as hardware resources, real estate, energy and human resources.

As noted by VMWare [2006], virtualization provides opportunities to reduce complexity, improve service levels in many areas. It also lowers capital and operating costs to produce and maintain better ICT infrastructure. During the last 5 years many organizations report 70-80% cost savings, and 3-6-month ROI (Return on Investment) periods while achieving surprising gains in operational flexibility, efficiency, and agility using virtualization.

Siebert [2009] states that virtualization is not a brand new technology, as noted earlier, however it has gained popularity in recent years and is employed to a point in the majority of datacenters. In addition to that, he says that an oversized quantity of vendors have written virtualization-specific applications or changed their applications and hardware to support virtual environments because they acknowledge that virtualization is here to remain.

Bittman et al. [2010] rightly point out that server virtualization for x86 architecture servers is one of the most popular trends in ICT nowadays, and will stay thus for many coming years. Competition is maturing, and therefore the market incorporates a variety of viable selections.

Bittman et al. [2015] make the point that about 75% of x86 server workloads are virtualized. Virtualization technologies are getting more powerful, supporting a lot of workloads and agile development. Price, cloud ability and specific use cases are driving enterprises to deploy multiple virtualization technologies.

In this thesis, going over available references, a life cycle for server virtualization is proposed in Chapter 5.

1.3. Objective of the Study

Objective of this study is to understand and apply virtualization as a tool to improve the use of hardware resources to reduce management and resource costs; improve business flexibility, improve security and reduce downtime. The study proposes a life cycle for server virtualization after studying the available references. Considering the wide spectrum of the topic, not every technical details of the subject will be given. Instead, a higher level of view, that is a logical view, is aimed in the study. Hence, it is expected that enterprises that are considering server virtualization technology can benefit from the result of this study as a first step of guidance.

1.4. Organization of the Study

The thesis starts with a brief Introduction. Second Chapter is about Virtualization Essentials. In that chapter, basic concepts of Virtualization are given. Next, a fictitious enterprise information system for a server virtualization application is defined. In the next chapter, namely Chapter 4, existing life cycles for server virtualization are examined. Chapter 5 is devoted to a proposed life cycle for server virtualization in an enterprise. The next chapter, Chapter 6, has an application of server virtualization for an enterprise. Enterprise is defined as a fictitious organization for security reasons. The last chapter is a summary and conclusions of the thesis including cloud computing and virtualization interaction as an extension of the study. Finally, appendices and references are placed at the end of thesis.

2. VIRTUALIZATION ESSENTIALS

Virtualization is one of the key services in the field of computing [Portnoy, 2012].

2.1. What is virtualization?

As a summary of developments in computing during the last fifty years, one notes the followings: mainframe computers dominated the field in the fifties and sixties. Desktop computers, client/server architecture and distributed systems had marked the following two decades. Starting at the end of nineties and still growing is the Internet, which is reshaping the whole world. 'Virtualization' became the name of another model-changing trend in that development process.

Virtualization, simply, is the isolation of some physical part into a logical object in computing. If an object is virtualized, one will then be able to acquire longer utility life from the resource that item offers. The major objective of this thesis is server virtualization.

"IBM (International Business Machines) mainframes in the 1960s are usually cited as the first mainstream virtualization application" [Popek and Goldberg, 1974]. Thus, a VM (Virtual Machine) was able to virtualize all of the hardware resources such as processors, memory, storage, and network connectivity of a mainframe computer of that time in a VM (Virtual Machine). The VMM (Virtual Machine Monitor), named the hypervisor, is the software that provides the environment where the VMs would run. Figure 2.1 is an example of a basic hypervisor [Portnoy, 2012].

A VMM must provide the three characteristics [Popek and Goldberg, 1974]:

Convenience - The VM is copy of the original (hardware) physical machine.
Abstraction or **Security** - The VMM should have full administration of the computer resources.
Performance - There should be very little or no distinction at all in performance between a VM and its physical identical.

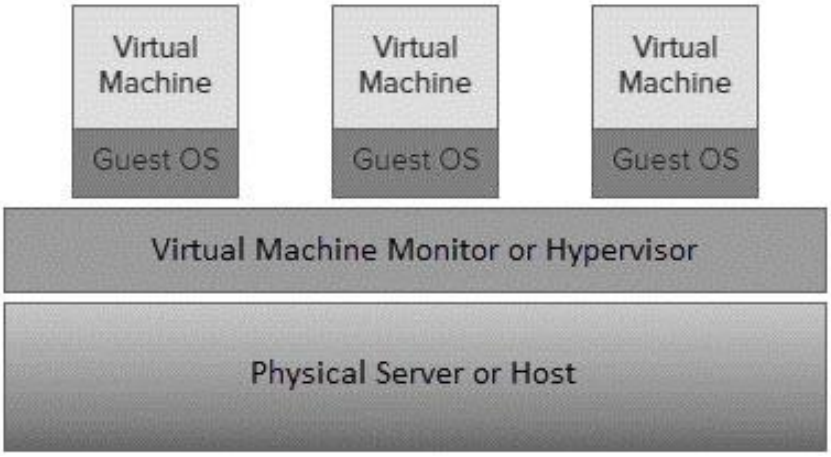

Figure 2.1 A basic hypervisor

In order to understand the reason behind virtualization, one may recall the developments in computing after 1980s: Microsoft Windows was released throughout the nineteen eighties especially as a private PC package. Windows's main purpose was to create a personal operating system; one application on one Windows Server ran fine, but often resulted with failures when more than one application was run. It thus, resulted "one server, one application" as the best practice. Another fact was that various organizations among one company did not want any common infrastructure because of information sensitivity and company policy.

Organizations had to use too many servers in "server farms" or "datacenters" because of dependency of technologies.

On top of all these, internet got into the picture at the end of 90s. In addition to "e-business or out of business" cries, newer devices like mobile phones were added to data center growth. Not solely that, however, total servers were consuming 2 percent of the whole electricity created within the USA and they consumed a similar quantity of energy to cool them. Data centers were reaching their physical limitations. Across a data center, it had been common an average 10-15 percent utilization. Massive bulk of servers was only 5 percent used. That means, most CPUs were unused for 95 percent or more.

2.2. The value of virtualization

Until recent years, there has been a huge increase in the quantity of servers in data centers. After some time, however, those servers' work decreased. The needed help came through the variety of virtualizations. As noted earlier, virtualization technology was used before. It was first used on IBM mainframes back in the early seventies. It was later renewed for contemporary PC systems. As noted by Popek and Goldberg [1974], "Virtualization permits several operating systems to execute on similar hardware at the same time, whereas saving every virtual machine functionally abstracted from other virtual machines".

VMware released a virtualization solution for x86 computers in 2001. Xen arrived two years later as an open source virtualization solution. These hypervisor solutions were placed between an operating system and the virtual machines (VMs) or put directly onto the hardware similar to a native operating system like Windows or Linux.

Consolidating physical servers into one server that might execute several virtual machines, permitting that physical server to run at a higher rate of utilization is the advantage of virtualization. It is *consolidation*. For instance, a server in Figure 2.2 has 5 VMs running on it, has a *consolidation ratio* of 5:1 [Portnoy, 2012].

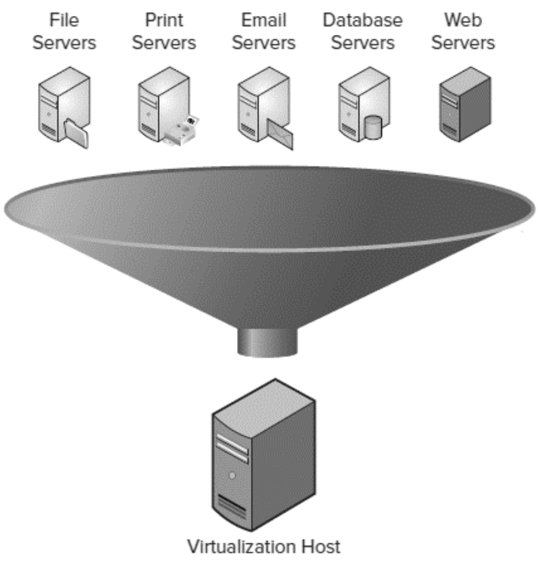

Figure 2.2 Server consolidation

In larger data centers, a whole bunch or maybe thousands of servers exist. Thus, virtualization provides some way to recover an outsized space of servers. Data center space use, cooling and energy costs decreased via virtualization. In addition to that, need of building more data centers has decreased. Furthermore, hardware maintenance costs and system administrators' time for performing different tasks decreased by using fewer servers. In addition all to that, total ownership cost of servers that consist of maintenance costs, power, cooling, cables, personnel costs, and others will decrease. Thus, “consolidation drives down costs” [Portnoy, 2012].

As enterprises began to envision the advantage of virtualization, they do not purchase new hardware once their leases were over, or if they had the equipment, they did not extend their hardware maintenance licenses. Instead, they prefer to virtualize those server workloads; this process is *containment*. Containment is very useful for enterprises in many ways. Figure 2.3 shows an example of server containment.

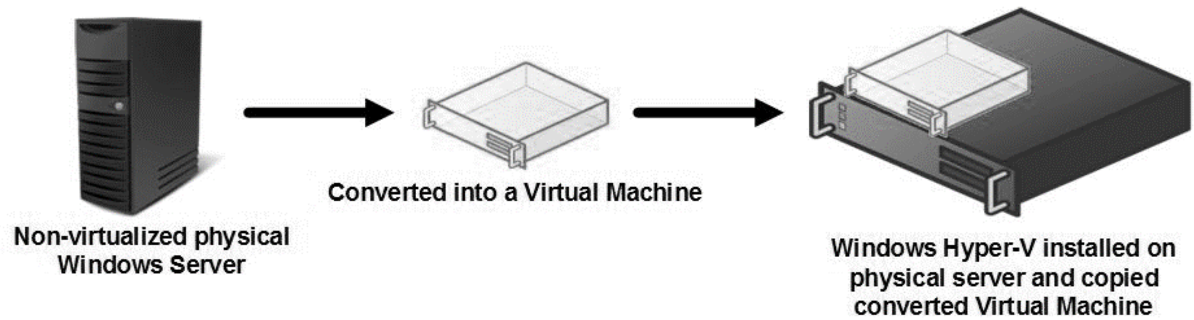

Figure 2.3 Server containment

According to Bailey [2009], more virtual servers were deployed in 2009, than physical servers. Server virtualization enhances an ancient server by consolidation and containment.

Virtual servers are simply a group of files. Convenience has become vital such as maintaining 7/24 operations through increased package and options, or disaster recovery capabilities because of evolution of internet. Virtual machines may be migrated from one physical host to a different one, without any interruption in service.

“Moving to a wholly virtualized infrastructure provides enterprises a far bigger degree of accessibility, agility, flexibility, and administration capability than they may have in some exceedingly entirely physical surroundings” [Portnoy, 2012]. One may add that a big advantage of virtualization is providing the inspiration for subsequent part of knowledge center evolution: cloud computing or solution.

2.3. Virtualization and cloud computing or solution

Roughly, 5 years ago, only a few individuals would have had any idea about “cloud computing or cloud solution” terms. “Nowadays it might be tough to search out somebody who is engaged within the worldwide business or client markets who has not detected the word cloud computing. Very like the push to the net throughout the first 2000s, several of today’s firms are acting on cloud enablement for his or her offerings” [Portnoy, 2012]. Mirroring their actions throughout the internet explosion, client services also are creating the move to the cloud. For instance, Apple lately offered a service named as iCloud wherever one will be able to store multimedia files such as photos, music, etc. Alternative firms, like IBM, Google, Microsoft, and Amazon are providing similar cloud services. By cloud technology, it is very easy to access and utilize resources.

Referring to Portnoy [2012],”Virtualization is the engine, which will drive cloud computing by turning the data center into a self-managing, highly scalable, highly available, pool of easily expendable resources. Virtualization and, by extension, cloud computing provide larger automation opportunities that cut back management costs and increase a company’s ability to dynamically deploy solutions. By being able to abstract the physical layer far from the actual hardware, cloud computing creates the concept of a virtual data center, a construct that contains everything a physical data center would. This virtual data center, deployed in the cloud, offers resources on an on-demand basis, much like a power company provides electricity. In other words, these new models of computing will dramatically simplify the delivery of new applications and permit organizations to speed up their deployments with scalability[1], resiliency[2], or availability[3]".

2.4. Understanding Virtualization Software Operations

"The hypervisor abstracts the physical layer and presents this abstraction for virtualized servers or virtual machines to use. A hypervisor is installed directly onto a server, without any operating system between it and the physical devices. Virtual machines are then installed. Virtual machine can see and work with a range of hardware resources. The hypervisor becomes the interface between the hardware devices on the physical server and the virtual devices of the virtual machines. The hypervisor presents only some set of the physical resources to every individual virtual machine and handles the particular I/O from VM to physical device and back once more. Hypervisors do more than simply offer a platform for running VMs; they provide increased convenience options and make new and higher ways for provisioning and management" [Portnoy, 2012].

"While hypervisors are the foundations of virtual environments, virtual machines are the engines that run the applications. Virtual machines contain operating systems, applications, network connections, access to storage, and other necessary resources. However, virtual machines are packaged in a set of data files. This packaging makes virtual machines much a lot of versatile and manageable with the normal file properties in a very new method. Virtual machines are copied, upgraded, and even migrated from server to another server, without ever having to disrupt the user applications" [Portnoy, 2012].

2.5. Understanding Hypervisors

In this section, a hypervisor is discussed briefly.

[1] **Scalability** is the degree to which a computer system is able to grow and become more powerful as the number of people using it increases.
[2] **Resiliency** is the ability to become strong or successful again after a difficult situation or event.
[3] **Availability** is the ratio between the time during which the circuit is operational and elapsed time.

2.5.1. Describing a hypervisor

The original virtual machine monitor or manager (VMM) serves to solve a specific problem. The term virtual machine manager has fallen out of favor and is known as hypervisor now. Today's hypervisors allow us to make better use of the ever-faster processors that regularly appear in the commercial market and more efficient use of the larger and denser memory offerings that come along with those newer processors. The hypervisor is a layer of software that resides below the virtual machines and above the hardware. Figure 2.4 illustrates where the hypervisor resides [Portnoy, 2012].

Without a hypervisor, an operating system communicates directly with the hardware beneath it. Disk operations go directly to the disk subsystem, and memory calls are fetched directly from the physical memory. Without a hypervisor, more than one operating system from multiple virtual machines would want simultaneous control of the hardware, which would result in chaos. The hypervisor manages the interactions between each virtual machine and the hardware that the guests all share.

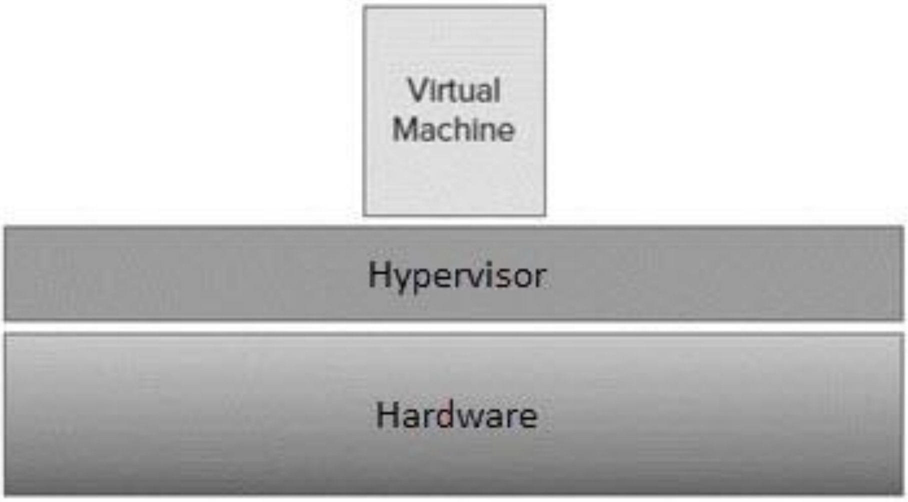

Figure 2.4 Location of the hypervisor

The first virtual machine monitors (VMM) were used for the development and debugging of operating systems because they provided a sandbox for programmers to test rapidly and repeatedly, without using all of the resources of the hardware. Soon they added the ability to run multiple environments concurrently, carving the hardware resources into virtual servers that could each run its own operating system. This model is what evolved into today's hypervisors.

Initially, the problem that the engineers were trying to solve on main frames was one of resource allocation, trying to utilize areas of memory that were not normally accessible to programmers. The code they produced was successful and was dubbed a hypervisor because, at the time, operating systems were called supervisors and this code could supersede them.

Twenty years passed before virtualization made any significant move from the mainframe environment. In the 1990s researchers began investigating the possibility of building a commercially affordable version of a VMM. The structure of a VMM is quite simple. It consists of a layer of software that lives in between the hardware, or host, and the virtual machines that it supports. These virtual machines are called 'guests'. Figure 2.5 is a simple illustration of the Virtual Machine Monitor architecture [Portnoy, 2012].

There are two classes of hypervisors, and their names, Type 1 and Type 2, give no clue at all to their differences. The only item of note between them is how they are deployed, but it is enough of a variance to point out.

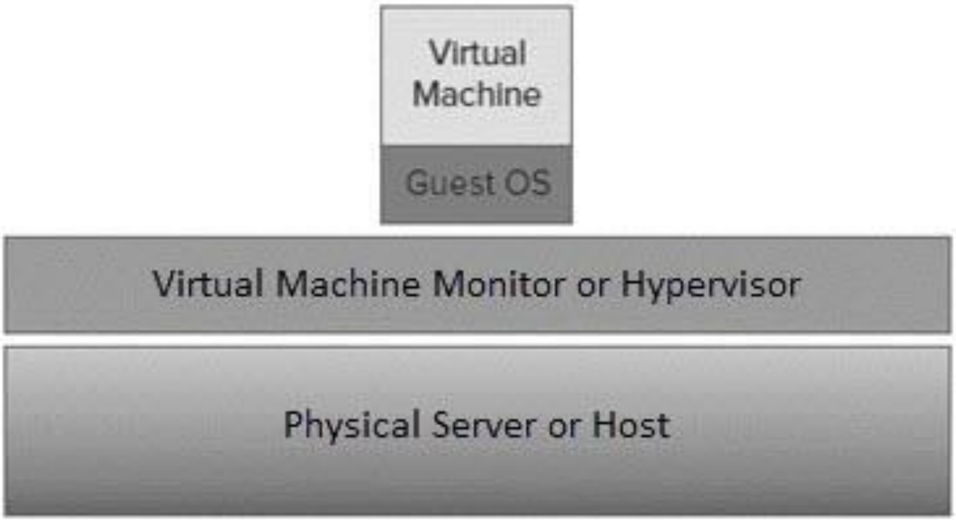

Figure 2.5 A Virtual machine monitor

2.5.2. Understanding type 1 hypervisors

Type 1 hypervisors run directly on the server hardware without an operating system beneath it. Because there is no intervening layer between the hypervisor and the physical hardware, this is also named as a bare-metal implementation. Without an intermediary, the Type 1 hypervisor can directly communicate with the hardware resources in the stack below it, making it much more efficient than the Type 2 hypervisor. Figure 2.6 illustrates a simple architecture of a Type 1 hypervisor [Portnoy, 2012].

Aside from having better performance characteristics, Type 1 hypervisors are also considered being more secure than Type 2 hypervisors.

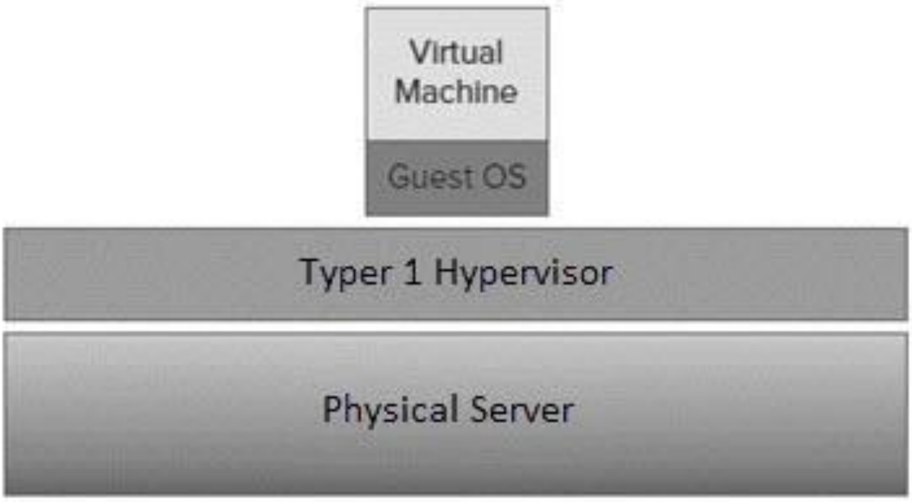

Figure 2.6 A Type 1 hypervisor

Less processing overhead is required for a Type 1 hypervisor. That simply means that more virtual machines can run on each host. From a pure financial standpoint, a Type 1 hypervisor would not require the cost of a host operating system, although from a practical standpoint, the discussion would be much more complex and involve all of the components and facets that comprise a total cost of ownership calculation.

2.5.3. Understanding type 2 hypervisors

A Type 2 hypervisor runs atop a traditional operating system as in Figure 2.7 [Portnoy, 2012].

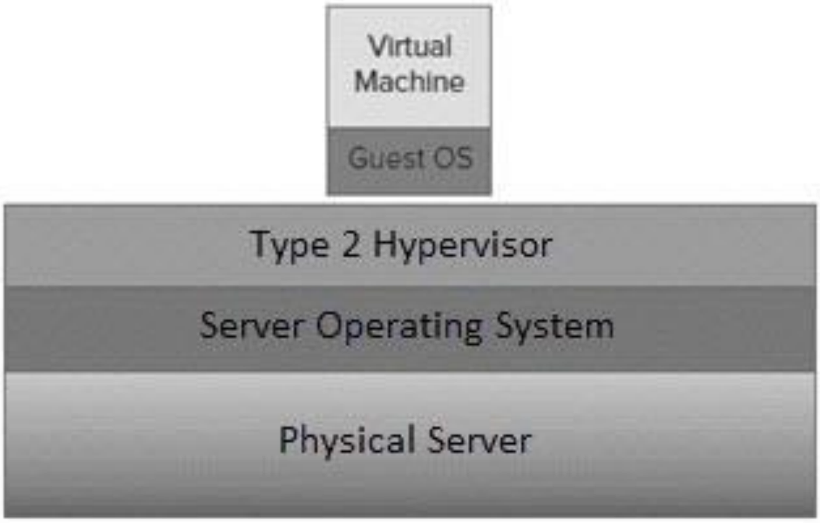

Figure 2.7 A Type 2 hypervisor

This model can support a large range of hardware. Often Type 2 hypervisors are easy to install and deploy because much of the hardware configuration work, such as networking and storage, has already been covered by the operating system.

Type 2 hypervisors are not as efficient as Type 1 hypervisors and also less reliable because of this extra layer between the hypervisor itself and the hardware.

2.5.4. Understanding the role of hypervisor

Hypervisor is simply a layer of software that sits between the hardware and the one or more virtual machines that it supports. It has also a simple job. The three characteristics illustrate these tasks [Popek and Goldberg, 1974]:

- Provide an environment identical to the physical environment;
- Provide that environment with minimal performance cost;
- Retain complete control of the system resources.

Hypervisors are the glue of virtualization. They connect their virtual guests to the physical world, as well as load balance the resources they administer. Their main function is to abstract the physical devices and act as the intermediary on the guests' behalf by managing all I/O between guests and device. There are two main hypervisor implementations: with and without an additional operating system between it and the hardware. Both have their uses. The commercial virtualization market is currently a growth industry, so new and old solution providers have been vying for a significant share of the business. The winners will be well positioned to support the next iteration of data center computing, support content providers for consumer solutions, and become the foundation for cloud computing.

2.6. Understanding Virtual Machines

Virtual machines are the fundamental components of virtualization. They are the containers for traditional operating systems and applications that run on top of a hypervisor on a physical server. Inside a virtual machine, things seem very much like the inside of a physical server but outside, things are quite different.

2.6.1. What is a virtual machine?

A virtual machine (VM), has many of the same characteristics as a physical server. Like an actual server, a VM supports an operating system and is configured with a set of resources to which the applications running on the VM can request access. Unlike a physical server (where only one operating system runs at any one time and few, usually related, applications run), many VMs can run simultaneously inside a single physical server, and these VMs can also run many different operating systems supporting many different applications. Also, unlike a physical server, a VM is in actuality nothing more than a set of files that describes and comprises the virtual server.

The main files that make up a VM are the configuration file and the virtual disk files. The configuration file describes the resources that the VM can utilize: it enumerates the virtual hardware that makes up that particular VM. Figure 2.8 is a simplified illustration of a virtual machine which consists of virtual hardware, operating system(s) and application(s) [Portnoy, 2012]. If one thinks of a virtual machine as an empty server, the configuration file lists which hardware devices would be in that server: CPU, memory, storage, networking, CD drive, etc. In fact, as one will see when we build a new virtual machine, it is exactly like a server just off the factory line—some (virtual) iron waiting for software to give it direction and purpose.

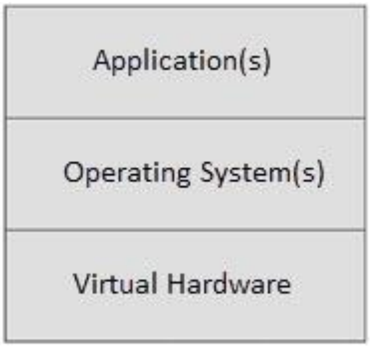

Figure 2.8 A virtual machine

2.6.2. Understanding how a virtual machine works

One way to look at how virtualization works is to say that a hypervisor allows the decoupling of traditional operating systems from the hardware. The hypervisor becomes the transporter and regulator of resources to and from the virtual guests it supports. It achieves this capability by fooling the guest operating system into believing that the hypervisor is actually the hardware.

Before going further, let's examine how an operating system manages hardware in virtualization host. Figure 2.9 will help illustrate this process. When a program needs some data from a file on a disk, it makes a request through a program language command, such as an fgets () in C, which gets passed through to the operating system [Portnoy, 2012]. The operating system has file system information available to it and passes the request on to the correct device manager, which then works with the physical disk I/O controller and storage device to retrieve the proper data. The data comes back through the I/O controller and device driver where the operating system returns the data to the requesting program. Not only are data blocks being requested, but memory block transfers, CPU scheduling, and network resources are requested too. At the same time, other programs are making additional requests and it is up to the operating system to keep all of these connections straight.

2.6.3. Working with virtual machines

Virtual machines exist as two physical entities: the files that make up the configuration of the virtual machines and the instantiation in memory that makes up a running VM once it has been started. In many ways, working with a running virtual machine is very similar to working with an actual physical server. Like a physical server, one can interact with it through some type of network connection to load, manage, and monitor the environment or the various applications that the server supports. Also like a physical server, one can modify the hardware configuration, adding or subtracting capability and capacity, though the methods for doing that and the flexibility for doing that are very different between a physical server and a virtual server.

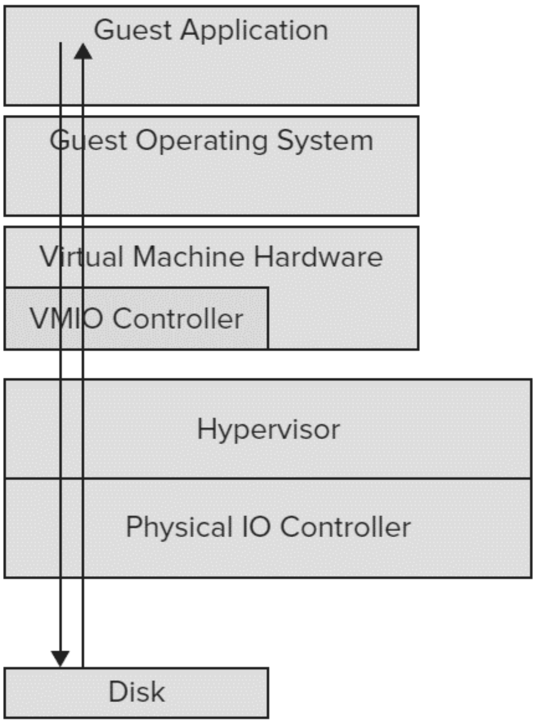

Figure 2.9 A simplified data request

In the following same relevant terms such as clones, templates, snapshots and OVF (Open Virtualization Format) will briefly be summarized.

a) Understanding virtual machine clones

Server provisioning takes considerable resources in terms of time, manpower, and money. Before server virtualization, the process of ordering and acquiring a physical server could take weeks, or even months in certain organizations, not to mention the cost, which often would be thousands of dollars. Once the server physically arrived, additional provisioning time was required. A server administrator would need to perform a wide list of chores, including loading an operating system, loading whatever other patches that operating system needed to be up-to-date, configuring additional storage, installing whatever corporate tools and applications the organization decided were crucial to managing their infrastructure, acquiring network information, and connecting the server to the network infrastructure. Finally, the server could be handed off to an application team to install and configure the actual application that would be run on the server. The additional provisioning time could be days, or longer, depending on the complexity of what needed to be installed and what organizational mechanisms were in place to complete the process.

Contrast this with a virtual machine. If one needs a new server, one can clone an existing one, as shown in Figure 2.10 [Portnoy, 2012]. The process involves little more than copying the files that make up the existing server. Once that copy exists, the guest operating system only needs some customization in the form of unique system information, such as a system name and IP address, before it can be instantiated. Without those changes, two VMs with the same identity would be running the network and application space, and that would wreak havoc on many levels. Tools that manage virtual machines have provisions built in to help with the customizations during cloning, which can make the actual effort itself nothing more than a few mouse clicks.

Now, while it may only take a few moments to request the clone, it will take some time to enact the copy of the files and the guest customization. Depending on a number of factors, it might take minutes or even hours. But, if we contrast this process with the provisioning of a physical server, which takes weeks or longer to acquire and set up, a virtual machine can be built, configured, and provided in mere minutes, at a considerable savings in both man hours and cost.

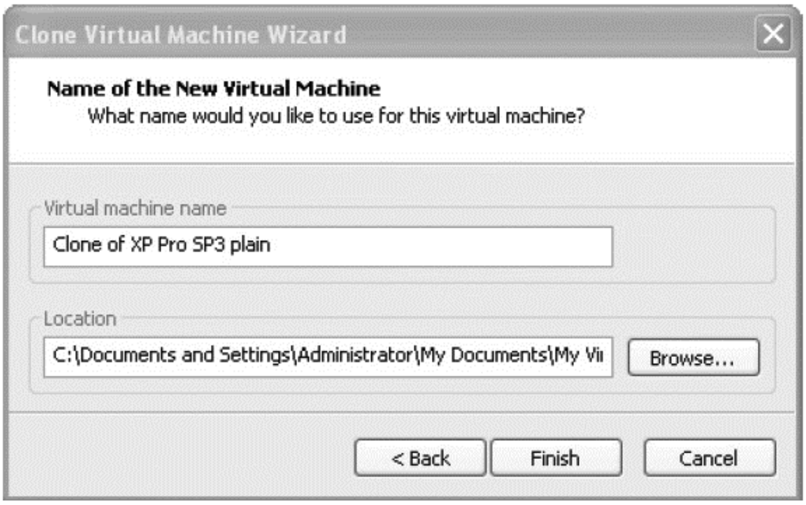

Figure 2.10 Cloning a VM

b) Understanding templates

Similar to clones, virtual machine templates are another mechanism to rapidly deliver fully configured virtual servers. A template is a mold, a preconfigured, preloaded virtual machine that is used to stamp out copies of a commonly used server. Figure 2.11 shows the Enable Template Mode checkbox to enable this capability [Portnoy, 2012]. The difference between a template and a clone is that the clone is running and a template is not. In most environments, a template cannot run, and in order to make changes to it (applying patches, for example), a template must first be converted back to a virtual machine. One would then start the virtual machine, apply the necessary patches, shut down the virtual machine, and then convert the VM back to a template. Like cloning, creating a VM from a template also requires a unique identity to be applied to the newly created virtual machine. As in cloning, the time to create a virtual machine from a template is orders of magnitude quicker than building and provisioning a new physical server.

Unlike a clone, when a VM is converted to a template, the VM it is created from is gone.

c) Understanding snapshots

Snapshots are pretty much just, what they sound like, a capturing of a VM's state at a particular point in time. They provide a stake in the ground that one can easily return to in the event that some change made to the VM caused a problem one would like to undo. Figure 2.12 is a basic illustration of how snapshots work [Portnoy, 2012]. A snapshot preserves the state of a VM, its data, and its hardware configuration. Once one snapshot a VM, changes that are made no longer go to the virtual machine. They go instead to a delta disk, sometimes called a child disk. This delta disk accumulates all changes until one of two things happens, another snapshot or a consolidation, ending the snapshot process. If another snapshot is taken, a second delta disk is created and all subsequent changes are written there. If a consolidation is done, the delta disk changes are merged with the base virtual machine files and they become the updated VM.

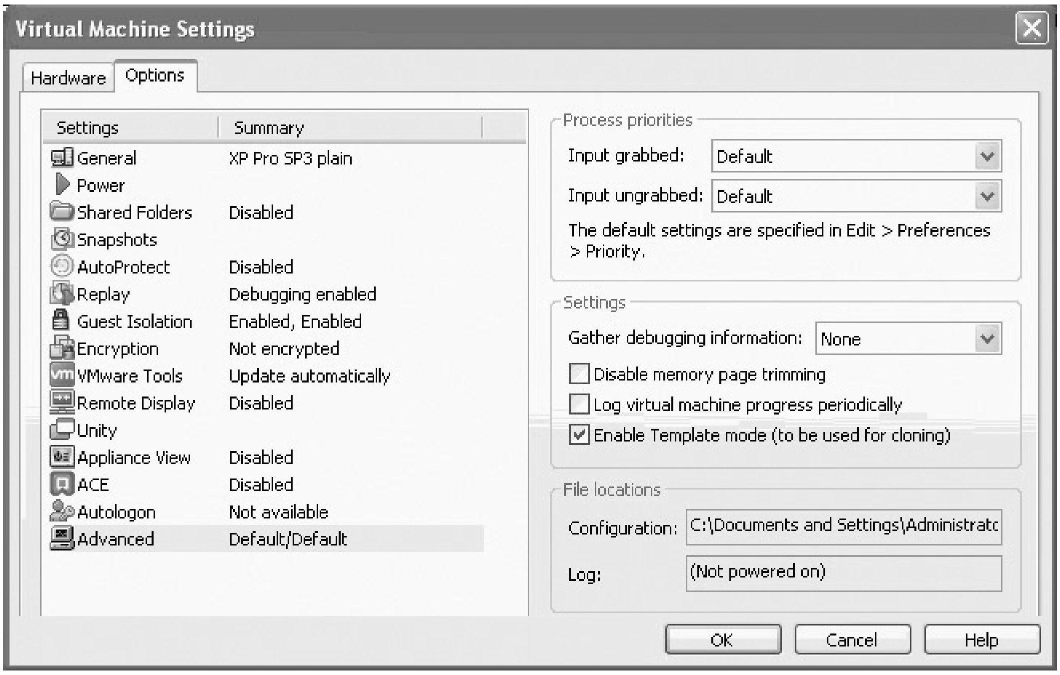

Figure 2.11 Creating a VM from a template

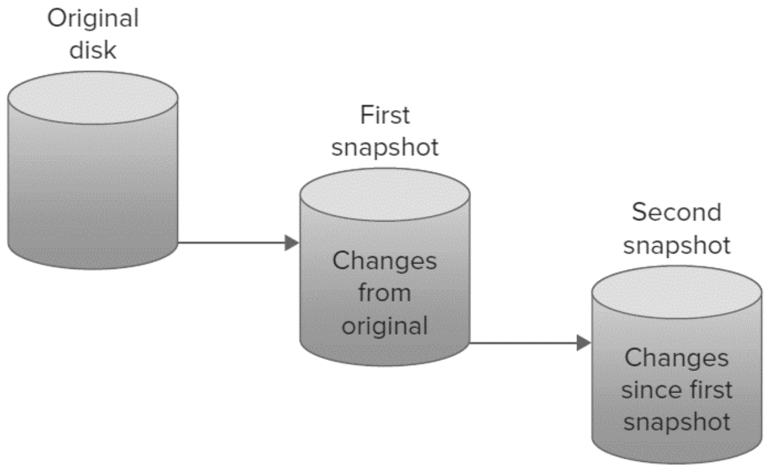

Figure 2.12 A snapshot disk chain

Finally, one can revert back to the state of a VM at the time when a snapshot was taken, unrolling all of the changes that have been made since that time. Snapshots are very useful in test and development areas, allowing developers to try risky or unknown processes with the ability to restore their environment to a known healthy state. Snapshots can be used to test a patch or an update where the outcome is unsure, and they provide an easy way to undo what was applied. Snapshots are not a substitute for proper backups. Applying multiple snapshots to a VM is fine for a test environment but can cause large headaches and performance issues in a production system.

d) Understanding OVF (Open Virtualization Format)

Another way to package and distribute virtual machines is using the OVF. OVF is a standard, created by an industry-wide group of people representing key vendors in the various areas of virtualization. The purpose of that standard is to create a platform and vendor-neutral format to bundle up virtual machines into one or more files that can be easily transported from one platform to another platform. Most virtualization vendors have options to export virtual machines out to OVF format, as well as have the ability to import OVF formatted VMs into their own formats.

The OVF standard supports two different methods for packaging virtual machines. The OVF template creates a number of files that represent the virtual machine and also supports a second format, OVA (Open Virtual Appliance/Application), which will encapsulate all of the information in a single file. In fact, the standard states, “An OVF package may be stored as a single file using the TAR (Tape Archive) format. The extension of that file shall be .ova (open virtual appliance or application).”

2.7. Creating a Virtual Machine

Virtual machines are the building blocks for today’s data centers. Once an infrastructure is in place, the process of populating the environment with workloads begins. There are two main methods of creating those virtual machines, the first being a P2V (Physical-to-Virtual) operation, and the second being a from-the-ground-up execution.

Virtualization technology and its use by corporations did not spring up instantly out of nowhere. Even though increasing numbers of virtual machines are being deployed today, there are still many millions of physical servers running application workloads in data centers. A large percentage of new workloads are still being deployed as physical servers as well. As data center administrators embraced virtualization, there are two strategies for using virtualization in their environments: containment and consolidation as explained earlier in Section 2.2. Containment is the practice of initially deploying new application workloads as virtual machines so additional physical servers will not be required, except to expand the virtual infrastructure capacity.

Consolidation is the practice of taking existing physical servers and converting them into virtual servers that can run atop the hypervisor. In order to achieve the significant savings that virtualization provided, corporate IT departments needed to migrate a significant portion of their physical server workloads to the virtual infrastructure. To effect this change, they needed tools and a methodology to convert their existing physical servers into virtual machines.

In the following investigating the physical-to-virtual converting process and hot and cold cloning will be summarized.

a) Investigating the physical-to-virtual (P2V) process

The process of converting a physical server to virtual, often shortened to the term P2V, can take a number of paths. The creation of a brand new virtual machine is usually the preferred method. A new virtual machine provides the opportunity to load the latest operating system version and the latest patches as well as rid the existing workload of all previous application files. Also, as part of creating a virtual machine, one will see that certain physical drivers and server processes are no longer needed in the virtual environment. When one creates a new VM, part of the process needs to make these adjustments as well.

If one was to perform this process manually for multiple servers, it would be tremendously repetitive, time-consuming, and error prone. To streamline the effort, many vendors have created P2V tools that automate the conversion of existing physical servers into virtual machines. Instead of creating a clean virtual server and installing a new operating system, the tools copy everything into the VM. Older servers that run applications or environments that are no longer well understood by their administrators are at risk when the hardware becomes no longer viable. The operating system version may no longer be supported, or even be capable of running on a newer platform. Here is a perfect situation for a P2V - the ability to transfer that workload in its native state to a new platform independent environment where it can be operated and managed for the length of its usable application life without concern about aging hardware failing. The other advantage of this process, is that a system administrator would not have to migrate the application to a new operating system where it might not function properly, again extending the life and usability of the application with minimal disruption.

One of the disadvantages of a P2V conversion can be an advantage as well. The P2V process is at its very core a cloning of an existing physical server into a virtual machine. There are some changes that occur during the process, translating physical drivers into their virtual counterparts and certain network reconfigurations, but ultimately a P2V copies what is in a source server into the target VM.

In addition to the various hypervisor solution providers, a number of third-party providers also have P2V tools. Here is a partial list of some of the available tools:

- VMware Converter;
- Novell Platespin Migrate;
- Microsoft System Center VMM (Virtual Machine Manager);
- Citrix XenConvert;
- Quest Software vConverter;
- Symantec System Recovery.

b) Hot and cold cloning

There are two ways of performing the P2V process: Hot and Cold. There are advantages and disadvantages to both. Cold conversions are done with the source machine nonoperational and the application shut down, which ensures that it is pretty much a straight copy operation from old to new. Both types involve a similar process:

- Determine the resources being used by the existing server to correctly size the necessary resources for the virtual machine;
- Create the virtual machine with the correct configuration;
- Copy the data from the source (physical) server to the target (virtual) server;
- Run post-conversion cleanup and configuration. This could include network configuration, removal of applications and services that are not required for virtual operations, and inclusion of new drivers and tools.

Hot cloning, as the name implies, performs the clone operation while the source server is booted and the application is running. One disadvantage to this is that data is constantly changing on the source machine, and it is difficult to ensure that those changes are migrated to the new VM (Virtual Machine). The advantage here is that not all applications can be suspended for the period of time that is necessary to complete a P2V. A hot clone allows the P2V to complete without disrupting the application. Depending on where the application data is being accessed from, it may be even simpler. If the data is already kept and accessed on a SAN (Storage Area Network) rather than local storage in the server, the physical server could be P2Ved, and when the post-conversion work and validation is completed, the physical server would be shut down, and the disks remounted on the virtual machine. This process would be less time-consuming than migrating the data from the local storage to a SAN along with the P2V.

The length of time to complete a P2V is dependent on the amount of data to be converted, which correlates directly to the size of the disks that need to be migrated to a virtual machine. More disks and larger disks require more time. Times can vary widely from just an hour to maybe a day. Most systems, however, can comfortably be done in just a few hours, and very often P2Vs are run in parallel for efficiency. Vendors and service organizations have years of experience behind them now with these conversions and have created P2V factories that allow companies to complete dozens or hundreds of migrations with minimal impact on a company's operation and with a high degree of success. At the end of the process, a data center has a lot fewer physical servers than at the start.

2.8. Understanding Applications in a Virtual Machine

Virtualization and all of its accompanying benefits are changing the way infrastructures are designed and deployed, but the underlying reasons are all about the applications. Applications are the programs that run a company's business, provide them with a competitive advantage, and ultimately deliver the revenue that allows a company to survive and grow. With the corporate lifeblood at risk, application owners are reluctant to change from the existing models of how they have deployed applications to a virtual infrastructure. Once they understand how a virtual environment can help mitigate their risk in the areas of performance, security, and availability, they are usually willing to make the leap. Hypervisors leverage the physical infrastructure to ensure performance resources. Multiple virtual machines can be grouped together for faster and more reliable deployments. As both corporate and commercial services begin to shift to cloud computing models, ensuring that the applications supported on the virtual platforms are reliable, scalable, and secure is vital to a viable application environment.

2.8.1. Examining virtual infrastructure performance capabilities

We have concentrated so far on virtual machines and the virtual environment that supports them. While this is valuable, the result has to be that the applications deployed on physical servers can be migrated to these virtual environments and benefit from properties one has already investigated. Applications are groups of programs that deliver services and information to their users. These services and information provide income for their companies. Fearful that the service will be degraded, the groups responsible for the applications (the application owners) are often reluctant to make changes to their environments. Application owners are unwilling to risk application changes that might impact the application's availability, scalability, and security. The ease of altering a virtual machine's configuration to add additional resources can make virtual machines more scalable than their physical counterparts. Other virtualization capabilities, such as live migration or storage migration, bring greater flexibility and agility to applications in a virtual environment. Another area where virtualization provides great benefits is the creation and manageability of the virtual machines through templates and clones, which can significantly reduce application deployment time and configuration errors, both areas that affect a company's bottom line. All of these are important, but probably most crucial is application performance.

Applications that perform poorly are usually short lived because they impact a business on many levels. Aside from the obvious factor that they extend the time it takes to accomplish a task and drive down efficiency, slow applications frustrate users, both internal and external to a company, and could potentially cost revenue. Again, it raises the topic of increased user expectations. Think about one's own experiences with online services. Would one continue to purchase goods from a website where the checkout process took 20 minutes, or would one find another vendor where it would be less cumbersome? This is one reason why application owners are hesitant about virtualization—they are unsure about sharing resources in a virtualization host when their current application platform is dedicated entirely to their needs, even though it might be costly and inefficient.

2.8.2. Deploying applications in a virtual environment

The best way to be sure that an application performs well is to understand the resource needs of the application, but more importantly, to measure that resource usage regularly. Once one understands the requirements, one can begin to plan for deploying an application in a virtual environment. There are some things that one can always count on. A poorly architected application in a physical environment is not necessarily going to improve when moved to a virtual environment. An application that is starved for resources will perform poorly as well. The best way to be sure an application will perform correctly is to allocate the virtual machine enough resources to prevent contention. Consider the following example [Portnoy, 2012]:

Many applications are delivered in a three-tier architecture, as shown in Figure 2.13. The configuration parameters in the figure are merely sample numbers. There is a database server where the information that drives the application is stored and managed. Usually, this will be Oracle, Microsoft SQL (Structured Query Language) Server, or maybe the open-source solution MySQL. This server is typically the largest one of the three tiers with multiple processors and a large amount of memory for the database to cache information in for rapid response to queries. Database servers are resource hungry for memory, CPU, and especially storage I/O throughput. The next tier is the application server that runs the application code—the business processes that define the application. Often that is a Java-oriented solution, IBM Websphere, Oracle (BEA) WebLogic, or open-source Tomcat. In a Microsoft environment, this might be the .NET framework with C#, but there are many frameworks and many application languages from which to choose. Application servers usually need ample CPU resources, have little if any storage interaction, and have average memory resources. Finally, there is the webserver. Webservers are the interface between users and the application server, presenting the application's face to the world as HTML (HyperText Markup Language) pages. Some examples of webservers are Microsoft IIS and the open-source Apache HTTP (HyperText Transfer Protocol) server. Webservers are usually memory dependent because they cache pages for faster response time. Swapping from disk adds latency to the response time and might induce users to reload the page.

When one visits a website, the webserver presents with the HTML pages to interact. As one selects functions on the page, perhaps updating one's account information or adding items to a shopping cart, the information is passed to the application server that performs the processing. Information that is needed to populate the web pages, such as one's contact information or the current inventory status on items one is looking to purchase, is requested from the database server. When the request is satisfied, the information is sent back through the application server, packaged in HTML, and presented to as a webpage. In a physical environment, the division of labor and the division of resources is very definite since each tier has its own server hardware and resources to utilize. The virtual environment is different.

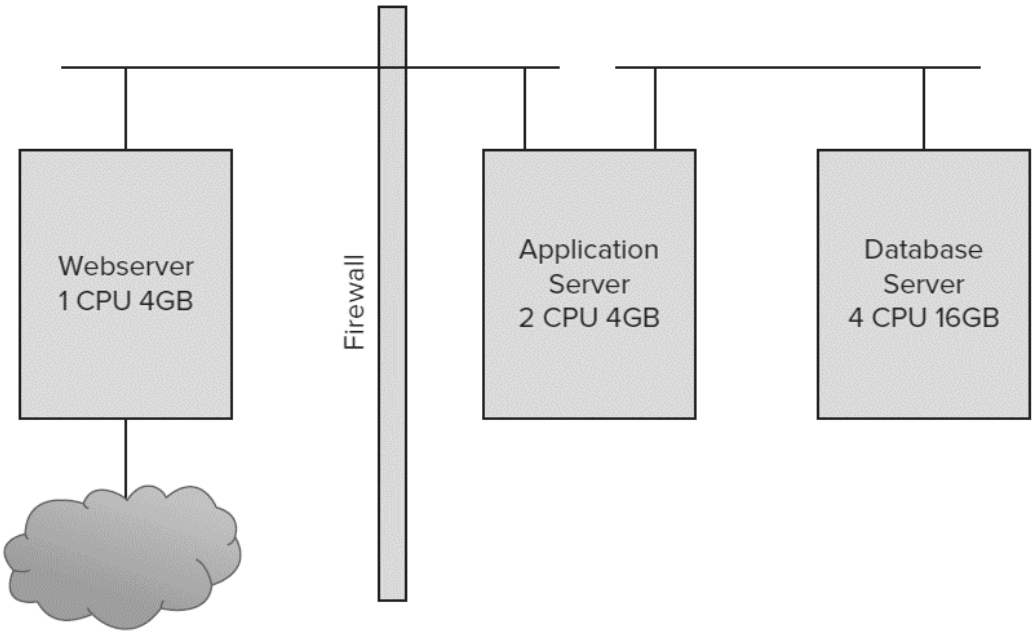

Figure 2.13 Three-tier architecture-physical

Figure 2.14 shows a possible architecture for this model. Here, all of the tiers live on the same virtual host [Portnoy, 2012]. In practice, that is probably not the case, but for a small site it is definitely possible. The first consideration is that the virtualization host now needs to be configured with enough CPU and memory for the entire application, and each virtual machine needs to have enough resources carved out of that host configuration to perform adequately. The virtual machine resource parameters discussed earlier -shares, limits, and reservations- can all be used to refine the resource sharing. Note that while in the physical model, all of the network communications occurred on the network wire; here, it all takes place at machine speeds across the virtual network in the virtualization host. Here also, the firewall separating the webserver in the DMZ from the application server can be part of the virtual network. Even though the application server and the database server are physically in the same host as the webserver, they are protected from external threats because access to them would need to breach the firewall, the same as in a physical environment. Because they do not have direct access to an external network, through the firewall is the only way they can be reached.

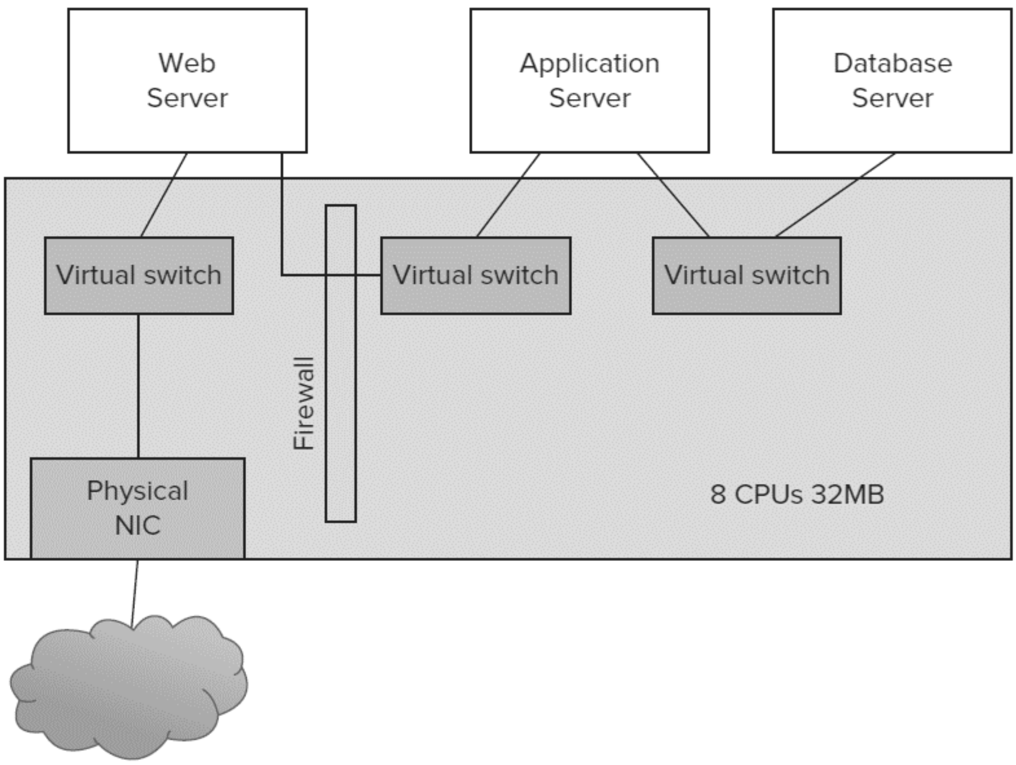

Figure 2.14 Three-tier architecture-virtual

As the application performance requirements change, the model can easily adjust. Applications that need to support many users run multiple copies of the webserver and the application server tiers. In a physical deployment, it would not be unusual to have dozens of blade servers supporting this type of application. Load balancers are placed between the tiers to equalize traffic flow and redirect it in the event of a webserver or application server failure. In a virtual environment, the same can be true when deploying load balancers as virtual machines. One large difference is that as new webservers or application servers are needed to handle an increasing load, in a virtual environment, new virtual machines can be quickly cloned from an existing template, deployed in the environment, and used immediately. When there are numerous cloned virtual machines on a host running the same application on the same operating system, page sharing is a huge asset for conserving memory resources. When resource contention occurs in a virtualization cluster, virtual machines can be automatically migrated, assuring best use of all of the physical resources. Live migration also removes the necessity of taking down the application for physical maintenance. Finally, in the event of a server failure, additional copies of the webserver and application server on other virtualization hosts keep the application available, and high availability will restore the downed virtual machines somewhere else in the cluster.

With all of the different layers and possible contention points, how do one know what is happening in an application? There are tools that will monitor the activity in a system and log the information so it will be available for later analysis and historical comparison. This information can be used to detect growth trends for capacity modeling exercises, allowing the timely purchase of additional hardware, and prevent a sudden lack of resources. Virtualization vendors supply basic performance management tools as part of their default management suite. Additional functionality is offered as add-on solutions for purchase. There is also a healthy third-party market of tools that supports multiple hypervisor solutions. As always, there are many tools developed as shareware or freeware and easily available as a download. All of these can be viable options, depending on one’s particular use case. The point is that measuring performance, and understanding how an application is functioning in any environment, should be a mandatory part of an organization’s ongoing application management process.

2.8.3. Understanding virtual appliances and vApps

A number of people in the IT departments probably created the three-tier application discussed earlier, although in smaller shops one person can wear all of the hats. A virtualization administrator created the original virtual machines and configured them according to some basic parameters. An operating system engineer provisioned the operating systems on each one and then updated it with the latest patches. Any corporate standard tools were added at this time. The application developer or specialist then installed and configured the application components—webserver, application server, application code, and database. Integration tests were run to validate and stress the virtual machines as an application unit. Once all the testing was complete and any resulting changes had been applied, they were converted into templates, the gold images to produce the production virtual machines. One learned earlier that virtual machines can speed this provisioning procedure by orders of magnitude, but it is still a transfer of the physical server provisioning process to a virtual infrastructure without fundamental operational changes. Virtual appliances change this model.

Virtual appliances are prebuilt virtual machines that already contain everything needed to deploy an application. Often the operating system is an open-source deployment or a specially developed thin OS, also called JeOS (Just enough Operating System) that only has what the application requires and no more. Because of this, virtual appliances require none of the traditional patching and maintenance of a traditional operating-system-based system. When a new version is available, the entire virtual machine is replaced, minimizing the time needed to deploy a new release. The installation in many cases consists of a download, unpacking the virtual machine on a virtualization host, powering it on, and performing some minimal configuration steps to connect it to the desired network and storage. Virtual appliances are often delivered in OVF format so they can be quickly deployed on any hypervisor solution.

3. A FICTITIOUS ENTERPRISE INFORMATION SYSTEM FOR A SERVER VIRTUALIZATION APPLICATION

3.1. General

As mentioned in previous chapters, server virtualization of an enterprise involves various topics. In this thesis, due to the complexity of such a process, an application is developed to demonstrate the technicalities of server virtualization for an enterprise as an example of a life cycle for future applications.

The proposed application aims to increase efficiency in the use of hardware resources, reduce management and resource costs, improve business flexibility, improve security and reduce downtime. Every enterprise differs from each other according to the environment it operates, needed infrastructure and desired level of security. Every type of organizations use different types and amount of information technologies to operate. It is clear that finding an enterprise to willingly share this type of information is very hard and even for some enterprises to share such information and discuss it in public may be unethical or prohibited.

Regarding the confidentiality issues and diversity of organizational needs for the proposed application a fictitious enterprise information system will be defined so that all the steps of server virtualization will be applied on that information system. Thus, the major aim in the study is to assume an architecture to deploy server virtualization infrastructure and show all the necessary technical steps for application.

3.2. Definition of a Fictitious Enterprise Information System

In order to define a sample enterprise information system one should first decide which aspects are necessary for the definition. In order to define the enterprise as simple as possible, answers to the following questions are used:

a) Is it a private or public organization?
b) How many internal users are present?
c) Is there any service exposed to internet?

d) How important confidentiality, integrity and availability to the organization?
e) How many external users are present? And how often do they use the system?
f) Is there any remote office(s)?
g) Are there any mobile internal users?

In the application, the assumed properties of the fictitious enterprise information system are summarized in the Table 3.1.

Table 3.1 Properties of the fictitious organization

Assumed Property	Attribute
Type of organization	Private
Number of daily internal users	250
Services over internet	Yes
Security categorization	High
Number of weekly external users	10
Remote offices	Yes
Mobile users	Yes

3.2.1. System roles

Roles of the assumed servers of the enterprise are given in Table 3.2.

Table 3.2 System roles of servers

System Name	Roles
Apollo	Active Directory / DNS (Domain Name System)
Charon	File Server
Dinlas	Printer Server
Helios	DHCP (Dynamic Host Configuration Protocol)
Hypnos	Web Server
Kratos	FTP (File Transfer Protocol) Server
Morpheus	Mailbox Server One
Oceanus	Mailbox Server Two
Plutus	CAS (Client Access Server) / Hub Transport
Triton	CAS / Hub Transport Server Two
Uranus	Application Server
Zelus	Database Server
Notus	Mobile Device Control Server
Dionysus	Proxy Log Server
Boreas	Antivirus Server

Based on that data, system consisting of 15 servers is depicted as Figure 3.1.

3.2.2. Server hardware specifications

In Table 3.3 hardware specifications of the servers in the assumed enterprise are given.

Table 3.3 Specifications of servers (*)

System	Make / Model	CPU	Ram	Disk
Apollo	HP ProLiant DL140 G3	2 x Intel Xeon 5140 @ 2.33GHz	3,00 GB	100 GB
Charon	HP ProLiant DL320 G4	2 x Intel Pentium D 820	2,00 GB	300 GB
Dinlas	HP ProLiant DL140 G2	2 x Intel Xeon	2,00 GB	150 GB
Helios	HP ProLiant DL140 G3	2 x Intel Xeon 5140 @ 2.33GHz	2,00 GB	100 GB
Hypnos	HP ProLiant DL120 G6	2 x Intel Xeon X3450 @ 2.67GHz	5,00 GB	1000 GB
Kratos	HP ProLiant DL140 G2	2 x Intel Xeon	2,00 GB	200 GB
Morpheus	HP ProLiant DL320 G4	2 x Intel Xeon	2,00 GB	500 GB
Oceanus	HP ProLiant DL140 G3	2 x Intel Xeon 5140 @ 2.33GHz	1,00 GB	80 GB
Plutus	HP ProLiant DL120 G6	2 x Intel Xeon X3450 @ 2.67GHz	4,00 GB	200 GB
Triton	HP ProLiant DL140 G3	2 x Intel Xeon 5140 @ 2.33GHz	1,00 GB	100 GB
Uranus	HP ProLiant DL140 G2	2 x Intel Xeon	2,00 GB	150 GB
Zelus	HP ProLiant DL380 G4	2 x Intel Xeon	4,00 GB	150 GB
Notus	HP ProLiant DL380 G5	2 x Intel Xeon E5320 @ 1.86GHz	3,00 GB	100 GB
Dionysus	HP ProLiant DL380 G4	2 x Intel Xeon	3,00 GB	150 GB
Boreas	HP ProLiant DL180 G6	2 x Intel Xeon E5520 @ 2.27GHz	8,00 GB	250 GB

(*) In all servers, windows server operating system is assumed.

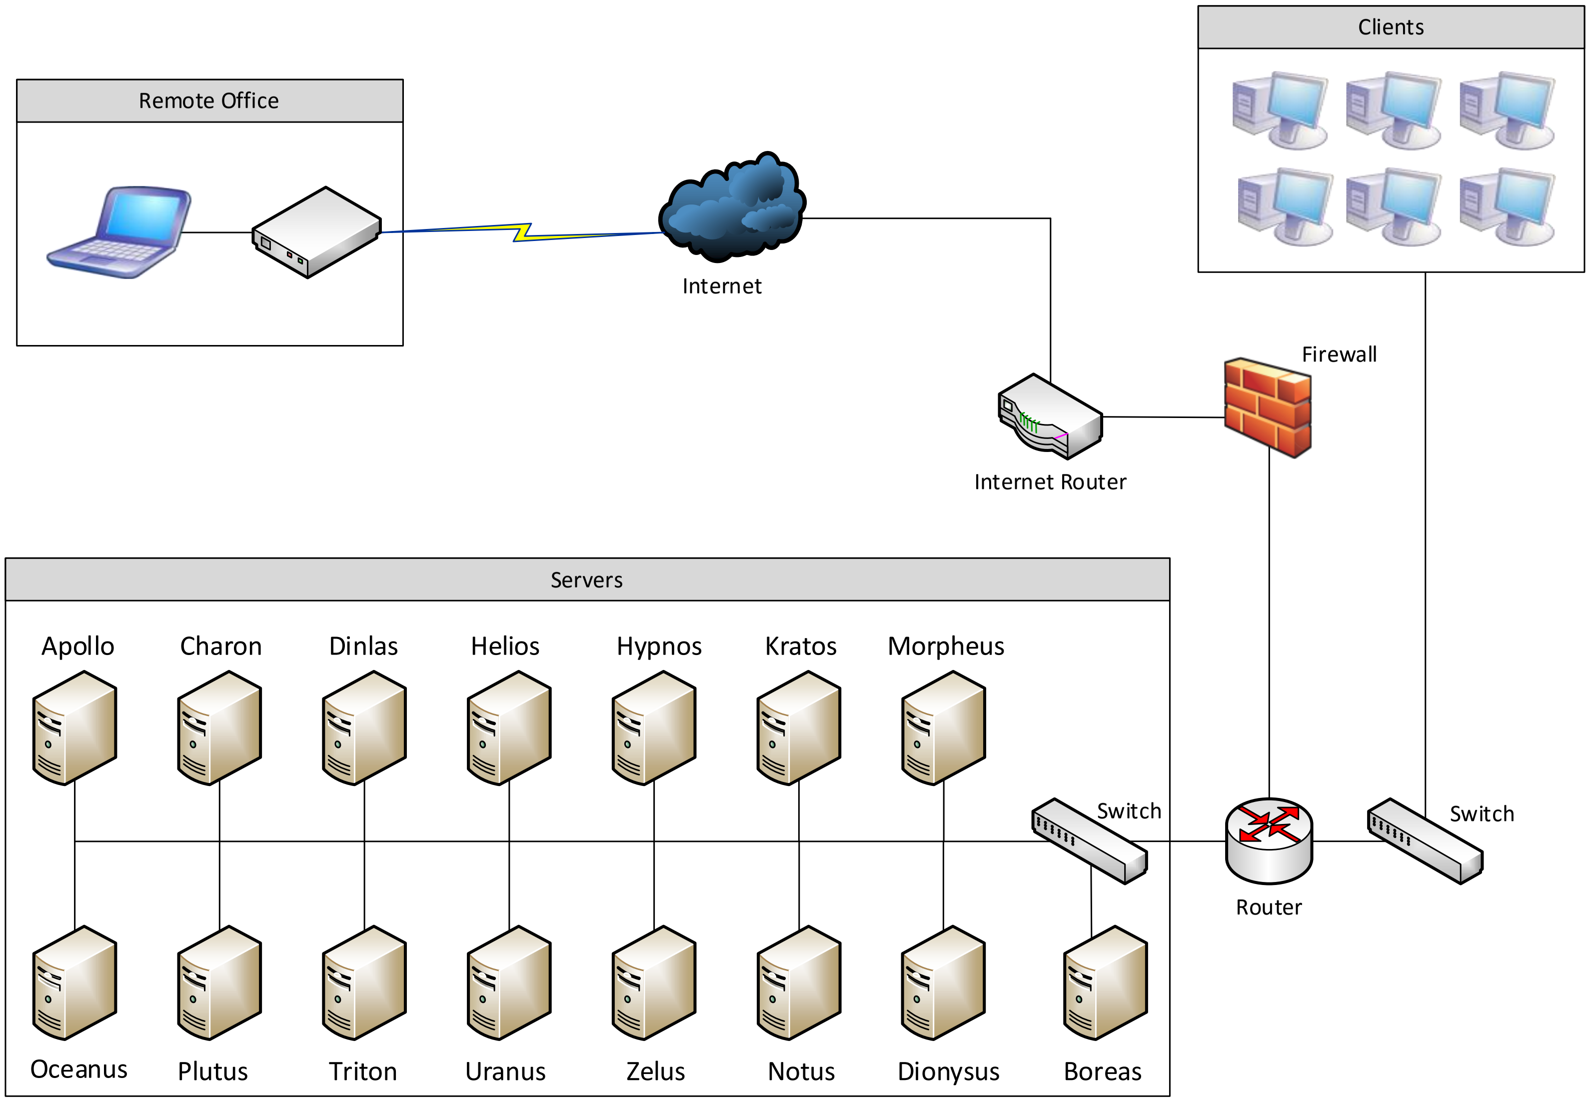

Figure 3.1 Assumed organization's network infrastructure

4. EXISTING LIFE CYCLES FOR SERVER VIRTUALIZATION

In order to apply server virtualization on the virtual system defined in Chapter 3, the available life cycles for this purpose are investigated. In the investigation it is found that there are three major references for this purpose given by VMware [2006], Dell and VMware or Siebert [2009], and Dell-Intel [2010]. The following sections have their brief summaries.

4.1. VMWare Life Cycle

In VMWare [2006], there are five phases such as “Build CoE”, “Assess” “Plan and Design”, “Build”, and “Manage”. These phases are shown in Figure 4.1.

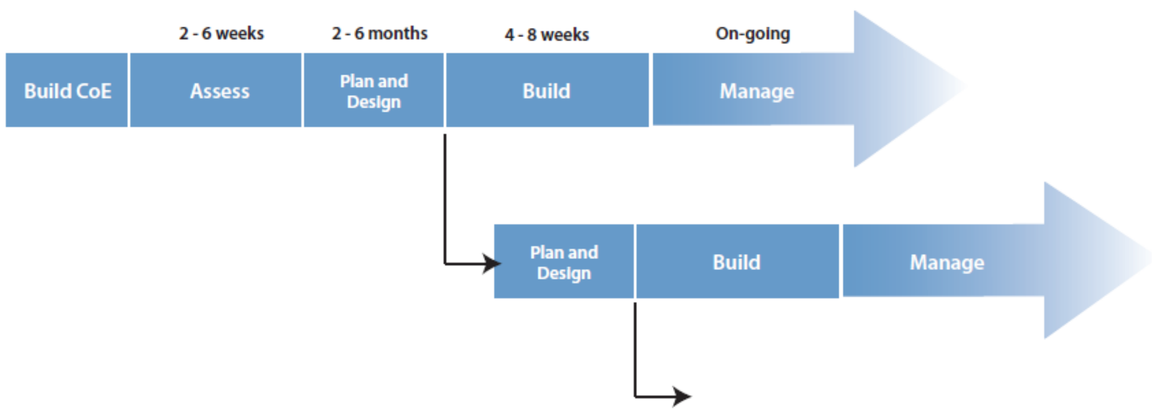

Figure 4.1 Phases of virtual infrastructure deployment and their approximate durations

The phases are defined as follows:

- Build a Virtualization Center of Excellence (Build CoE) (2 to 4 weeks);
- Phase 1: Conduct Operational Assessment (2 to 6 weeks);
  - Step 1. Determine operational readiness;
  - Step 2. Assess how well the IT infrastructure environment is understood and documented;
  - Step 3. Review current/planned projects to assess impact and identify candidate projects for collaboration;

- Step 4. Document current application portfolio and fit with virtualization;

- Phase 2: Plan/Design (2 to 6 months);
- Phase 3: Build (4 to 8 weeks);
- Phase 4: Manage (ongoing).

4.2. Dell and VMWare or E. Siebert Life Cycle

The next life cycle is developed together by Dell and VMWare in discovering the stages of a virtualization project. These phases are summarized as follows [Siebert, 2009]:

a. Assessing current environment;
b. Planning virtual environment;
c. Building virtual environment;
d. Configuring virtual environment;
e. Securing virtual environment;
f. Populating virtual environment;
g. Monitoring virtual environment;
h. Maintaining virtual environment;
i. Backing up virtual environment;
j. Troubleshooting virtual environment.

4.3. Dell and Intel Life Cycle

Another life cycle for Virtualization Lifecycle is proposed by Dell and Intel [2010] joint work and it is summarized in Table 4.1.

Table 4.1 Lifecycle phases of virtualization by Dell and Intel

Phase	**Major Activities/Concerns**
Planning	Test lab/QA Discovery/Inventory Capacity Planning
Provisioning	Bare Metal Hypervisor Deployment Workload Migration
Protect	Backup/Disaster Recovery High Availability
Operations	Charge Back Workload Orchestration SLA/Policy Management VM Performance VM Update/HW Update
Compliance	Audit & Compliance Antivirus, Firewall

Note: In the above table QA stands for Quality Assurance, SLA stands for Service Level Agreement and VM stands for Virtual Machine.

5. A PROPOSED LIFE CYCLE FOR SERVER VIRTUALIZATION IN AN ENTERPRISE

5.1. General

In Chapter 4, three different references for the Life Cycle for Server Virtualization were reviewed: VMware [2006], Siebert [2009] and Dell and Intel [2010]. In these references virtualization deployment projects consist of several phases. Although all references describe a life cycle for server virtualization, they have some similarities and differences between each other. Because of that, a new life cycle created by combining three different references is proposed in this chapter.

5.2. A Proposed Life Cycle

The steps of the proposed new life cycle is defined as follows:

a. Build Center of Excellence;
b. Perform Operational Assessment;
c. Plan and Design;
d. Build;
e. Configure;
f. Provide Security;
g. Populate;
h. Monitor;
i. Maintain;
j. Backup;
k. Troubleshooting.

In the following section, each of these steps are briefly summarized.

5.3. Description of the Steps in the Proposed Life Cycle

a. Build Center of Excellence [VMware, 2006]

VMware [2006] defined a center team that consist of members that has working knowledge and experience on virtualization environment, knowledge on current ICT base operations and has reliability with business application owners. The major members for the team are defined as follows [RHI, 2008]:

- ***Project Manager*** needs a combination of subject-matter expertise and project management skills. He/she must possess excellent communication, interpersonal and team leadership abilities as well as the capacity to work with cross-functional team to accomplish overall project goals.
- ***Virtualization Expert*** has working knowledge and experience of a virtualized environment.
- ***System analyst*** must be good analytical thinkers and problem solvers as well as impressive communicators. They need a broad understanding and experience about hardware and software systems such as installation, maintenance and life cycles.
- ***Network engineers*** must have in-depth knowledge of networking hardware and software. They have experiences about network design and implementation, LAN/WAN interfacing, security, TCP/IP, server and network infrastructure.
- ***Data security analysts*** must possess a thorough understanding of all aspects of computer and network security including such areas as firewall administration, encryption technologies and network protocols.
- ***Database Administrators*** need a strong technical foundation in database structure, configuration, installation and practice.
- ***Storage Expert*** is someone who has a common task such as disk, tape, SAN, and storage software management.
- ***Application Developers*** must have strong analytical and problem-solving abilities about applications from both a technical/programming perspective and a business point of view.
- ***Backup Expert*** is someone who has a common task such as backing up, restoring and monitoring backup jobs.

b. Operational Assessment [Siebert, 2009]

Evaluation of the environment is a standout amongst the most basic strides in a virtualization venture. While executing a virtualization infrastructure, one will most likely change over existing physical servers to virtual machines (VMs). Therefore, it is vital to comprehend the execution attributes of the used servers and applications.

To evaluate present environment, utilize execution checking devices to gauge normal and crest usage of CPU, memory, organize and disk assets on the servers to be virtualized. Accumulate this information for at least a week, in spite of the fact that a month is best.

It is additionally essential to accumulate execution measurements during basic business cycles (e.g., week-by-week finance handling or a month-to-month reporting process). One ought to likewise consider particular instruments, for example, VMware's Capacity Planner[4] to gather this data.

c. Plan and Design [Siebert, 2009]

While arranging a virtual environment, one has a few choices to make. In the first place, pick a virtualization stage on which to have VMs. There are a few virtualization vendors in the commercial center, and one has to pick a stage that fits his financial plan and has the elements he requires.

VMware[5] is the present market leader in x86 virtualization[6], and the vSphere[7] release is compelling because it is component rich and has strong execution measurements and expansive visitor OS support. However, it should be noted that these elements include some significant downfalls. Some of VMware's releases

4 **VMWare Capacity Planner** is a virtualization estimate tool that assesses an IT infrastructure and its usage.

5 **VMware, Inc.** is an American company that provides cloud and virtualization software and services.

6 In computing, **x86 virtualization** refers to hardware virtualization for the x86 architecture.

7 **VMware vSphere** is VMware's cloud computing virtualization operating system.

can be exorbitant, particularly those with a more noteworthy number of components.

All things considered, be that as it may, be careful about seller asserts – particularly with regards to examinations of their items with those of different suppliers. To that end, one must perform own particular item assessments and examination to figure out which item best suits his requirement.

d. Build [Siebert, 2009]

Planning and estimating one's equipment is another vital stride in a virtualization venture on the grounds that one will in all probability need to buy new server equipment for virtual hosts.

Deciding capacity for one's virtual hosts is another imperative choice on the grounds that it will significantly affect one's VM execution and the sorts of components that one can utilize. Shared storage solutions such as FC (Fiber Channel)[8], iSCSI (Internet Small Computer System)[9] and NFS (Network File System)[10] are necessary to take advantage of certain advanced features such as HA (High Availability) and LM (Live Migration).

While estimating one's virtual hosts, consider leaving space to oblige extra VMs if a host disappointment occurs. What is more, leave space accessible for future development of the base.

e. Configure [Siebert, 2009]

When environment is ready, it's an ideal opportunity to design it to set it up for one's virtual machines (VMs). Organizing and stockpiling are normally the greatest

8 **Fibre Channel**, or FC, is a high-speed network technology (commonly running at 2-, 4-, 8- and 16-gigabit per second rates) primarily used to connect computer data storage.

9 **iSCSI** is an acronym for Internet Small Computer Systems Interface, an Internet Protocol (IP)-based storage networking standard for linking data storage facilities.

10 **Network File System (NFS)** is a distributed file system protocol, allowing a user on a client computer to access files over a computer network much like local storage is accessed.

and most essential arrangement steps, yet there are likewise numerous little steps that one ought to design appropriately. While arranging a system, comprehend one's prerequisites and environment heretofore. It's a smart thought to counsel with one's system gathering to guarantee that it comprehends the requirements of one's host servers while joining them to a physical system.

Designing capacity is an undertaking that one needs to get right the first run through. It can be troublesome and tedious to change stockpiling arrangements after they have been set up. On the off chance that one settles on the wrong decisions while designing capacity, they can undermine host and VM execution. So set aside time to comprehend one's prerequisites, and guarantee that one altogether comprehends one's choices before arranging stockpiling on one's hosts.

Further, storage can be sophisticated to configure and generally involves working with a storage area network (SAN)[11] administrator to accurately connect to the host servers. It is ought to be worked closely with SAN administrators to ensure that they understand needs and properly configure logical unit numbers, or LUNs[12].

While designing a virtualized environment, the main issue is to ensure one completely comprehends what one is attempting to arrange.

f. Provide Security [Siebert, 2009]

In a virtual environment, administrators frequently neglect to set aside the opportunity to appropriately secure the environment, which is a major oversight. Many exposed hypervisors are sensibly secure out of the container, yet there is dependably opportunity to get better. Moreover, it is very easy to make a hypervisor unsecure by misconfiguration and changing native settings. Security is

[11] **A storage area network (SAN)** is a network which provides access to consolidated, block level data storage.

[12] In computer storage, a logical unit number, or **LUN**, is a number used to identify a logical unit, which is a device addressed by the SCSI protocol or Storage Area Network protocols which encapsulate SCSI, such as Fibre Channel or iSCSI.

enhanced in a virtual situation, where a solitary physical server runs numerous virtual servers, and lacking security on a physical server can straightforwardly influence the security of every single virtual server running on that host.

Numerous security executives are careful about virtual hosts due to the expanded security dangers furthermore some misperceptions about what makes VMs unreliable. Set aside an ideal opportunity to disclose to one's security group how security functions in a virtual situation. Additionally, plot the additional steps one has taken to assist secure hosts and virtual machines. When one instructs them about virtual security, they ought to be more agreeable and willing to work with one.

g. Populate [Siebert, 2009]

Virtual machines are anything but difficult to make - truth be told, too simple. Be that as it may, indeed, VMs can bring about considerable issues once they engender. One of the most serious issues in virtual situations is VM sprawl[13], or the uncontrolled development of virtual machines in a virtual domain. VM sprawl is comparative, in that virtual machines regularly get made without respect for the assets that they expend and, hence, these VMs can overpower the host server's assets.

In order to address sprawl, one may execute one of the numerous items that backing chargeback and making writes about asset use for virtual situations. Likewise, restricting the quantity of individuals who can make virtual machines and building up a formal procedure for asking for new virtual machines can avoid sprawl and unmonitored virtual machines. One ought to consider requiring legitimization for solicitations for any new virtual machines and organize an endorsement procedure to compel clients to mull over whether they have to make another VM. At long last, making asset pools can confine the measure of assets accessible on one's host servers for new virtual machines.

[13] **Virtualization sprawl** is a phenomenon that occurs when the number of virtual machines (VMs) on a network reaches a point where the administrator can no longer manage them effectively.

h. Monitor [Siebert, 2009]

Observing a virtualization domain is essential to guarantee that it stays solid and works legitimately. Frequently, issues may not be self-evident, and a decent observing framework cautions one to issues so they can be determined. In virtual situations, even little issues can have real impacts in light of the fact that such a large number of virtual machines keep running on a solitary host, and all these VMs fight for that host's assets. So it's imperative not to overlook observing; without it, one’s virtual surroundings may attempt to let one knows something that one can't hear in light of the fact that one is not tuning in.

There are a few things that one ought to screen, for example, execution, server equipment and virtualization programming particular cautions and occasions. Hardware failures of hosts may be disturbing for virtual environments despite advantages such as HA (High Availability) and FT (Fault Tolerance) designed to minimize system downtime. When a fan, disk or ram module has broken so one can make a move can minimize interruption to one’s environment.

i. Maintain [Siebert, 2009]

Ignore patching and updates at one’s own peril. Software patches and upgrades are necessary to resolve defects and security vulnerabilities and to update software with new features. With virtual hosts, it is even more important to apply patches in a timely manner because a defect or vulnerability can affect many virtual machines on that host. Each virtualization vendor has a different patch schedule, so be aware of them, and sign up for alerts when new patches are released. Pay special attention to security patches, and apply them as quickly as possible. One should establish a regular patch routine, and after patches and updates have been released, don not delay in applying them in one’s environment.

j. Backup [Siebert, 2009]

Good backups are an absolute necessity for any environment, physical or virtual. Customary reinforcement systems for physical situations can in any case be utilized as a part of virtual situations; however, there are choices that can go down one's virtual machines all the more productively and rapidly. Virtual situations take into account distinctive ways to deal with reinforcements due to their snapshot capacities and remarkable construction modeling. Taking a snapshot of a VM's virtual disk suspends keeps in touch with it and considers reinforcements to happen without the likelihood of documents being changed amid the reinforcement. Further, many third-party products such as Veeam[14] and Commvault[15] are specifically developed to backup virtual machines. They are worth considering for even better integration, efficiency and backup speed.

k. Troubleshooting [Siebert, 2009]

Trouble is a given in even the best-designed environments: Hardware can fail, bugs will strike, and unintended outcomes happen. Troubleshooting virtual environments can be more complicated than physical ones because of the architectural differences between them. When a problem occurs in one's virtualization environment, there are more components to check, and trying to look for a problem's cause can be more difficult. Troubleshooting can often be frustrating, and knowing where to look and what to do is the key to quickly finding and resolving problems. Familiarize himself with the many different locations in which log files are located, and know which log files to check for specific problems.

[14] **Veeam** Software is a privately held information technology company that develops backup, disaster recovery and virtualization management software for VMware and Hyper-V virtual environments. [Veeam, 2016]

[15] **Commvault** Systems, Inc. (NASDAQ: CVLT), founded in 1996, is a publicly traded data management and information management software company headquartered in Tinton Falls, New Jersey. [Commvault, 2016]

6. APPLICATION

In Chapter 5, a life cycle for server virtualization in an enterprise was proposed. In the following, steps of that life cycle will be followed sequentially for the server virtualization of the virtual information system defined already in Chapter 3.

6.1. Build Center of Excellence

Building an in-house team group who have experience of virtualization environment, know current IT base operations and have acceptability with business application proprietors and also IT system pioneers. After that, virtualization team is created. Team members were already defined in Section 5.3.

6.2. Operational Assessment

Enterprise is engaged to run a server consolidation analysis to determine how a server consolidation would be beneficial for the enterprise and to evaluate a selected server consolidation scenario. The VMware Capacity Planner assessment tool ran for approximately 30 days from November 1, 2015 through November 30, 2015. Data was collected on 15 servers specified by the enterprise for review and analysis during the installation. The performance data were collected and analyzed.

This assessment has been developed to establish and define the technical workload that exists for the identified 15 servers. Data were collected for 4 weeks on these servers. After analysis, report is automatically created from the tool. According to this report, inventory and performance data of servers are generated.

The list of 15 servers in the enterprise was reviewed by the enterprise to verify that all servers were consolidation candidates. Using the workload of these consolidation candidates as an input to a capacity model, several consolidation scenarios were built and then compared against the existing costs. Upon review of these possible scenarios, the enterprise can then determine the feasibility of any of these as potential consolidation projects.

This operational assessment report demonstrates how one's business can realize dramatic savings in server hardware, storage, network, data center power and cooling, real estate space, and disaster recovery by implementing server virtualization solutions. In the following subsections results in the report are presented.

6.2.1. VMWare capacity planner

This report is generated, using a VMware product called Capacity Planner. The report generated is a plan based on the virtual enterprise's specific environment and goes beyond the results that an interview process can provide.

VMWare Capacity Planner Collector gathers inventory and performance data that is sent to the Data warehouse, analyzed by the Data Analyzer, and presented.

The data collected from environment for analysis includes:

- Hardware and software inventory to provide capacity and system purpose;
- Hardware resource utilization;
- Application specific utilization.

6.2.2. Summary of environment

Enterprise's environment was monitored for 4 weeks. The data show a significant underutilization of current capacity with opportunities to save money using virtualization.

6.2.2.1. System count

Number of CPUs in the system is given in the Table 6.1 below.

6.2.2.2. Processor summary

CPU Utilization of systems information is given in the Table 6.2 below.

Table 6.1 Number of CPUs in the systems

CPU(s)	Systems
2	10
4	2
8	3
Total Systems	15

Table 6.2 CPU utilization percentage of systems

Average CPU Utilization	3,54
Peak CPU Utilization	5,57

6.2.3. Assessment summary results

The assessment of the fictitious organization includes two scenarios. Each scenario is created to show the potential consolidation by using different rules. The scenario rules are described in detail in later sections.

One scenario provided the best results. This was determined using the lowest number of systems required. If one has other goals, a different scenario can be used.

Total power, cooling, and space requirements are possible alternative goals. The analysis engine cannot always determine total power, cooling, and space requirements of used systems. Examine the table carefully to determine if existing systems have enough chassis information to be included in the table.

6.2.4. Scenario comparison

This assessment included two analysis scenarios. The aggressive scenario has been chosen as the optimum path for enterprise. It has the desired qualities for a successful virtualization plan. Two Scenario are given in the Table 6.3.

Table 6.3 Two alternative scenario results

Scenario Name	System Analyzed	Host Required	System Exceptions	Overall Consolidation Ratio %	Total Storage (TB)	Rack Used	KW Used	Tons BTU/hr Used
Consolidation Scenario (Moderate)	15	2	0	87	5	4.0	2.00	0.85
Consolidation Scenario (Aggressive)	15	2	0	87	10	8.0	4.00	1.70

6.2.5. Inventory results

Capacity Planner found 15 systems on all domains at the enterprise. All servers were inventoried for this consolidation study.

6.2.5.1. Overview

The inventoried and monitored time and count of servers are given in Figure 6.1 below.

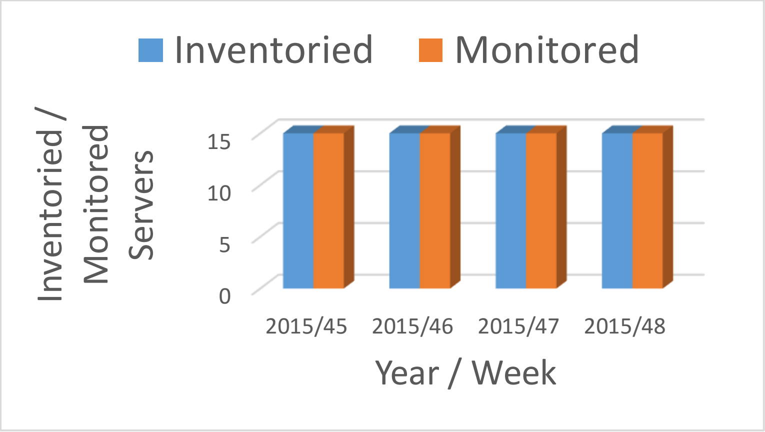

Figure 6.1 Inventoried and monitored servers

6.2.5.2. <u>Consolidatable servers</u>

Consolidatable servers' numbers are given in Figure 6.2 below.

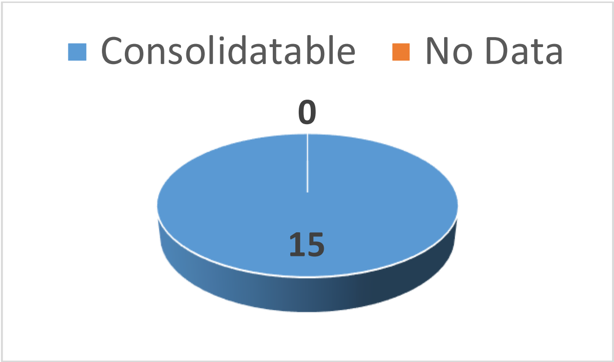

Figure 6.2 Consolidatable server count

6.2.5.3. <u>System roles</u>

Roles of servers are given in Table 6.4.

Table 6.4 Roles of servers

Server Name	Roles
Apollo	Active Directory / DNS Server
Charon	File Server
Dinlas	Printer Server
Helios	DHCP Server
Hypnos	Web Server
Kratos	Ftp Server
Morpheus	Mailbox Server One
Oceanus	Mailbox Server Two
Plutus	CAS / Hub Transport Server One
Triton	CAS / Hub Transport Server Two
Uranus	Application Server
Zelus	Database Server
Notus	Mobile Device Control Server
Dionysus	Proxy Log Server
Boreas	Antivirus Server

6.2.5.4. <u>Source systems</u>

Server specifications were already given in Chapter 3, in Table 3.3. Additional characteristics of servers are produced in VMware Capacity Planner Tool and given as Table 6.5.

Table 6.5 Source systems capacity data

System Name	Make / Model	Capacity										
		Processors		Memory	Disk	Network		Physical				
		Count	Speed (MHz)	Size (MB)	Size (GB)	Count	Speed (Mb/sec)	Rack Units	Weight (lbs)	Power (KW)	Thermal (BTU/hr)
Apollo	HP/ProLiant	2	2,327	3,071	100	1	1024	1	35	0.65	2,228.00
Charon	HP/ProLiant	2	2,793	2,048	298	1	1024	1	27	0.45	1,710.00
Dinlas	HP/ProLiant	2	3,600	2,047	150	1	1024	1	35	0.50	2,371.00
Helios	HP/ProLiant	2	2,327	2,047	100	1	1024	1	35	0.65	2,228.00
Hypnos	HP/ProLiant	8	2,660	5,120	1000	1	1024	1	31	0.40	1,710.00
Kratos	HP/ProLiant	2	3,600	2,048	200	1	1024	1	35	0.50	2,371.00
Morpheus	HP/ProLiant	2	2,739	2,048	500	1	1024	1	27	0.45	1,710.00
Oceanus	HP/ProLiant	2	2,327	1,024	80	1	1024	1	35	0.65	2,228.00
Plutus	HP/ProLiant	8	2,660	4,096	200	1	1024	1	31	0.40	1,710.00
Triton	HP/ProLiant	2	2,327	1,024	100	1	1024	1	35	0.65	2,228.00
Uranus	HP/ProLiant	2	3,600	2,048	150	1	1024	1	35	0.50	2,371.00
Zelus	HP/ProLiant	4	3,000	4,096	150	1	1024	2	60	0.57	2,508.00
Notus	HP/ProLiant	4	1,866	3,072	100	1	1024	2	60	0.85	3,390.00
Dionysus	HP/ProLiant	2	3,600	3,072	150	1	1024	2	60	0.57	2,508.00
Boreas	HP/ProLiant	8	2,267	8,192	250	1	1024	2	45	0.46	1,773.00

6.2.5.5. <u>Utilization</u>

Utilization data of servers are given in Table 6.6.

Table 6.6 Source systems utilization data

	Utilization										
	Processor		Memory				Disk				Network
	% Used	Queue per CPU per GHz	% Used	File Sys Cache (MB)	Page File %	Paging File %	I/O (Trans/sec)	I/O (MB/sec)	Read Speed (MB/Sec)	Write Speed (MB/Sec)	Speed (MB/Sec)
Apollo	5.43	0.16	83.43	103.04	7.85	529.16	60.88	2.29	2.09	0.28	0.01
Charon	2.46	0.06	79.01	284.95	33.74	99.32	15.44	0.51	0.37	0.27	0.02
Dinlas	5.27	0.05	47.02	193.32	10.96	77.12	42.75	0.73	0.28	0.45	0.31
Helios	0.67	0.44	57.02	333.68	43.21	42.15	9.63	0.21	0.12	0.09	0.00
Hypnos	5.31	0.02	64.48	1,306.22	18.89	91.82	27.46	0.42	0.37	0.15	0.01
Kratos	5.63	0.64	69.13	158.55	19.16	15.65	5.07	0.14	0.07	0.08	0.05
Morpheus	24.14	0.94	75.47	192.07	34.11	26.63	18.91	0.50	0.15	0.35	0.02
Oceanus	10.32	6.43	86.97	135.29	55.70	84.80	24.24	0.42	0.19	0.23	0.04
Plutus	2.64	0.01	89.76	373.97	10.60	55.54	14.29	0.44	0.19	0.25	0.02
Triton	5.87	0.15	65.54	390.05	27.98	32.00	8.14	0.15	0.08	0.07	0.16
Uranus	4.55	0.05	37.98	235.68	3.55	376.46	152.76	2.74	1.38	1.36	0.01
Zelus	10.71	0.08	57.71	481.64	3.55	16.19	10.70	0.12	0.02	0.11	0.03
Notus	13.84	0.04	45.67	353.36	4.26	221.85	68.87	1.07	0.87	0.32	0.00
Dionysus	2.47	0.18	63.78	132.56	5.50	465.60	169.37	3.09	2.28	1.55	0.01
Boreas	0.33	0.00	52.89	1,109.78	22.82	438.16	9.30	0.11	0.03	0.08	0.05

Table values show that some servers' CPU usage rates are very low. For example, Boreas usage is 0.33 percent. It shows that some servers are underutilized. By virtualizing these servers, server resources will be used more efficiently.

6.2.5.6. Exception systems

No server is an exception.

6.3. Plan and Design

The optimum consolidation scenario selected was Scenario (Aggressive HW Template). The result of this scenario are:

- 15 systems were analyzed;
- 2 Host servers were needed to host 15 virtual systems;
- The consolidation ratio for eligible systems is 15 to 2, a 87% decrease in systems.

6.3.1. Selected virtualization vendor

There are quite a number of virtualization technology companies in the sector. VMware, Microsoft and Citrix are well-known in virtualization. VMware vSphere and Microsoft Hyper-V are the most popular and successful products in server virtualization. VMware was a leader and Microsoft was a challenger in the market in May 2010 [Bittman et al., 2010]. In July 2015, both VMware and Microsoft were at the leaders' side [Bittman et al., 2015]. Because of license cost advantages, technology improvements and large number of supported Windows Operating Systems, Microsoft Hyper-V is selected in the thesis study.

6.3.2. Selected host platform

According to the consolidation scenario (Aggressive HW Template), specifications for each virtualization host is given in Table 6.7 below.

The total storage required by the scenario is not based on disk storage of individual hosts.

Table 6.7 Virtualization host specification

Hardware Component	**Number**	**Size, Speed**
CPUs	2	3500 MHz
RAM	8	8 GB
Network Interface Cards	8	1 Gb/sec
Disk Storage	3	300 GB

6.3.3. Next steps

Operational assessment report is the beginning of the consolidation project. VMware Capacity Planner provides an initial plan for the following items:

- The number of hosts of the chosen configuration that are needed to start a virtualization project;
- The optimal placement of systems to hosts;
- Expected host and based on a 4-week average.

Environment change rapidly. The validity of this assessment diminishes over time. The initial recommendation is good for project budgeting. If this is one's first assessment, it is a good idea to take the following actions:

- Validate this consolidation estimate with technology team;
- Adjust the consolidation analysis to improve the consolidation ratio;
- Leverage these results to build a business case with financial analysis.

After initial assessment and planning, assessment should be repeated as needed to monitor the change of the environment in order to add more resources to hosts such as RAMs (Random Access Memory).

6.4. Build

According to VMware Capacity report two servers are enough to host 15 virtual machines and one storage to store virtual machines. The following tables, namely Table 6.8, 6.9 and 6.10 provides an overview of the type of servers recommended as virtualization host and storage according to report.

Table 6.8 Host server 1 properties

	Model	**Quantity**
Server Brand	HP ProLiant DL380p	1
Processors	Intel® Xeon® Processor E5-2637 v2	2
Memory	8GB UDIMM @1600MHz	8
Internal Storage	HP 300GB 12G SAS 10K rpm SFF (2.5-inch)	3
Networking	HP Ethernet 1Gb 4-port 331FLR Adapter	2
Form Factor	2U*	1

* A rack unit, U as a unit of measure describes the height of electronic equipment designed to mount.

Table 6.9 Host server 2 properties

	Model	**Quantity**
Server Brand	HP ProLiant DL380p	1
Processors	Intel® Xeon® Processor E5-2637 v2	2
Memory	8GB UDIMM @1600MHz	8
Internal Storage	HP 300GB 12G SAS 10K rpm SFF (2.5-inch)	3
Networking	HP Ethernet 1Gb 4-port 331FLR Adapter	2
Form Factor	2U	1

Table 6.10 Storage properties

Attributes	**Model**
Storage Brand	HPE MSA 1040 2 Port 1G ISCSI DC
Product Type	SAN Array
SAN Interface	iSCSI
Controller Type	6Gb/s SAS
Total Number of 2.5" Bays/Disk	24
Used 2.5" Bays/Disk	13 x 900GB 6G SAS 10K RPM SFF
Ethernet Technology	Gigabit Ethernet
Used Hard Drive Capacity	10.8 TB
Form Factor	2U

6.5. Configure

In this section, Windows Server 2012 R2 Hyper-V, Failover Clustering and iSCSI storage technologies are used.

6.5.1. Failover clustering overview

“A failover cluster is a group of independent computers that work together to increase the availability and scalability of clustered roles (formerly called clustered applications and services). The clustered servers (called nodes) are connected by physical cables and by software. If one or more of the cluster nodes fail, other nodes begin to provide service (a process known as failover)” [Microsoft Technet(a), 2016].

6.5.2. Windows Server 2012/R2 Hyper-V role overview

“The Hyper-V server role in Windows Server lets one creates a virtualized server computing environment where one can create and manage virtual machines. One can run multiple operating systems on one physical computer and isolate the operating systems from each other. With this technology, one can improve the efficiency of his computing resources and free up his hardware resources” [Microsoft TechNet (b), 2016].

6.5.3. iSCSI overview

“iSCSI is a way of connecting storage devices over a network using TCP/IP (Transmission Control Protocol/IP Control). It can be used over a LAN (Local Area Network), a WAN (Wide Area Network), or the Internet” [Microsoft, 2016].

“iSCSI devices are disks, tapes, CDs (Compact Disk), and other storage devices on another networked computer that one can connect to. Sometimes these storage devices are part of a network called a SAN” [Microsoft, 2016].

“In the relationship between his computer and the storage device, his computer is called an initiator because it initiates the connection to the device, which is called a target” [Microsoft, 2016].

6.5.4. Topology

There are two host servers that were mentioned in section 6.4 as Host Server 1 (hv-node1) and Host Server 2 (hv-node2) and one shared iSCSI storage.

In the topology in Figure 6.3, there are two other physical servers and nine switches that were not mentioned in section 6.4. One of physical server is for Domain Controller (DC) and the other physical server for Virtual Machine Management (MN). Two switches are used for LAN network. Two switches are used for VM network. One switch is used for heartbeat. Two switches are used for DMZ (DeMilitarized Zone) network. Two switches are used for iSCSI Network. In addition to that, two LAN switches are used for clients, other servers and management of hosts.

Following steps are explained in the Appendices section:

A. Topology Overview;
B. Hyper-V Failover Clustering Installation Requirements;
C. Hyper-V Installation;
D. Hyper-V Host Configuration;
E. Storage Configuration;
F. Hyper-V Nodes iSCSI Initiator Advanced Configuration;
G. Hyper-V Failover Cluster Installation and Configuration.

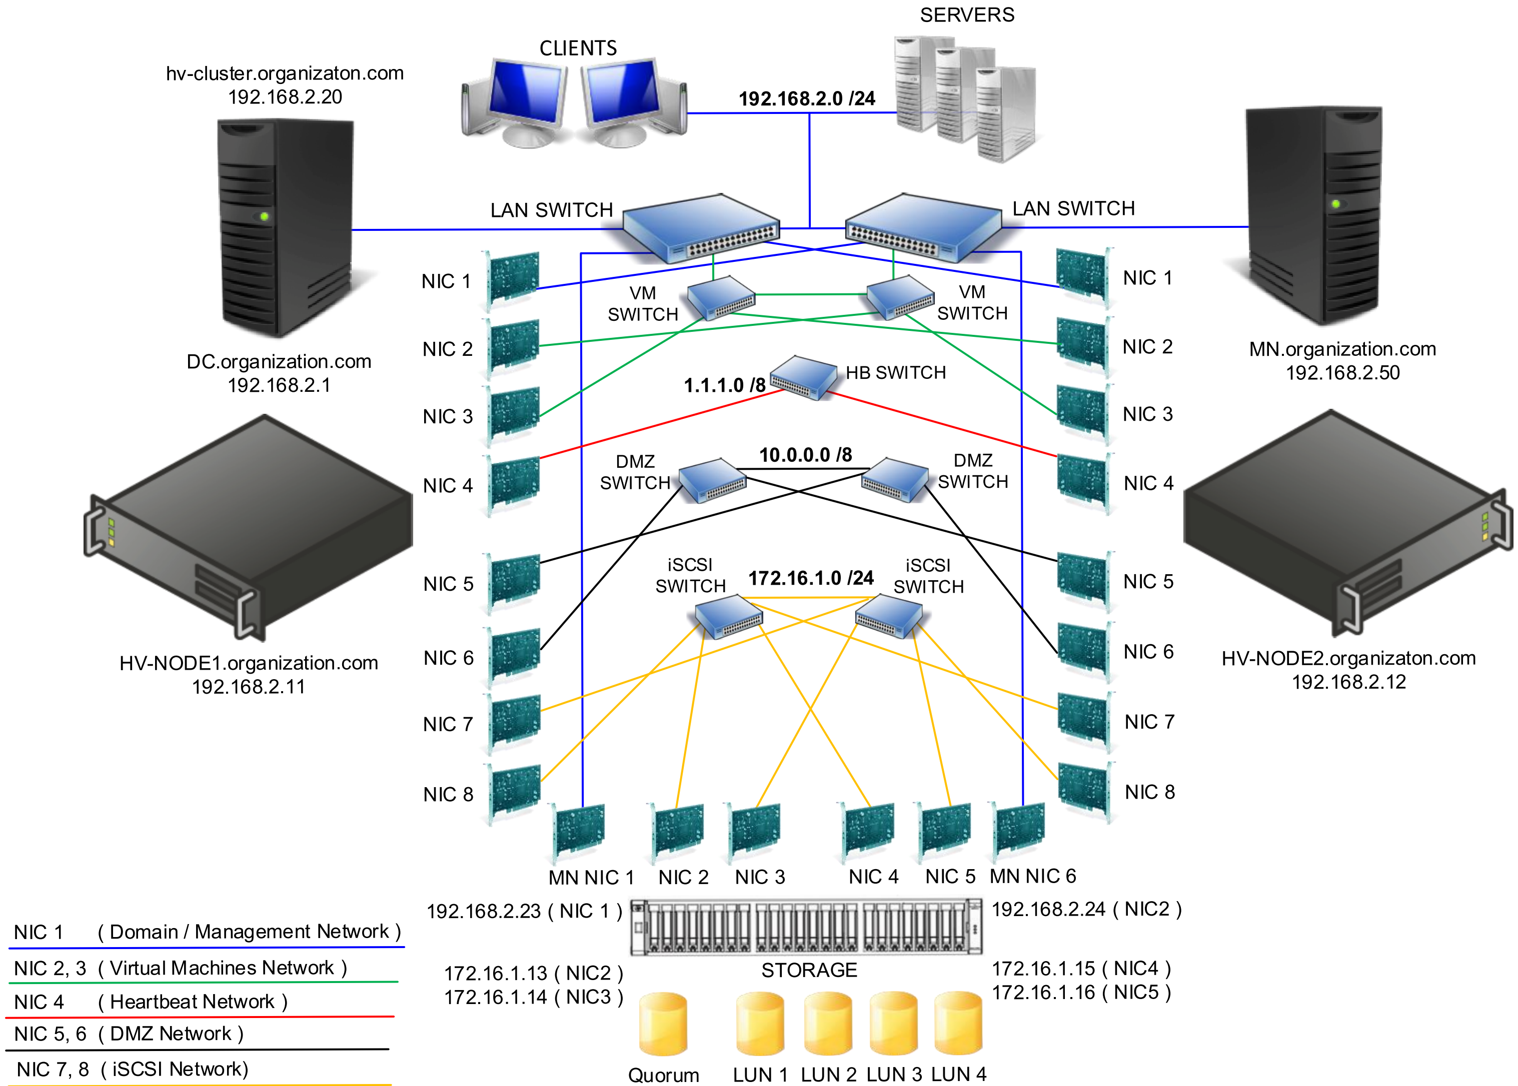

Figure 6.3 Windows Server 2012 Hyper-V iSCSI based failover cluster

6.6. Provide Security

"Hyper-V Administrators" has been included in Windows Server 2012 R2 (Release 2). Members from the Hyper-V Administrators have complete access to all features and configuration of Hyper-V. Hyper-V Administrators are not the same as system administrators in that they just perform Hyper-V particular administrators' operations. This guarantees the client part in charge of performing Hyper-V Administrators does not need to be a system administrator. It is suggested to include users with administrative responsibilities to this group instead of the local Administrators group to provide them with access to Hyper-V.

Securing Hyper-V includes every one of the measures that are required to protect any Windows Server 2012 R2 server part, in addition to securing the VMs, configuration files, and data are very important. There is a checklist to give support with enhancing the security of one's Hyper-V environment. There are two parts for security. One of them is for host and the other one is for virtual machine.

6.6.1. Host security

The following list gives the needed information about securing Hyper-V host:

- Management operating system must be up to date with the most recent security updates;
- Physical Hyper-V must be managed by dedicated physical network adapter;
- Virtual machines resource files where resides within the storage devices must be secure;
- Windows Server 2012 Security Baseline must be applied [Windows Server 2012 Security Baseline, 2016];
- Antivirus must not be installed in Hyper-V host. If Antivirus is installed in host, exclusion settings must be set for the Hyper-V system files;
- Every virtual machines administrator should manage own machine. Virtual machines administrator must not manage Hyper-V host;
- Bit locker Drive Encryption technology must be used to protect resources;

- Private or secure network must be used for VM configuration and virtual hard disk files;
- Live migration traffic must be used in private or dedicated network. IPsec (IP Security) technology may be considered for this.

6.6.2. Virtual machine security

The following list gives information about securing virtual machines that reside in Hyper-V host:

- Virtual hard disks and snapshot files must be stored in a secure location;
- Virtual switches must be configured correctly and virtual network adapters must be connected to correct virtual switches;
- Only necessary storage devices must be configured for a virtual machine;
- Windows Server 2012 Security Baseline must be applied for virtual machines that use windows operating system;
- Firewall, antivirus and intrusion detection software must be used and configured for each virtual machine;
- All virtual machines must be up to date with the most recent security updates;
- Virtual machines' integration services must be installed and up to date.

6.7. Populate

Hyper-V hosts are installed and configured as detailed in Appendices A-G. After that, 15 physical machines will be converted to virtual machines in this part. In order to convert physical machines to virtual machines, MVMC (Microsoft Virtual Machine Converter) 3.0 will be used [Microsoft Virtual Machine Converter 3.0, 2016].

MVMC is a Microsoft-supported, stand-alone solution for the Information and Communication Technology (ICT) pro or solution provider who wants to:

- Convert virtual machines and disks from VMware hosts to Hyper-V hosts and Microsoft Azure;
- Convert physical machines and disks to Hyper-V hosts.

Before one installs MVMC (Microsoft Virtual Machine Converter), one must install the following software on the computer on which MVMC will run:

- Windows Server 2012 R2, Windows Server 2012, or Windows Server 2008 R2 with SP1 (Service Pack 1) operating systems;
- Microsoft .NET Framework 3.5 and .NET Framework 4 if one installs MVMC on Windows Server 2008 R2 with SP1;
- Microsoft .NET Framework 4.5 if one installs MVMC on Windows Server 2012 R2, Windows Server 2012, or Windows 8;
- Visual C++ Redistributable for Visual Studio 2012 Update 1;
- Windows Management Framework 3.0.

In addition, enable the following features:

- BITS (Background Intelligent Transfer Service) feature in Windows Server on the destination host for physical machine conversion;
- Hyper-V on the destination host.

In order to install MVMC, follow these steps:

1. Download the MVMC Installer package to one's local server;
2. Run the Windows Installer package;
3. Enter a location where one wants to install MVMC or use the default location and continue with the installation.

It is important that the account that runs the Setup program must be the local administrator.

6.7.1. Physical to Hyper-V virtual machines conversion by using the graphic user interface

It is advised that MVMC should be installed on the destination host. In this example, hv-node1 will be used for MVMC installation.

It is assumed that MVCM is installed on hv-node1. After the installation is completed, use the following steps in order to convert Apollo physical server machine to Hyper-V:

1. Open MVMC, open the Machine Type page, click Physical machine conversion, and then click Next;

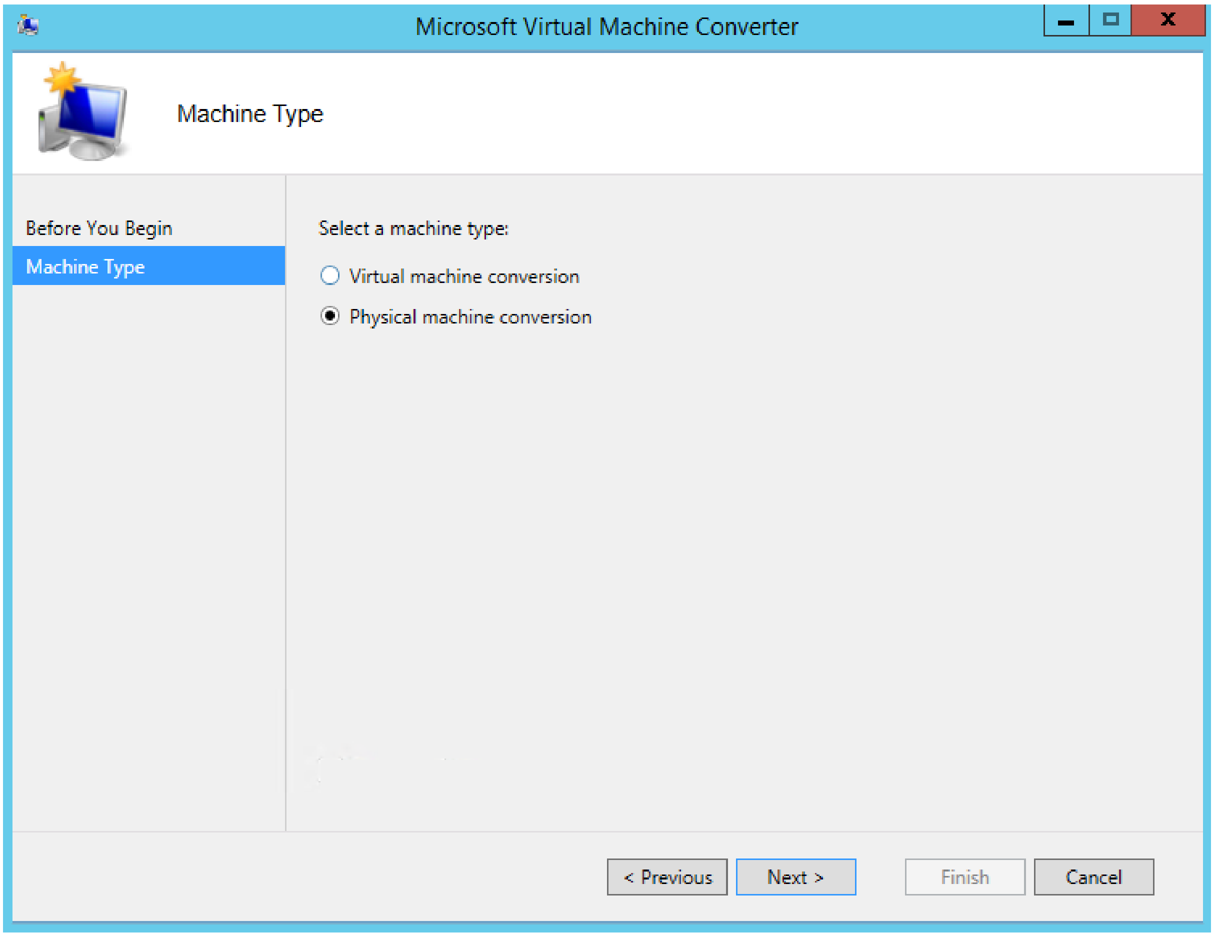

Figure 6.4 The machine type page

2. On the Source page, type apollo.organization.com for address, type administrator@organization for user name and type password for the administrator@organization;

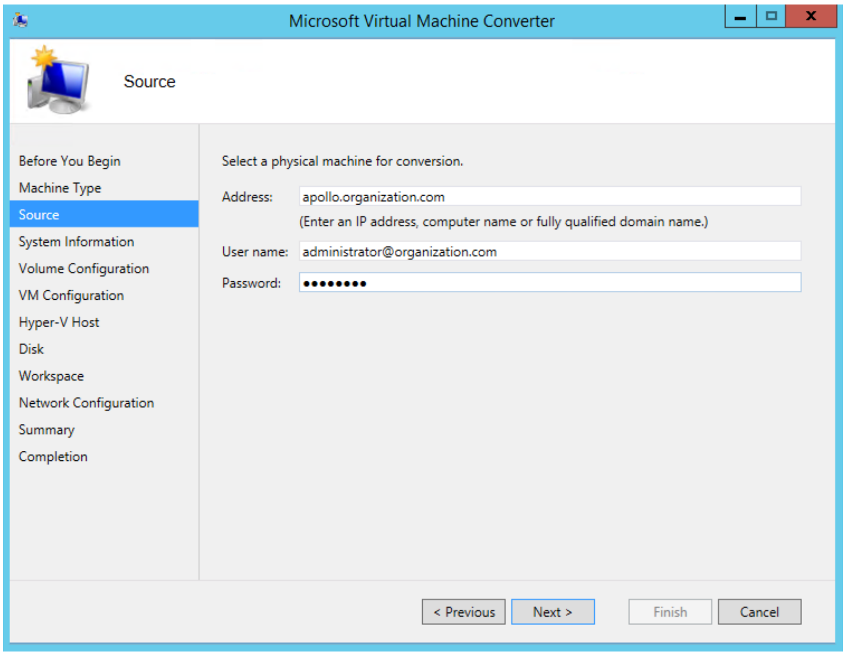

Figure 6.5 The source page

3. On the System Information page, click Scan System to temporarily install an agent on the source physical machine;

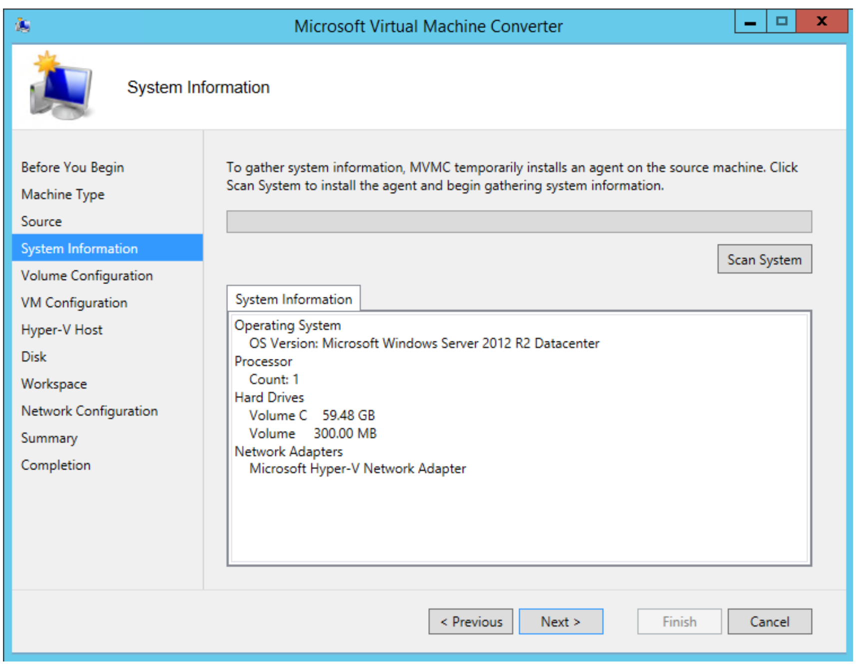

Figure 6.6 The system information page

4. On the Volume Configuration page, select the volumes that one wants to include in the conversion. Each converted virtual hard disk can be either fixed size or dynamically expanding. Click Next;

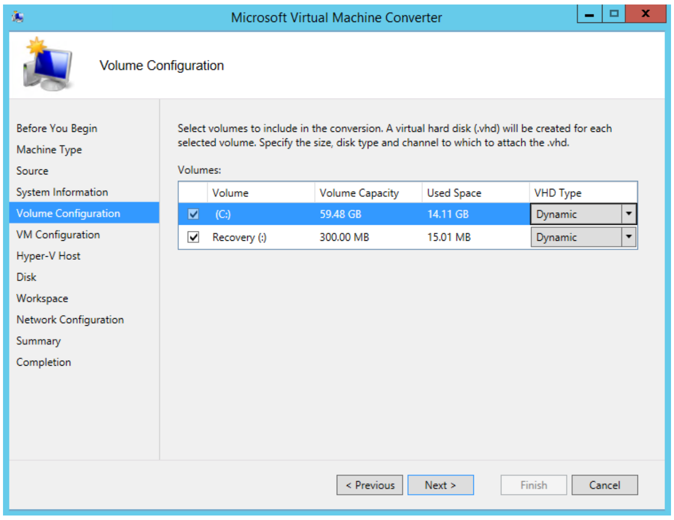

Figure 6.7 The volume configuration page

5. On the VM Configuration page, type APOLLO for virtual machine name, type 1 for number of processors, and type 2048 for memory for the target virtual machine, and then click Next;

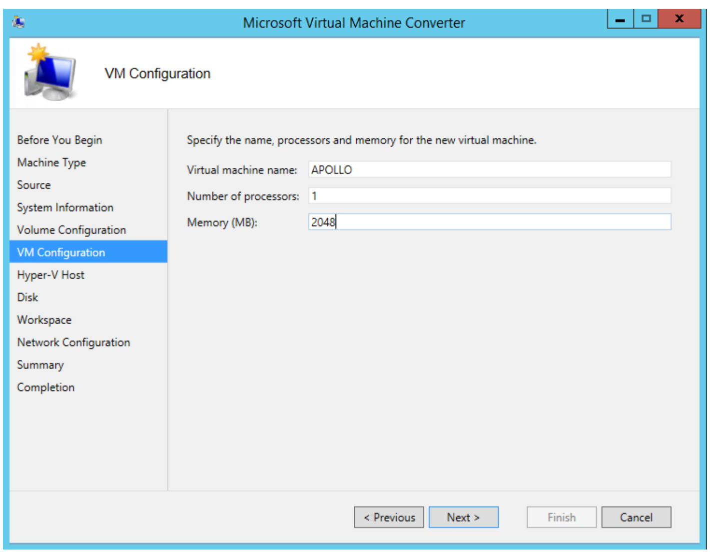

Figure 6.8 The VM configuration page

6. On the Hyper-V Host page, type HV-NODE1 for address, and then click Next;

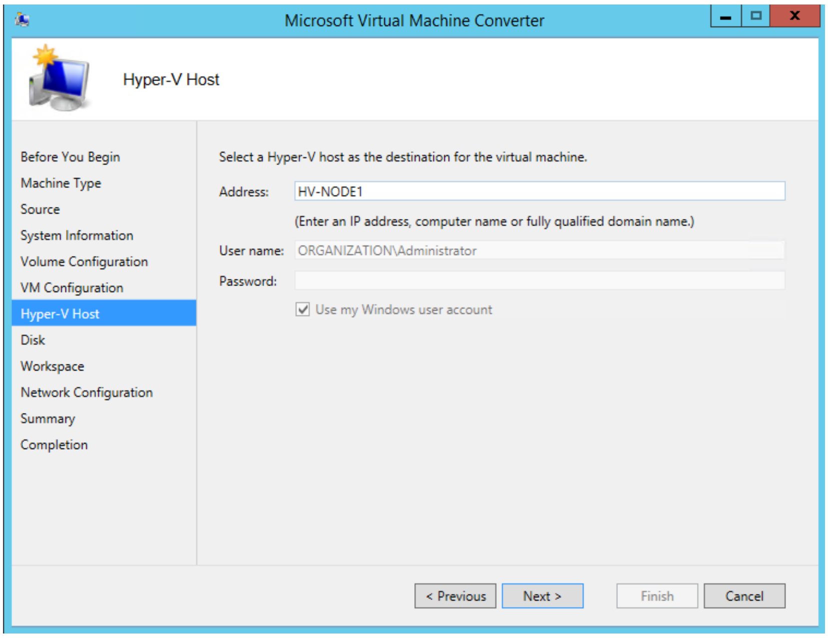

Figure 6.9 The Hyper-V host page

7. On the Disk page, in the Path box, click Browse to select the HV-NODE1 to where the VHD file is copied on the destination server running Hyper-V, and then click Next;

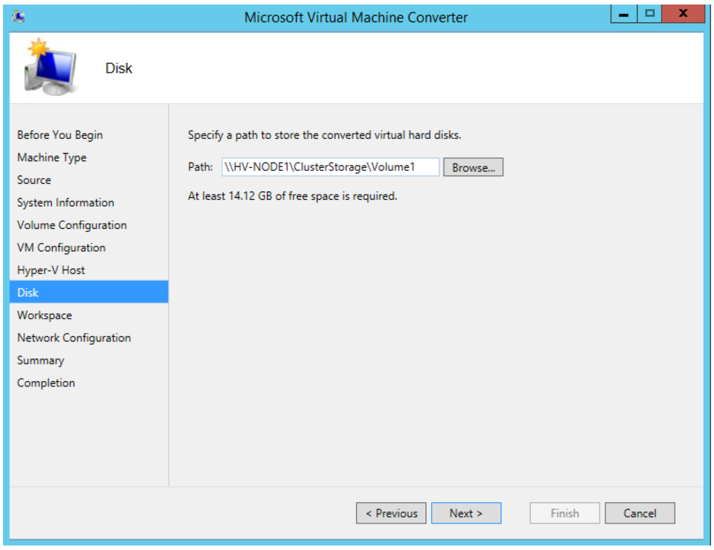

Figure 6.10 The disk page

8. On the Workspace page, click Browse to select the path to a workspace folder where the converted VHDs can be temporarily stored, and then click Next;

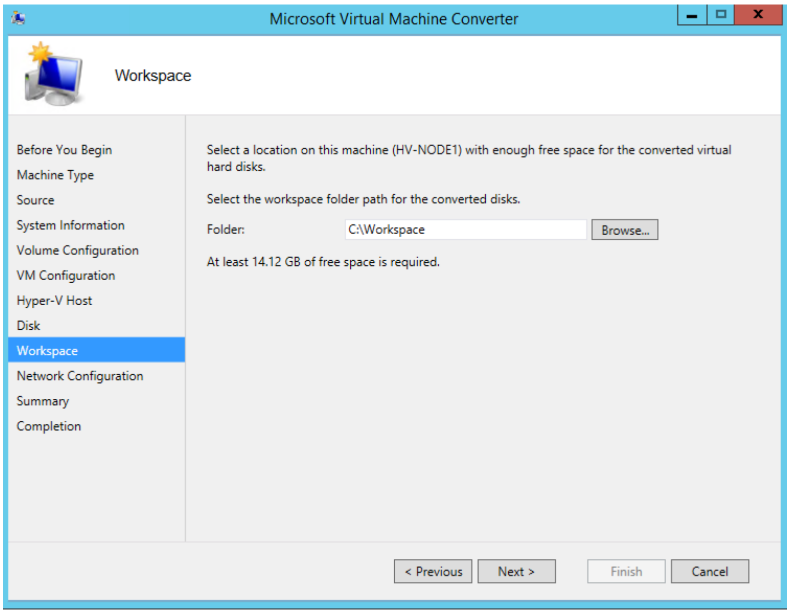

Figure 6.11 The workspace page

9. On the Network Configuration page, select any virtual switches to use for the target virtual machine, and then click Next;

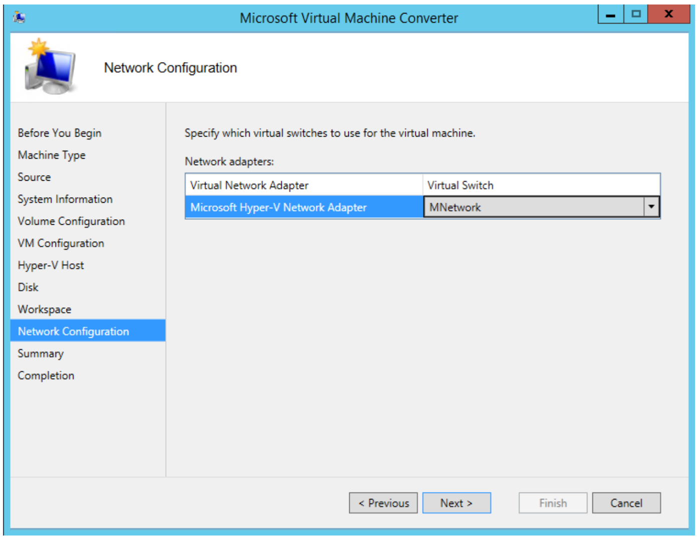

Figure 6.12 The network configuration page

10. On the Summary page, review the details, and then click Finish to complete the conversion;

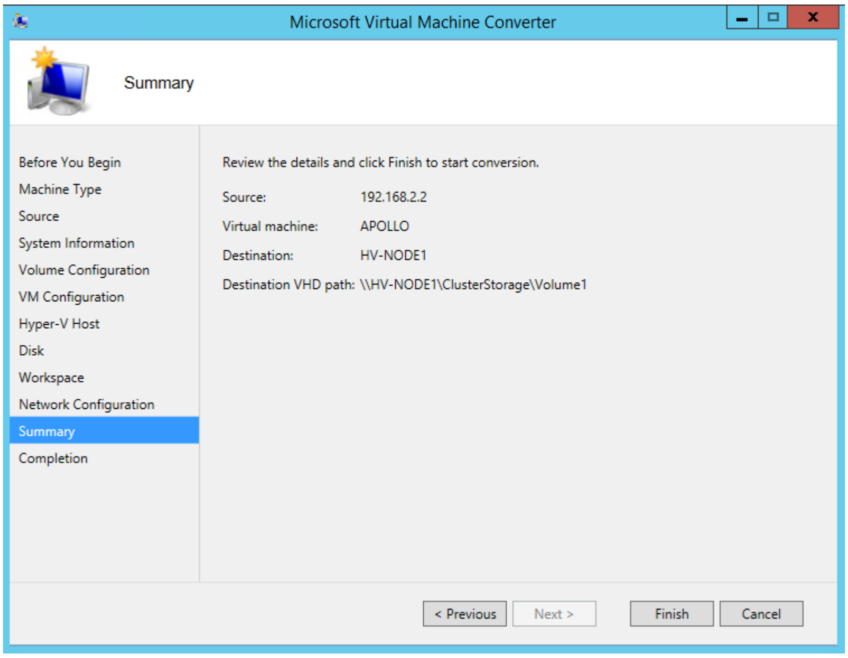

Figure 6.13 The summary page

11. When the virtual machine conversion has finished successfully, the Completion page appears;
12. Open Hyper-V Manager on hv-node1 and Apollo is converted to virtual machine as in Figure 6.14;

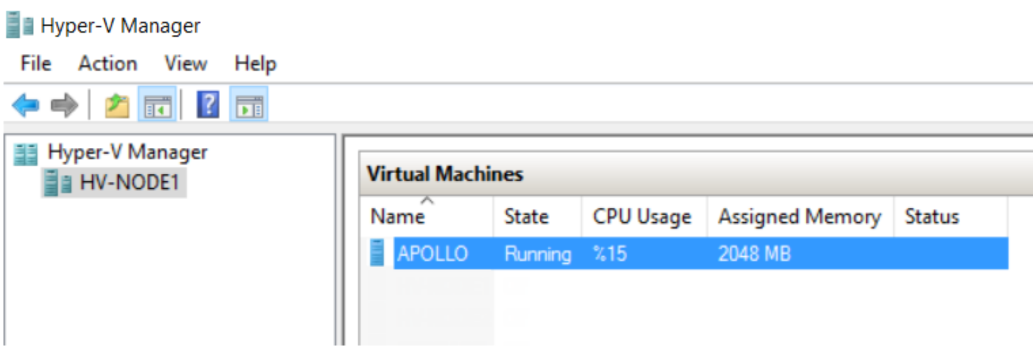

Figure 6.14 The Apollo virtual machine

13. It is checked that APOLLO folder is created under ClusterStorage\Volume1 in Figure 6.15 and Virtual Hard Disks and Virtual Machines folders are created under ClusterStorage\Volume1\APOLLO in Figure 6.16;

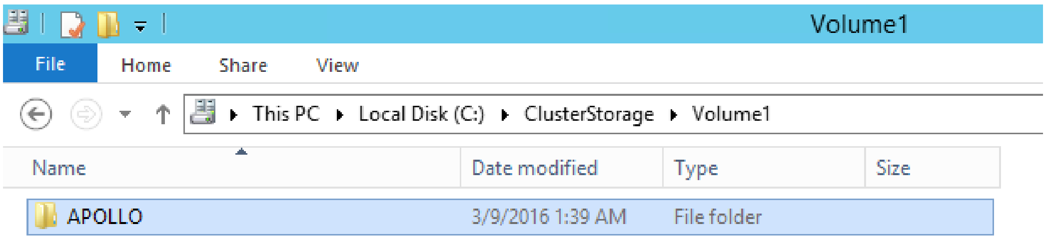

Figure 6.15 The Apollo folder

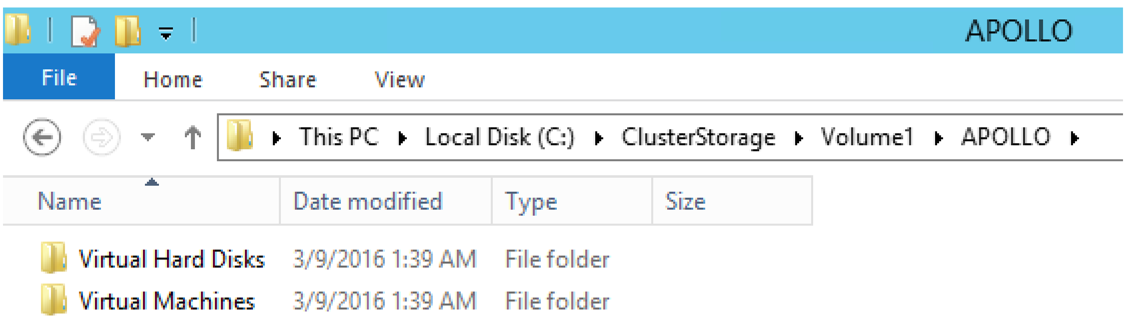

Figure 6.16 The Apollo folder details

14. Repeat the steps between 1 and 13 in order to convert all the other physical machines to machine to Hyper-V.

In order to prevent virtual machine sprawl, new example virtual server request form is created as follows in Form 1 and Form 2.

Form 1. Virtual server request form sample:

Information of Request Owner	
Name and Surname	Onur Milli
Department Name	Information System
Request Date	21.8.2015
Title	System and Network Expert
Manager	Prof. Dr. Ziya Aktaş
Room Number	6432
E-mail address	onurmilli@onurmilli.com

Form 2. Information about the requested virtual server form sample:

Information of Request Virtual Server	
Server Hostname	APOLLO
Server RAM (GB)	2 GB
Server vCPU	1
Server Storage (GB)	60
Server Operating System	Windows Server Standard 2012 R2
Service Level	7/24
IP Address	192.168.2.2
Backup	Yes

All virtual machine requests will be accepted after evaluation by responsible department. It is a best way to prevent VM sprawl.

6.8. Monitor

It is assumed that all physical servers are converted into virtual servers. In order to monitor and manage virtual machines, one of the alternative of monitoring tool is System Center 2012 R2 Microsoft Virtual Machine Manager.

Virtual Machine Manager permits one to put together and manage the virtualization hosts, host clusters, and infrastructure resources used to produce and deploy virtual machines and services to private clouds. These infrastructure resources embody host teams, networking resources, storage resources, and library servers and shares. Different resources represent the fabric from that private clouds are often deployed and managed through the System Center family of product [Tulloch, M. and Perriman, S., 2013].

All Virtual Machine Manager Roles and SQL Database Server are installed on a management server mentioned in 6.5.4 and named as mn.organization.com:

- Windows Server 2012 R2 is installed as the base operating system for Virtual Machine Manager;
- SQL Server 2012 with SP1 is installed for all SQL database services.

In order to install Virtual Machine Manager, follow the steps:

1. In AD (Active Directory), create the following accounts and groups, according to one's naming convention:

- ORGANIZATION\vmmsa (Service account for Virtual Machine Manager);
- ORGANIZATION\vmmadm (Runas account for Virtual Machine Manager);
- ORGANIZATION\sqlsa (Sql service account);
- ORGANIZATION\vmmadms (Administrators security group for Virtual Machine Manager).

2. Add the “vmmsa” and “vmmadm” account to the “vmmadms” global group;
3. Add the privileged accounts to the vmmadms group;
4. Install Windows Server 2012 R2 to all server role servers;
5. Install Prerequisites and SQL 2012 with SP1;
6. Install the SCVMM Server and Console;
7. Deploy SCVMM Agent to Hyper-V hosts.

After SCVMM installation is completed, all converted virtual machines on hv-node1 shown as in Figure 6.17.

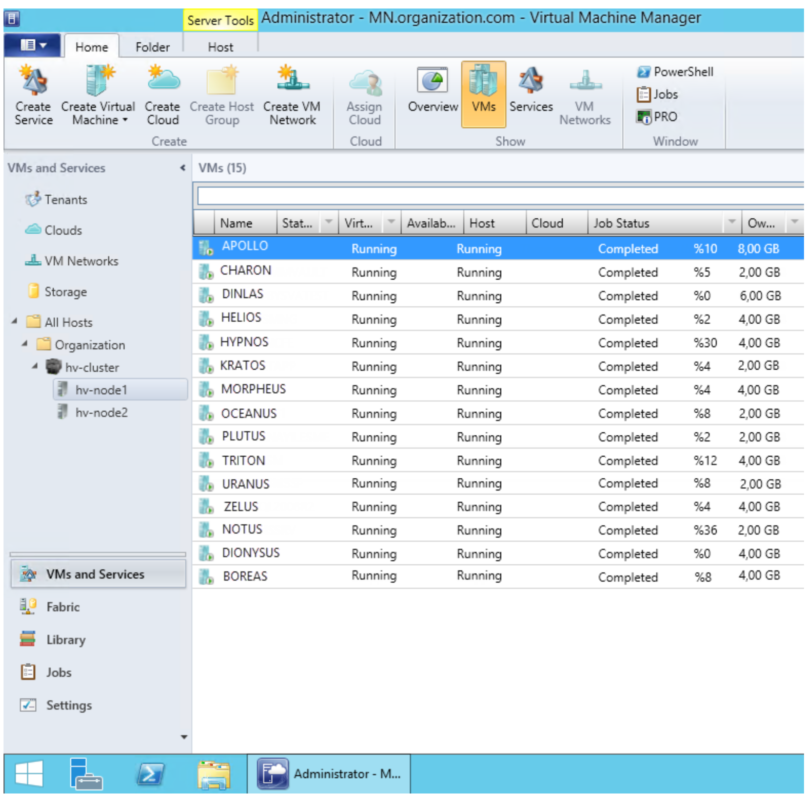

Figure 6.17 The virtual machine manager

Virtual machine usage summary is shown in Figure 6.18. In that figure, performance measures such as CPU performance, memory performance, storage I/O performance and network speed performance are given.

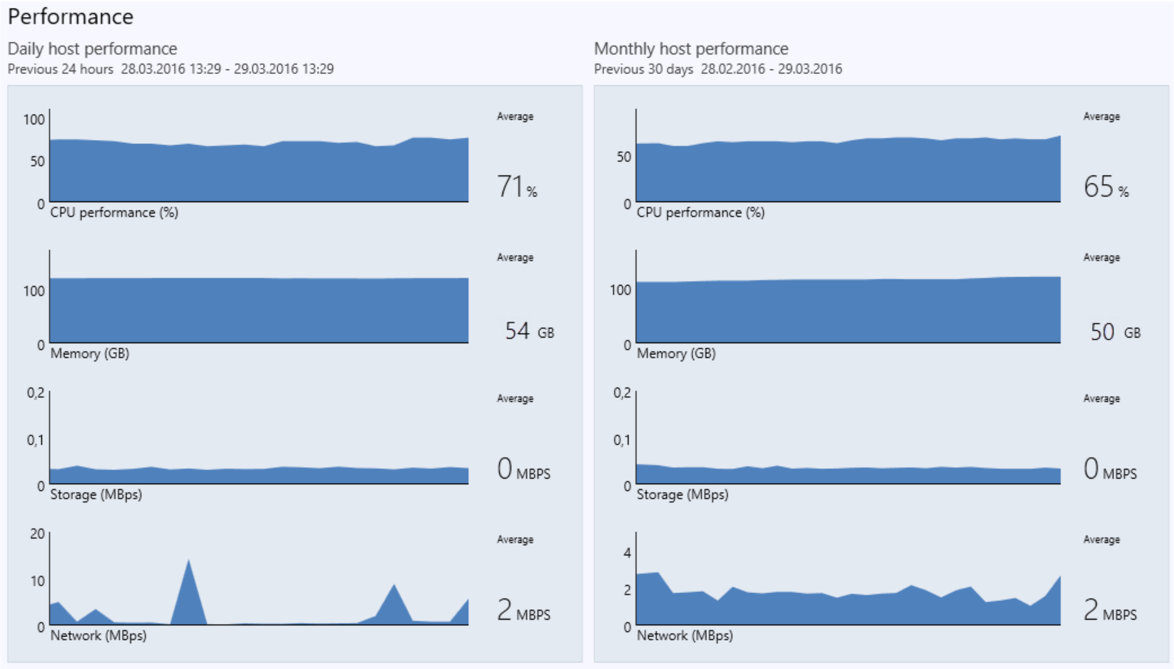

Figure 6.18 The virtual machine manager performance metrics

6.9. Maintain

In virtualization, platform hosts are very important because of hosting lots of guest virtual machines. Therefore, all hosts must be patched and updated regularly. In this thesis, CAU (Cluster-Aware Updating) feature is used in Hyper-V failover cluster environment. In CAU, Hyper-V host servers are patched with minimal downtime while the updates are being deployed.

It is assumed that all virtual machines are working on hv-node1. When one wants to patch a Hyper-V cluster node, one needs to do is shutting down the host. For example, follow the steps in order to update hv-node1 node below:

1. Open Failover Cluster Manager in hv-node1;
2. Click Nodes;

3. In the left pane, right click hv-node1;
4. In the opened window, click Pause and then click Drain Roles. All virtual machines will be migrated to hv-node2;
5. In hv-node1 install all the necessary patches and windows updates. After installation, restart the hv-node1;
6. After login to hv-node1, repeat the steps between 1 and 3;
7. In the opened windows, click Resume and then click Fail Roles back. All virtual machines will be migrated to hv-node1 again;
8. Repeat all the steps again for hv-node2.

6.10. Backup

One of the traditional maintenance tasks is backing up. In order to backup virtual machines, usually backup agent is installed on Hyper-V host machines. There are many virtualization backup solutions in the world. For example, one of them is DPM (Data Protection Manager) that is a part of System Center 2012 R2 family product. However, in this thesis, WSB (Windows Server Backup) which is a built-in tool for backing up Hyper-V in Windows Server 2012 R2 is used.

In order to install WSB feature and backup virtual machines, the next steps should be followed:

1. Connect to hv-node1;
2. Open PowerShell command with an administrative account and type “Install -WindowsFeature Windows-Server-Backup;
3. After WSB is installed, click Server Manager > Local Server > Tools, click Windows Server Backup to open the MMC (Microsoft Management Console) snap-in;
4. Right-click Local Backup in the left tree pane; then click Backup Schedule from the context menu to open the Backup Schedule Wizard;
5. Click Next in the Getting Started dialog box;
6. Select Custom for the backup type; then click Next;
7. Click the Add Items button to open the Select Items dialog box;

8. Expand Hyper-V and then choose the VM that one wishes to backup, as shown in Figure 6.19.

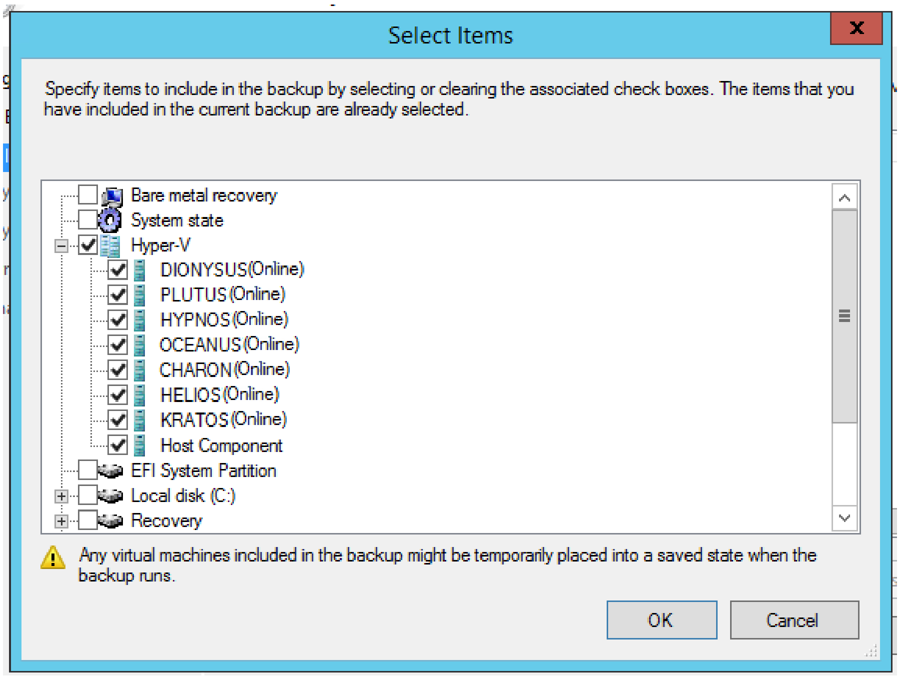

Figure 6.19 Select items

Note that the Host Component is also selected because this protects the virtual switches, resource pools, and Windows Authorization Manager on the host.

9. Back in the Select Items for Backup dialog box, click the Advanced Settings button, and then select the VSS (Volume Shadow Copy Service) Settings tab;
10. Choose “VSS full backup” if one is not using any other backup product to protect one’s VMs or “VSS copy backup” if one is using another backup solution and want to retain the application log files;
11. Click OK to close the dialog box; then click Next to continue;
12. Set the time of day to start the backup;
13. Click Add to move it to the “Scheduled time” list;
14. Click Next;
15. Select the option “Backup to a hard disk that is dedicated for backups (recommended);

16. Click Next;
17. Select the disk to use for the backup volume and click Next;
18. Click Yes to proceed;
19. Review one's selected options; then click Finish to format the disks and schedule the backup.

6.11. Troubleshooting

"Troubleshooting requires a logical and often systematic search for the source of the problem" [Finn et al., 2013]. Different techniques may be used in order to troubleshoot issues. There are two major troubleshooting tools for Windows Failover Clustering and Windows Server 2012 R2 Hyper-V.

The primary troubleshooting tools for Windows Failover Clustering are the cluster log, the Windows system event log, and the logs located under Applications and Service Logs\Microsoft\Windows.

The following logs in Figure 6.20 are used for Windows Failover Clustering:

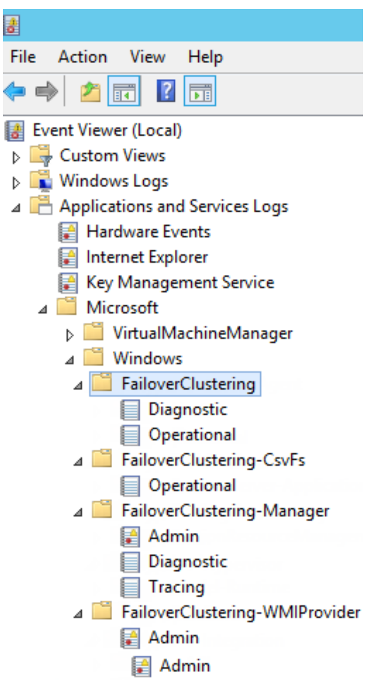

Figure 6.20 Windows failover clustering troubleshooting

The primary troubleshooting tools for Windows Server 2012 R2 Hyper-V are the Windows system event log and the logs located under Applications and Service Logs\Microsoft\Windows.

The following logs given as Figure 6.21 are used for Windows Server 2012 R2 Hyper-V:

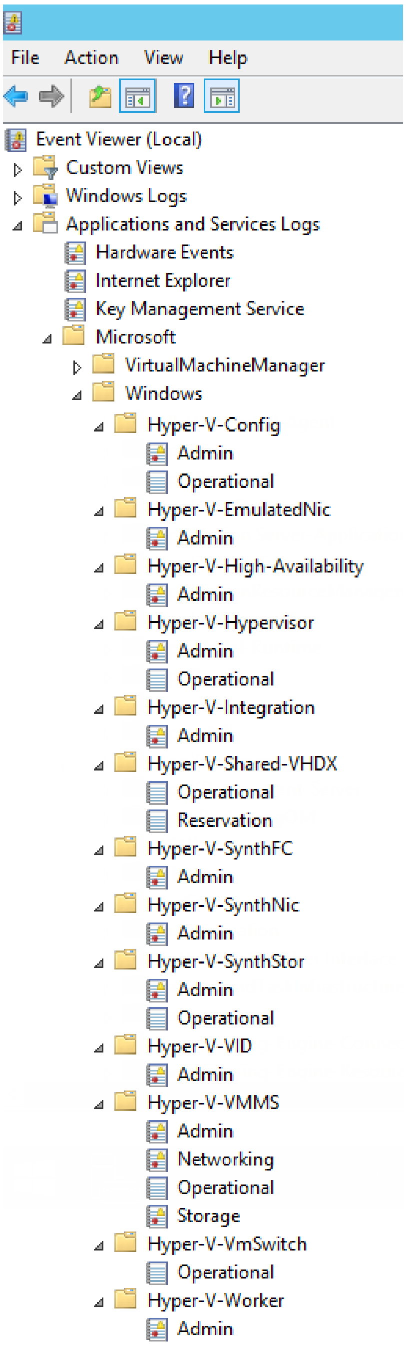

Figure 6.21 Windows server 2012 R2 Hyper-V troubleshooting

7. SUMMARY AND CONCLUSIONS

7.1. Summary

Main objective of this thesis is an application of server virtualization for an enterprise information system. Another objective of this study has been to demonstrate server virtualization advantages. In order to clearly explain the steps of virtualization process, a virtual information system is defined in the thesis. Considering the wide spectrum of the topic, some of the technical details of every subject is not given in the thesis.

The thesis started with a brief Introduction and continued with server virtualization relevant topics. Fundamental concepts of virtualization are explained in a sufficient detail first. After that, a fictitious enterprise information system for a server virtualization application is defined. Existing life cycles for server virtualization followed. Then, a server virtualization life cycle is proposed for a fictitious enterprise information system for security reasons. Lastly, proposed life cycle is applied on a fictitious enterprise information system and all the steps of the development are placed as Appendices A-G.

7.2. Discussion of the Results and Conclusions

Because of this study, management and resource costs decreased by consolidation of 15 servers into 2 servers. Business flexibility increased by creating virtual servers in minutes. Security improved by decreasing number of physical servers and configuring hosts servers by security baseline checklists. It is also reduced downtime by live migration and high availability features of virtualization technologies.

The major purpose and key contribution of this thesis is to provide a first step of guidance for organizations that are considering server virtualization.

Because of this thesis, one may conclude that one of the best ways to improve resource utilization, and at the same time simplify data center management, is through server virtualization. Data centers today use virtualization techniques to make abstraction of the physical hardware, create large aggregated pools of logical resources consisting of CPUs, memory, disks, file storage, applications and networking.

7.3. Extension of the Thesis

“One of the primary technology innovations that influenced and inspired key distinguishing features and aspects of cloud computing is server virtualization” [Erl et al., 2013].

“Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. Cloud model is composed of five essential characteristics such as on-demand self-service, broad network access, resource pooling, rapid elasticity and measured service, three service models such as SaaS (Software as a Service), PaaS (Platform as a Service) and IaaS (Infrastructure as a Service), and four deployment models such as private cloud, community cloud, public cloud and hybrid cloud” [Mell and Grance, 2011].

As a next study to this thesis, Private Cloud can be implemented with an enterprise server virtualization. Private clouds enable an organization to use cloud computing technology as a means of centralizing access to IT resources by different parts, locations, or departments of the organization.

“Many clouds are built on virtualized infrastructure technology. Cloud computing originated as a new way to deliver IT services by providing a customer interface to automated self-service catalogs of standard services, and by using auto scaling to respond to increasing or decreasing user demand. From an IT perspective, a private cloud offers the key advantages of speed, agility, and efficiency while maintaining control of sensitive workloads” [Intel IT Center, 2013].

LIST OF REFERENCES

A Virtualization Roadmap from Dell and Intel, Last Access: http://marketing.dell.com/Global/FileLib/V-bundle/virt-dell-intel-roadmap.pdf, 2016.

Bailey, M., The Economics of Virtualization: Moving Toward an Application-Based Cost Model, IDC, November 2009.

Bittman, T.J., Dawson, P., Weiss, G.J., Magic Quadrant for x86 Server Virtualization Infrastructure, May 2010.

Bittman, T.J., Dawson, P., Warrilow, M., Magic Quadrant for x86 Server Virtualization Infrastructure, July 2015.

Commvault, https://en.wikipedia.org/wiki/Commvault, Last Access: 2016.

Erl, T., Ricardo P., and Zaigham M., Cloud Computing: Concepts, Technology, Architecture, 1st, Prentice Hall, 2013.

Finn, A., Lownds, P., Luescher, M., Flynn, D., Windows Server 2012 Hyper-V Installation and Configuration Guide, Sybex, March 2013.

HPE MSA 1040 SAN Storage, http://www8.hp.com/us/en/products/disk-storage/product-detail.html?oid=6923856, Last Access: 2016.

Hyper-V Virtual Switch Overview, https://technet.microsoft.com/en-us/library/jj945275.aspx, Last Access: 2016.

Intel IT Center, Planning Guide, Virtualization and Cloud Computing, August 2013.

Mell, P., Grance T., The NIST Definition of Cloud Computing, NIST U.S. Department of Commerce, September 2011.

Microsoft Technet(a), https://technet.microsoft.com/en-us/library/hh831579.aspx, Last Access: 2016.

Microsoft Technet(b), https://technet.microsoft.com/en-us/library/mt169373.aspx, Last Access: 2016.

Microsoft Technet(c), https://technet.microsoft.com/en-us/library/dd759255.aspx, Last Access: 2016.

Microsoft Technet(d), https://technet.microsoft.com/en-us/library/hh831435.aspx, Last Access: 2016.

Microsoft Virtual Machine Converter 3.0, https://technet.microsoft.com/en-us/library/dn873998.aspx, Last Access: 2016.

Microsoft, http://windows.microsoft.com/en-us/windows-vista/what-is-internet-small-computer-system-interface-iscsi, Last Access: 2016.

Panek, W., MCTS Windows Server Virtualization Configuration Study Guide: Exam 70-652, June 2009.

Popek, G.J., and Goldberg, R.P., Formal Requirements for Virtualizable Third Generation Architectures, Communication of the ACM, p.412-121, 1974.

Portnoy, M., Virtualization Essentials, 1st, Sybex, 2012.

RHI, Glossary of Job Descriptions for Information Technology, Robert Half International (RHI), 2008.

Siebert, E., VMware VI3 Implementation and Administration, Prentice Hall, May 2009.

Silva, F.D., Maverick* Research: Peer-to-Peer Sharing of Excess IT Resources Puts Money in the Bank, September 2013.

Tulloch, M. and Perriman, S., Introducing Microsoft System Center 2012 R2 Technical Overview, 2013.

Veeam, https://en.wikipedia.org/wiki/Veeam_Software, Last Access: 2016.

VMware, Berc, L, Byun, B, Chambers, S, Easton, R, Eshenbach, C, Eubanks, A, Hogan, R, Jennings, J, Levin, G, Nydam, F, Pither, R, Sathyanarayan, N, Scott, C, Wong, E, W and Wroble, E, The Roadmap to Virtual Infrastructure: Practical Implementation Strategies, VMWare, 2006.

What Is PowerShell? https://blogs.technet.microsoft.com/heyscriptingguy/2015/01/02/what-is-powershell/, Last Access: 2016.

Windows Server 2012 Security Baseline, https://technet.microsoft.com/en-us/library/jj898542.aspx, Last Access: 2016.

APPENDICES

APPENDICES

APPENDIX A - Topology Overview

In this section, it will be given more information about topology.

A.1 - Network Infrastructure Overview

Blue network is used for domain, management and monitoring of hosts and iSCSI storage. It works on Local Area Network in this topology.

Green network is used for VMs that run on hosts servers to communicate with each other and other physical servers that are in the Local Area Network. It is named as External Virtual Network. Real physical network card that is bind as External Virtual Network is connected to Local Area Network switches. Both blue network and green network are connected to same switch. If VMs get an IP addresses from 192.168.2.0/24 network that is in the same network with blue one, they can communicate with blue network, other servers and clients on blue network. There is no need to give an IP addresses to hosts servers that are bind to host's physical network cards.

Red Network is used for heartbeat of cluster. This network controls the state of hosts. It is advised that it should be isolated from other networks. In this topology, it is directly connected on hosts servers with cross cable.

Black Network is used for DMZ network. This network separates the internal LAN from untrusted external networks.

Yellow Network is used for data traffic between hosts servers and iSCSI Storage. It must be isolated from other networks for performance and data security. It must be used on dedicated switch for only iSCSI Storage.

A.2 - Other Server Infrastructure Overview

In addition to host servers, there are two other servers. One of them is used for Active Directory and the other one is used for host management in order to control VMs. It will be given more information about two servers below;

First server full name is dc.organizaton.com. Domain name of this server is organization.com. The operating system of this server is Windows Server 2012 R2 Standard. Physical resources are 1 x Quad Core CPU, 6 GB Ram, 1 x NIC and 200 GB Storage. The role of this server is Domain Controller and DNS Server. IP address of this server in the topology is 192.168.2.1/24.

Second server full name is hv-node1.organization.com. Domain name of this server is organization.com. The operating system of this server is Windows Server 2012 R2 Datacenter. Physical resources are given at the beginning of build section. The role of this server is Hyper-V host. Server has eight NICs. NIC 1 is used for Domain Network/Management which is 192.168.2.11/24. NIC 2, 3 are used for VM Network as team. There is no need to give any TCP/IP. VM Network can obtain IP addresses automatically. NIC 4 is used for Heartbeat (Cluster) Network which is 1.1.1.1/8. NIC 5, 6 are used for DMZ Network which is 10.0.0.1/8. NIC 7, 8 are used for iSCSI Network as teamed which is 172.16.1.11/24.

Third server full name is hv-node2.organization.com. Domain name of this server is organization.com. The operating system of this server is Windows Server 2012 R2 Datacenter. Physical resources are given at the beginning of build section. The role of this server is Hyper-V host. Server has eight NICs. NIC 1 is used for Domain Network/Management which is 192.168.2.12/24. NIC 2, 3 are used for VM Network as team. There is no need to give any TCP/IP. VM Network can obtain IP addresses automatically. NIC 4 is used for Heartbeat (Cluster) Network which is 1.1.1.2/8. NIC 5, 6 are used for DMZ Network which is 10.0.0.2/8. NIC 7, 8 are used for iSCSI Network as teamed which is 172.16.1.12/24.

There is also one cluster object. It is named as hv-cluster. Domain name of this object is organization.com. IP address of this object is 192.168.2.20/24.

APPENDIX B - Hyper-V Failover Cluster Installation Requirements

There are some requirements for hv-cluster object, domain controller, hv-node1, and hv-node2.

B.1 - Cluster Object Requirements

Requirements are listed below:

- Active Directory Services must be used;
- Hyper-V nodes must be member of domain;
- It is advised that servers and NIC cards should be supported by Windows Server 2012 R2;
- iSCSI, FC, Fiber Channel over Ethernet (FCoE) and Serial Attached SCSI (SAS) may be used for data traffic. iSCSI is used in this study.

B.2 - Domain Controller Requirements

Requirements are listed below:

- Windows Server 2012 R2 Standard or Datacenter must be used;
- It is advised that servers hardware should be supported by Windows Server 2012 R2;
- Active Directory services role must be installed.

B.3 - Nodes Requirements

Requirements are listed below:

- Windows Server 2012 R2 Standard/Datacenter or Hyper-V Server 2012 R2 must be used;
- Nodes must be domain member;
- It is advised that servers hardware should be supported by Windows Server 2012 R2;

- At least four or five NICs should be used;
- Hardware must support 64bit Processor;
- Processor must support Intel Virtualization Technology (Intel-VT) or AMD Virtualization (AMD-V) and it must be enabled;
- Processor must support Intel Execute Disable (Intel-XD) or AMD No-Execute (AMD-NX) and it must be enabled;
- Hyper-V role must be installed.

APPENDIX C - Hyper-V Installation

Follow the installation steps of Windows Server 2012 R2 for both hv-node1 and hv-node2 host servers. Installation screens are given as Figures C.1 to C.12.

1. Insert Windows Server 2012 media, restart the server and Press any key to start the install;

Figure C.1 Boot installation

2. Select Language, Time and currency format, Keyboard or input method and click Next;

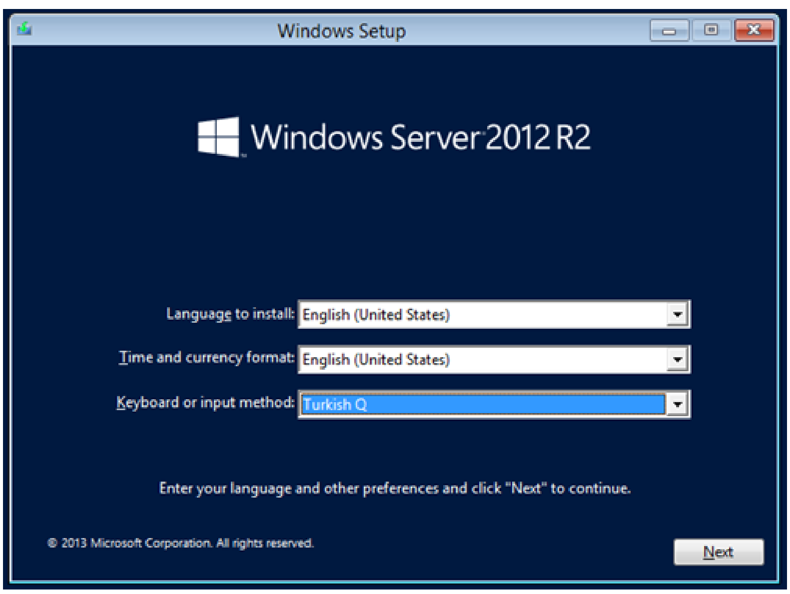

Figure C.2 Windows setup

3. Click Install now;

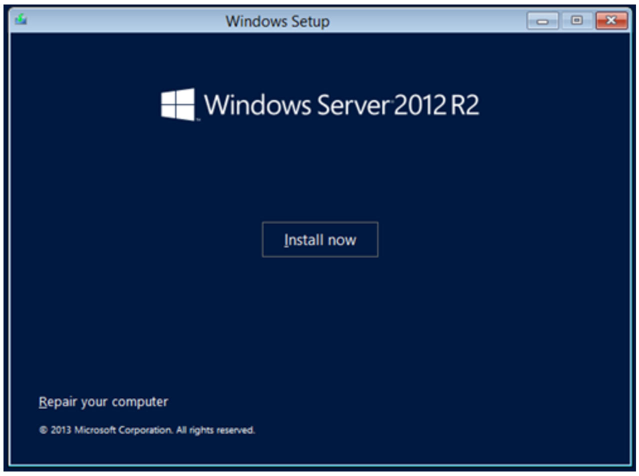

Figure C.3 Windows Server 2012 R2 installation

4. If it is necessary, enter product key. This part doesn't have any screenshot;
5. Select Windows Server 2012 Datacenter (Server with a GUI) and click Next;

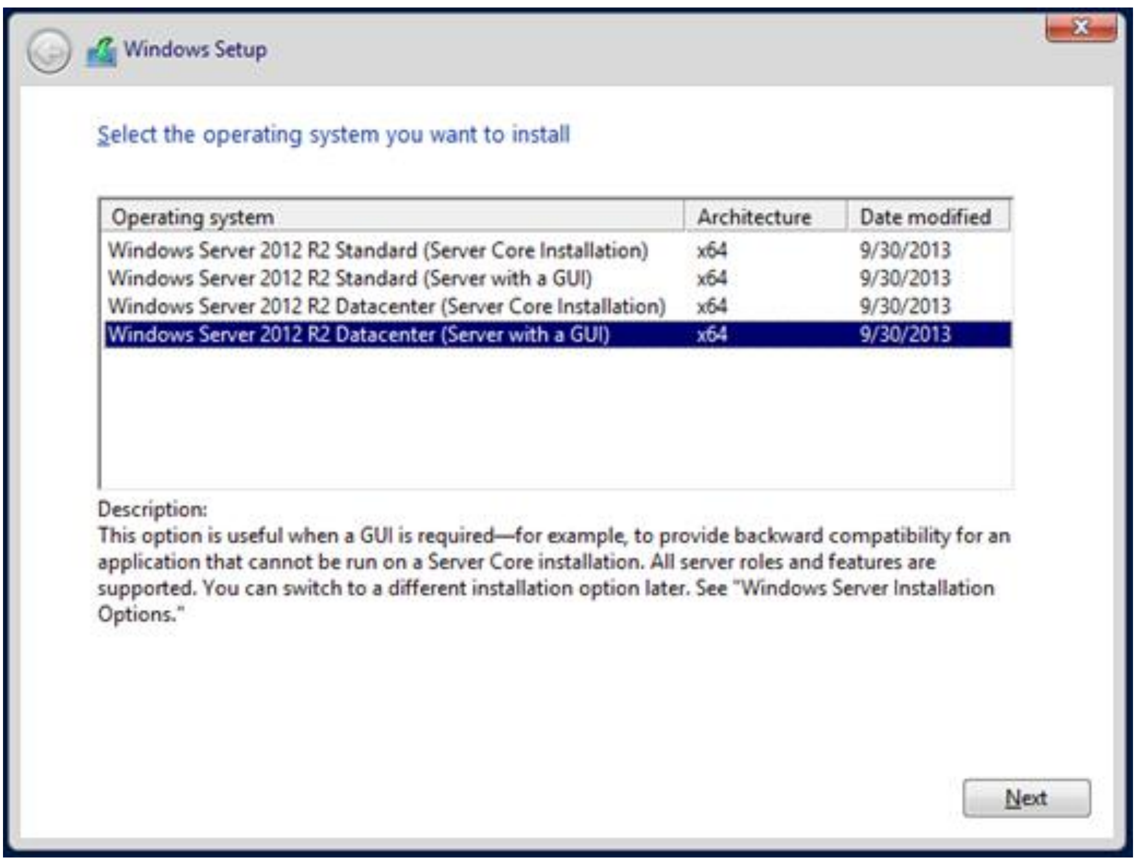

Figure C.4 Windows Server 2012 R2 datacenter with gui

6. Click I accept the license terms;

Figure C.5 License terms

7. Click Custom: Install Windows only (advanced) for new setup;

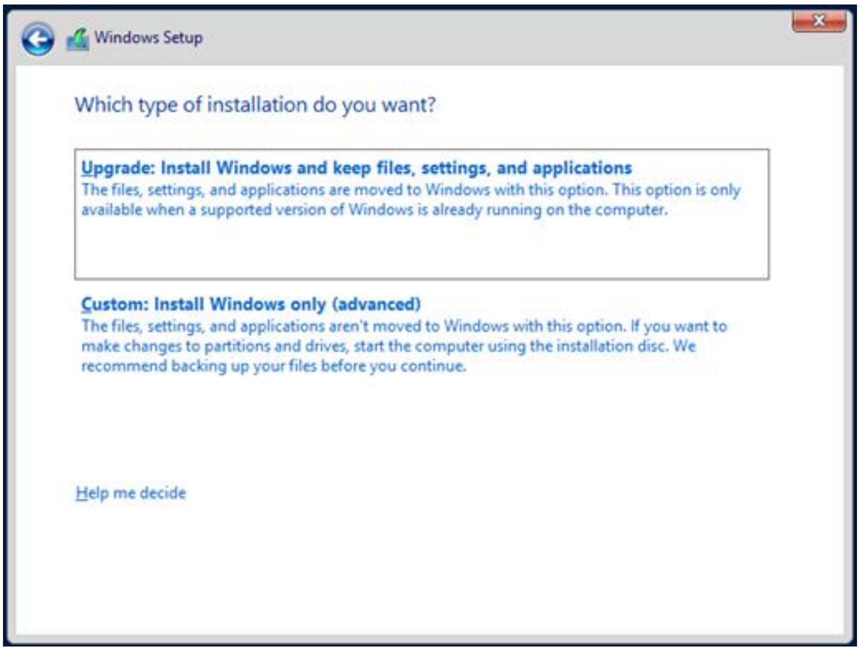

Figure C.6 Type of installation

8. Click Next;

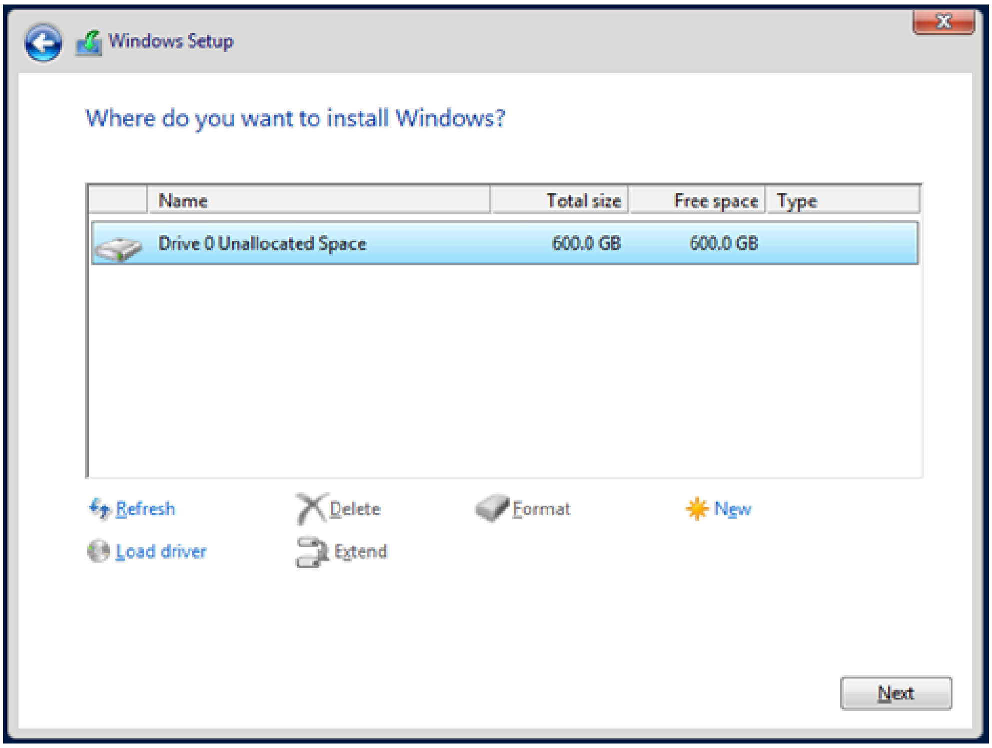

Figure C.7 Disk selection

9. Installation starts and progresses as the screenshot shows;

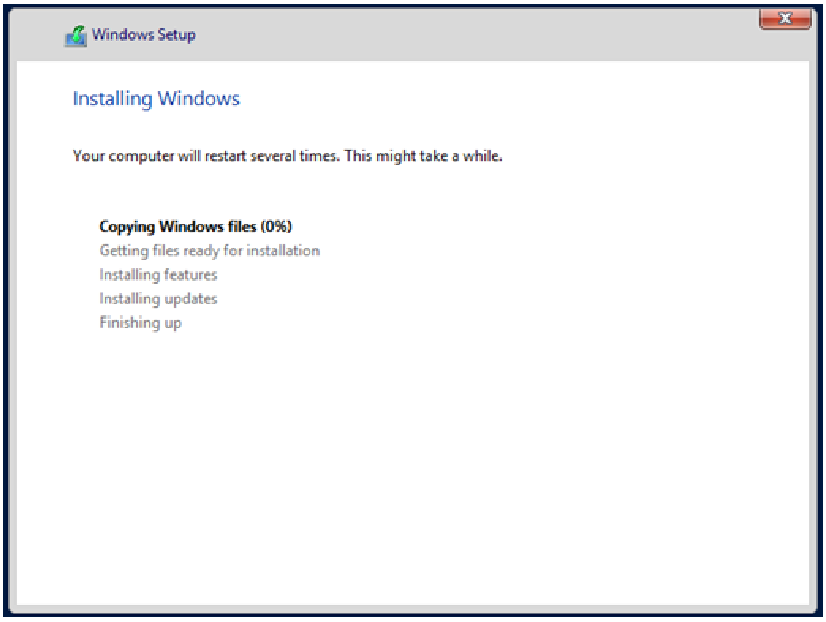

Figure C.8 Installation progress

10. After installation completed, enter password and click Finish;

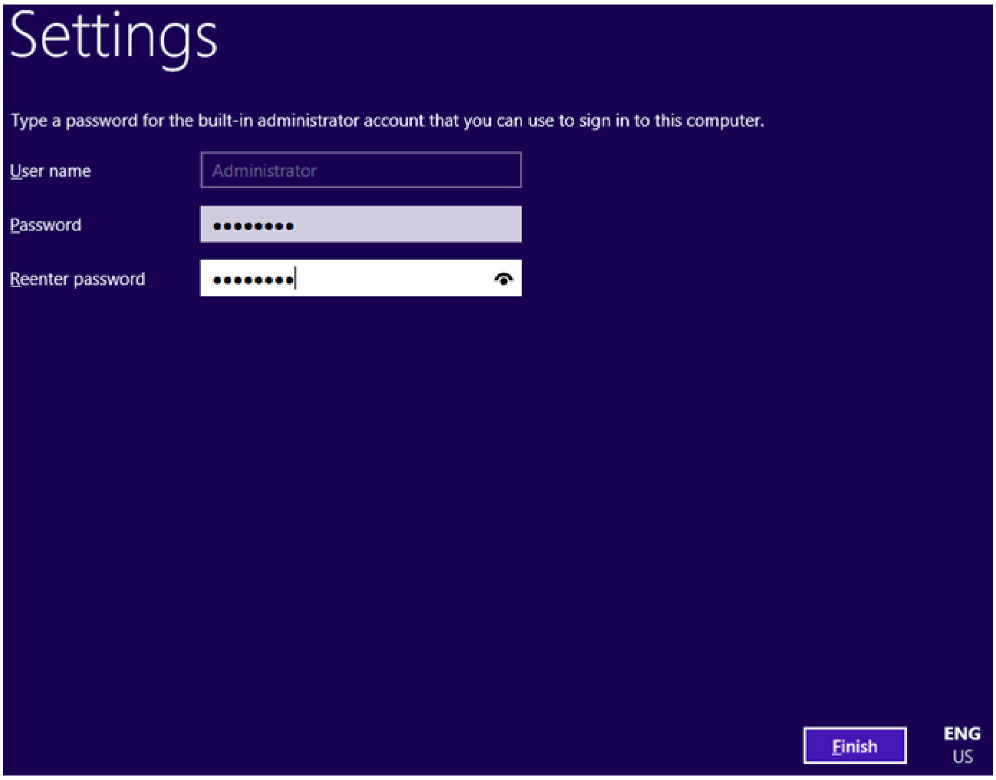

Figure C.9 User name and password

11. Press Ctrl+Alt+Delete to sign in;

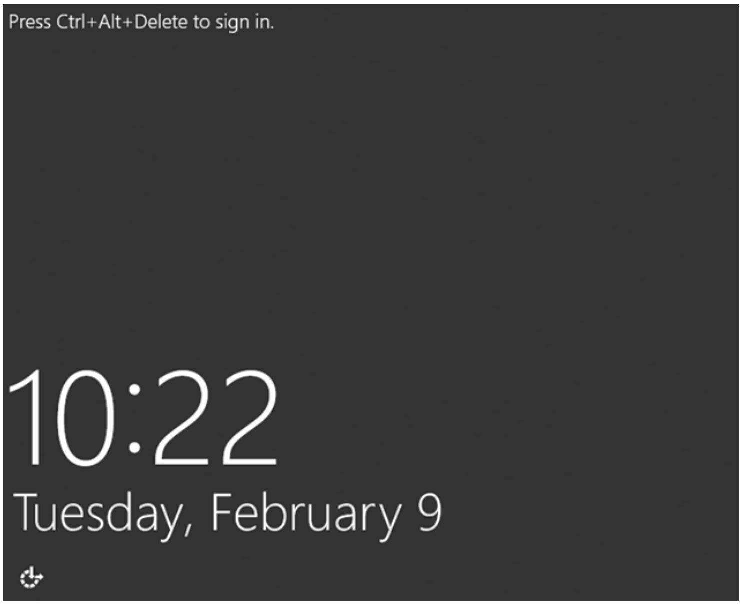

Figure C.10 Log on screen

12. Enter password and press Enter key;

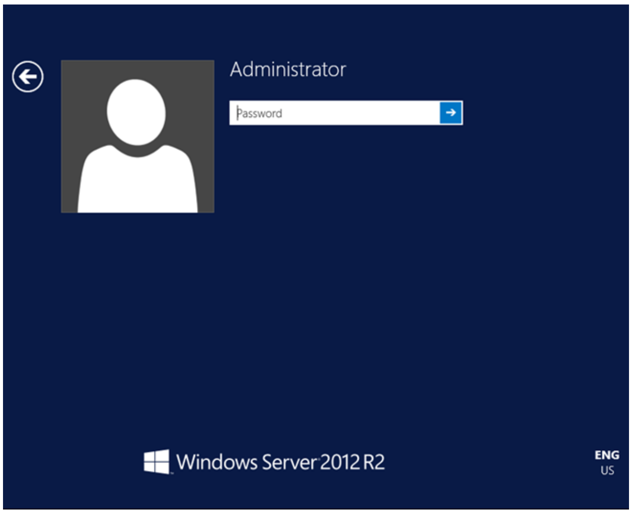

Figure C.11 Password

13. After first login, Server Manager Console opens.

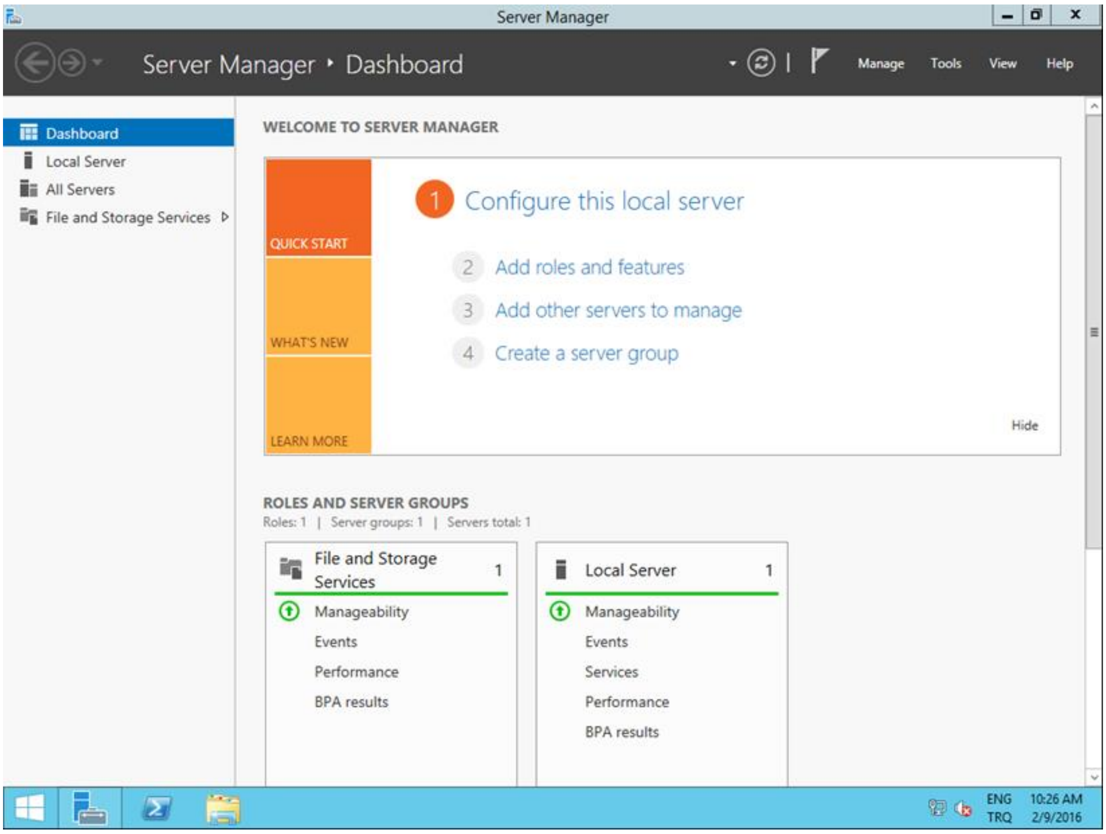

Figure C.12 Server manager console

After installation of hv-node1 and hv-node2 completed, do the following steps for hv-node1 and hv-node2 below:

- Install all the necessary hardware drivers and updates;
- It is advised to select English/US for language and regional settings;
- Control time zone and verify that it is correct;
- Control time and date information;
- Enter necessary IP addresses for Domain/Management communication as blue network in the topology;
- Enter hostnames;
- Join nodes to domain;
- Install all the necessary Windows updates.

APPENDIX D - Hyper-V Hosts Configuration

It is now assumed that domain server is already installed in the system. Hv-node1 and hv-node2 must be joined to organization.com domain. Make sure that all nodes are member of the domain as seen in Figure D.1.

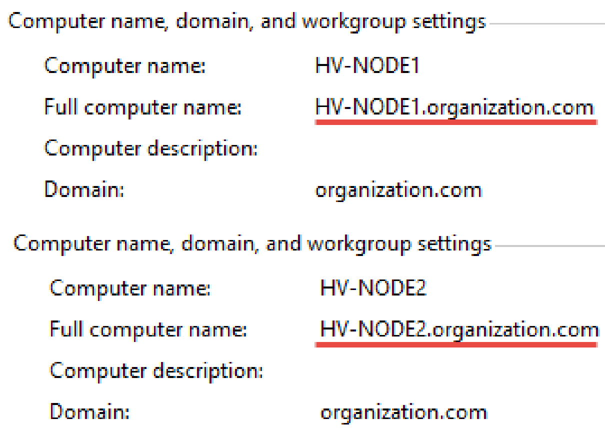

Figure D.1 Domain joined

D.1 - Hyper-V Network Connection Configuration

Hv-node1 and Hv-node2 have eight network cards. One NIC is used for Local Area Network. One NIC is used for Heartbeat. Two NICs are used for VM network. Two NICs are used for iSCSI network and two NICs are used for DMZ network. Eight NICs are shown as Figures D.2.

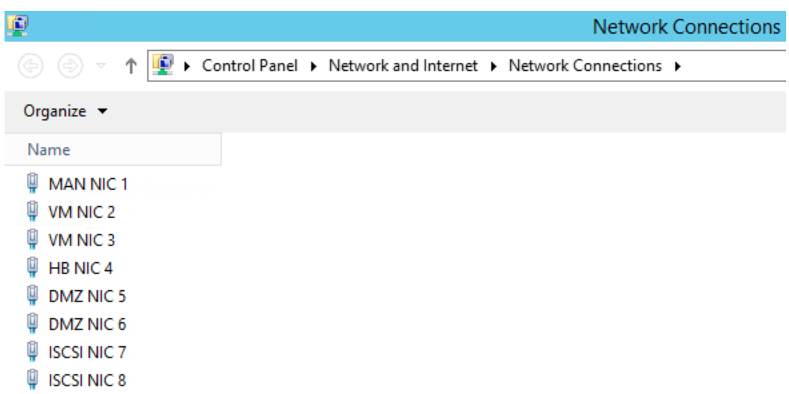

Figure D.2 Network cards

MAN NIC 1 is teamed as MNet. VM NIC 2 and VM NIC 3 are teamed as VMNet. HB NIC 4 is teamed as HBNet. DMZ NIC 5 and DMZ NIC 6 are teamed as DMZNet. ISCSI NIC 7 and ISCSI NIC 8 are teamed as iSCSINet. It is shown how to team NICs below. Follow the Figure D.3 - D.8 to configure TEAM for Hv-node1. After complete the steps, repeat all steps again for the Hv-node2.

1. Open Server Manager and Click Enabled;

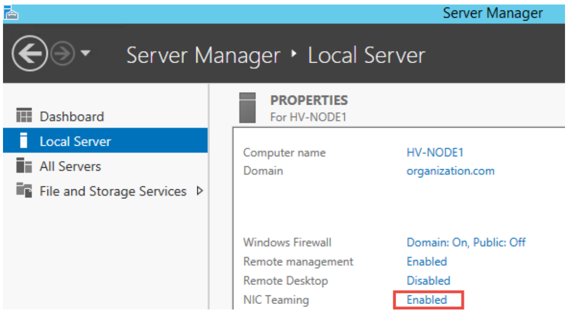

Figure D.3 Nic teaming 1

2. NIC Teaming windows opens;

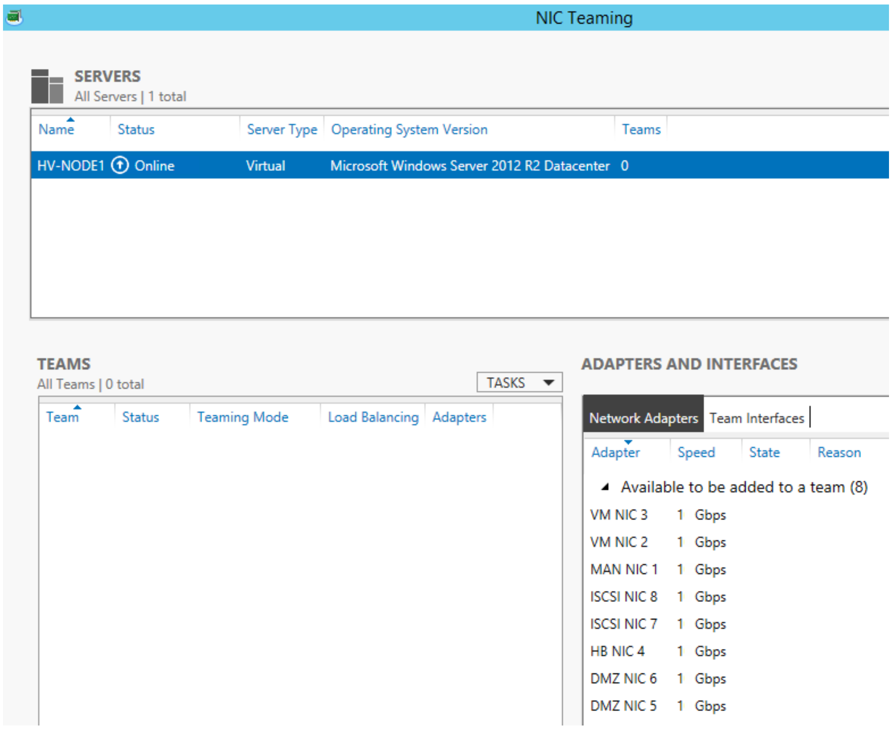

Figure D.4 Nic teaming 2

3. Select VM NIC 2 and VM NIC 3 and right click. Then, click Add to New Team;

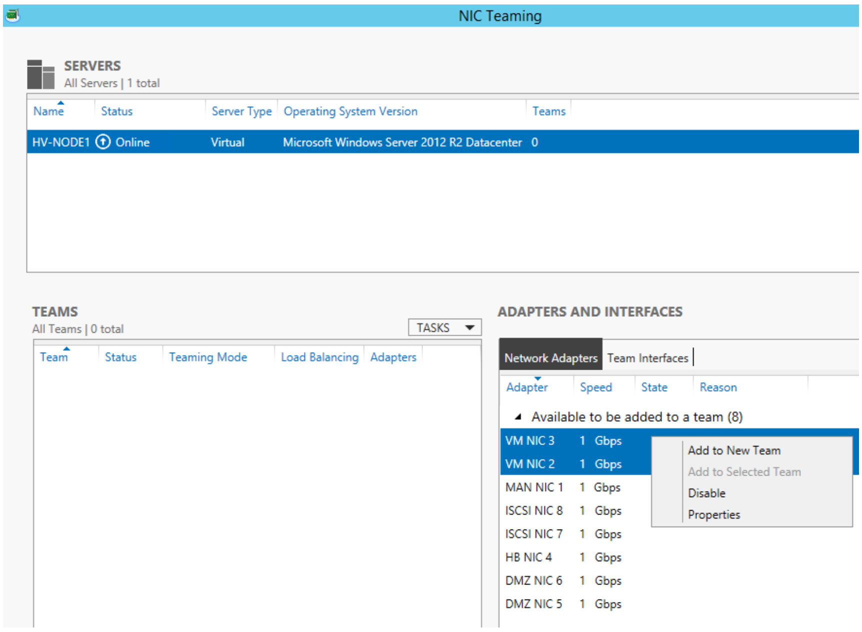

Figure D.5 Nic teaming 3

4. NIC Teaming window opens. Write VMNet as Team name and Click Ok;

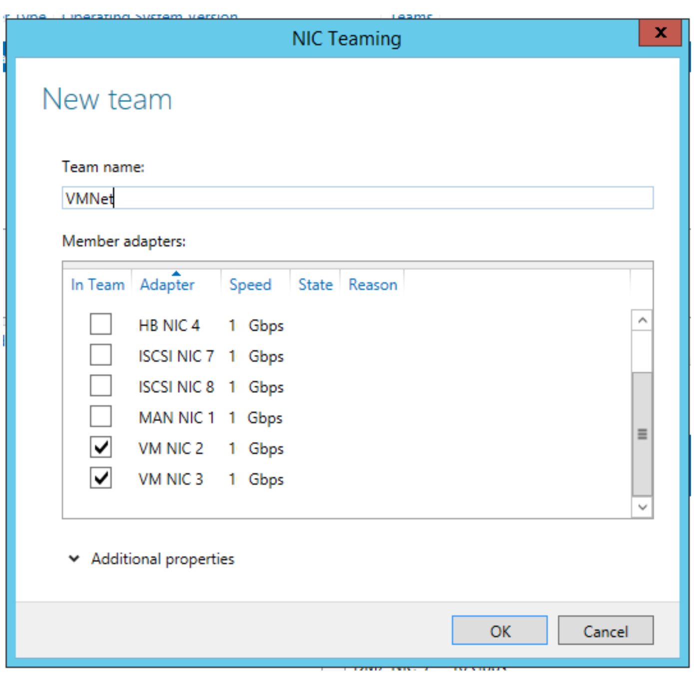

Figure D.6 Nic teaming 4

5. VM NIC 2 and VM NIC 3 are teamed as VMNet;

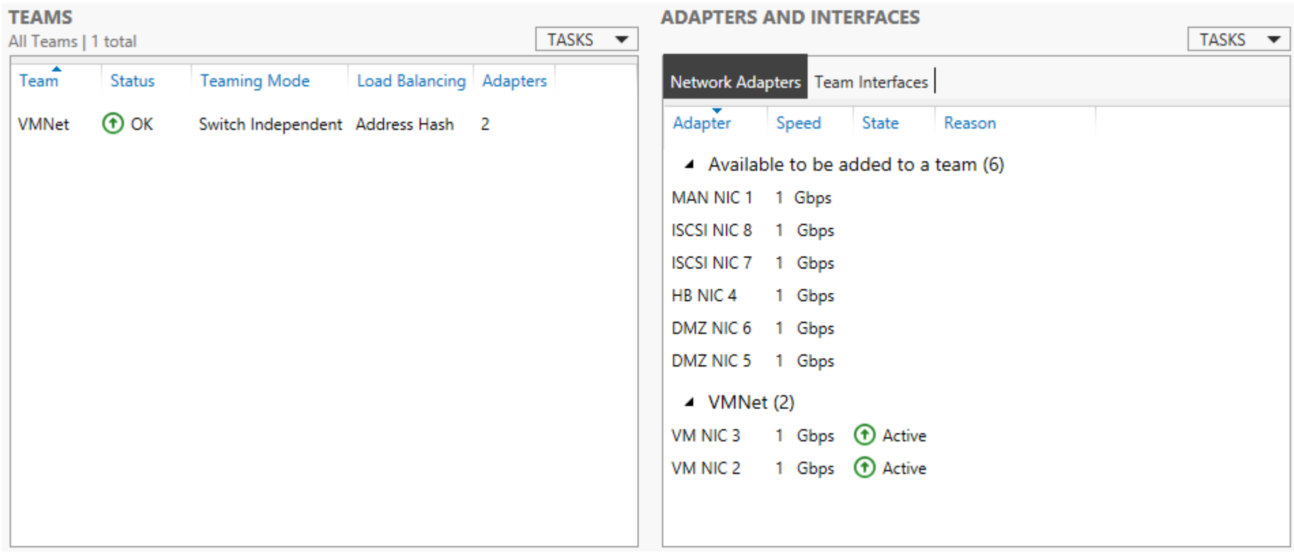

Figure D.7 Nic teaming 5

6. Repeat the same steps for MAN NIC 1 as MNet team, ISCSI NIC 7 and ISCSI 8 as ISCSINet team, HB NIC 4 as HBNet team, DMZ NIC 5 and DMZ NIC 6 as DMZNet team;

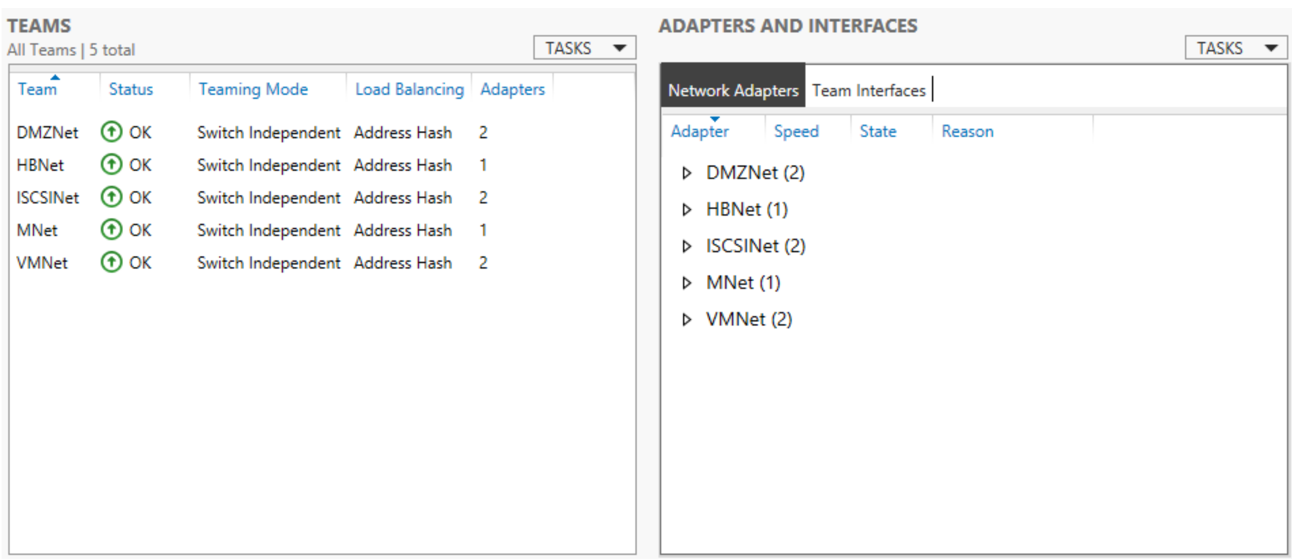

Figure D.8 Teamed nics

7. Configure MNet, VMNew, HBNet, iSCSNet and DMZNet as Figure D.9 - D.13 for hv-node1 and Figure D.14- D.18. for hv-node2.

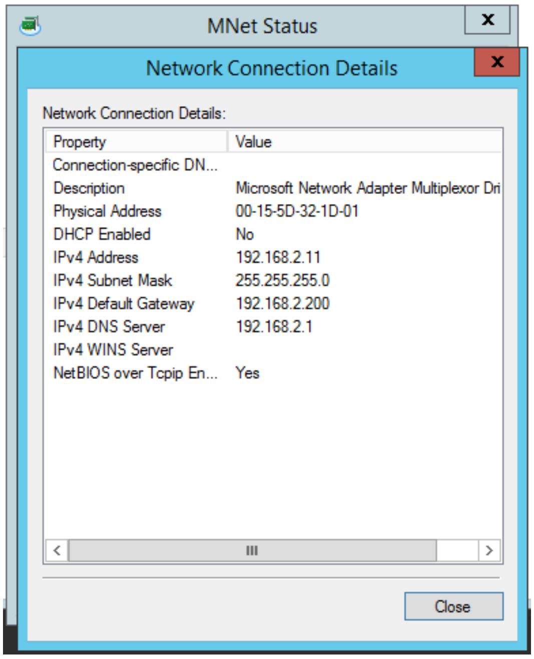

Figure D.9 Node 1
Domain/Management Blue Network

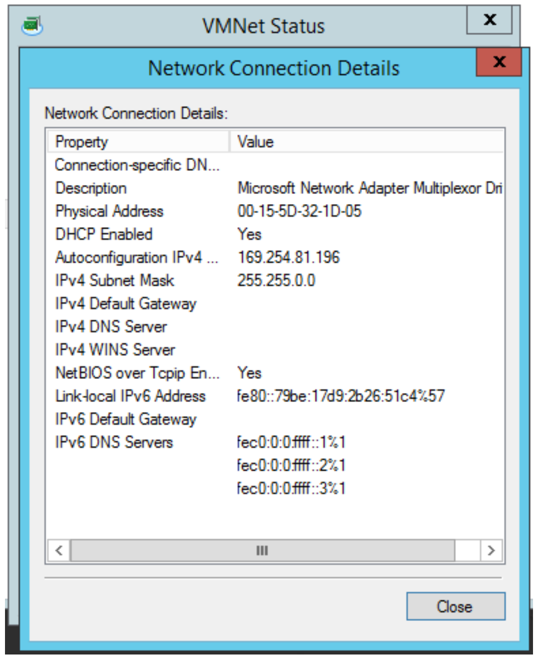

Figure D.10 Node 1
VM Green Network

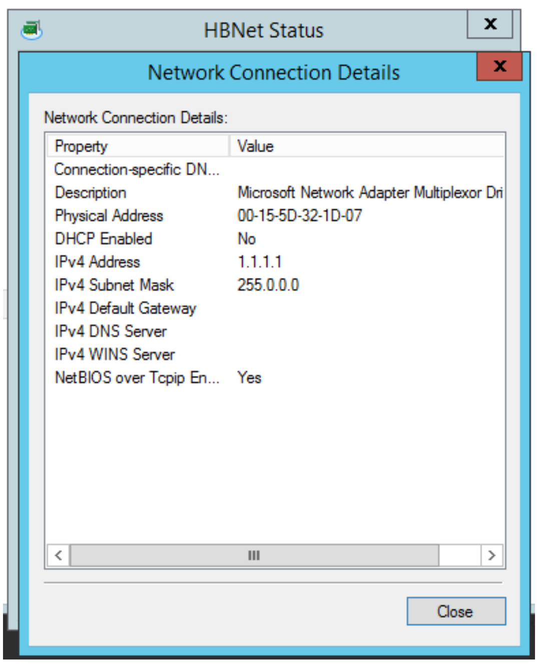

Figure D.11 Node 1 Heartbeat Red
Network

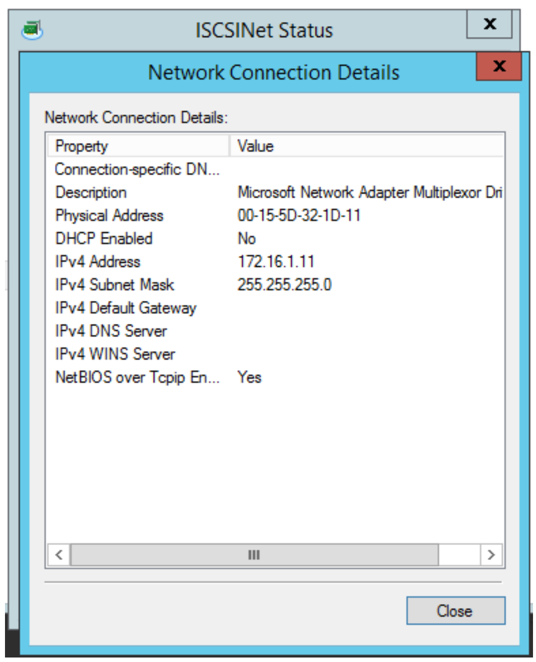

Figure D.12 Node 1 iSCSI Orange
Network

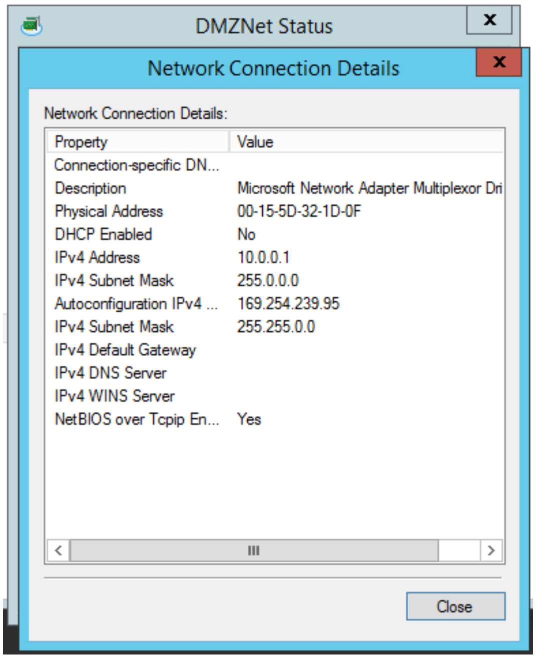

Figure D.13 Node 1 DMZ Black Network

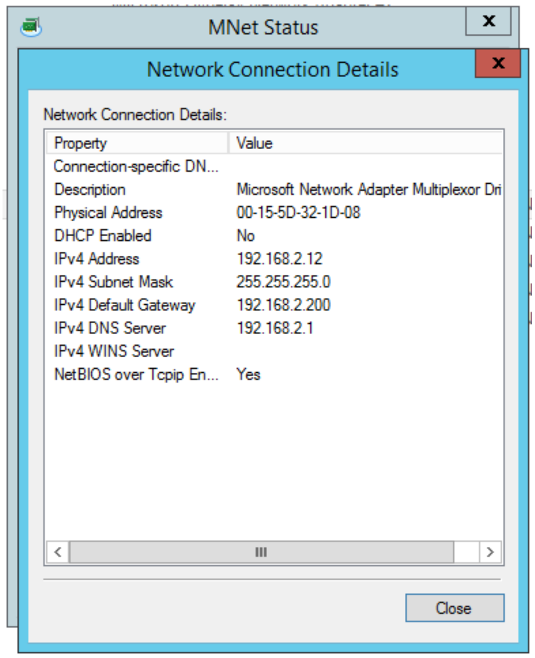

Figure D.14 Node 2 Domain/Management Blue Network

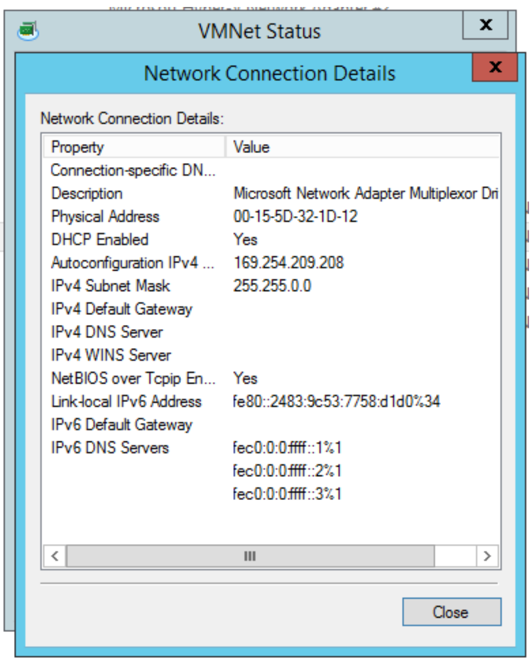

Figure D.15 Node 2 VM Green Network

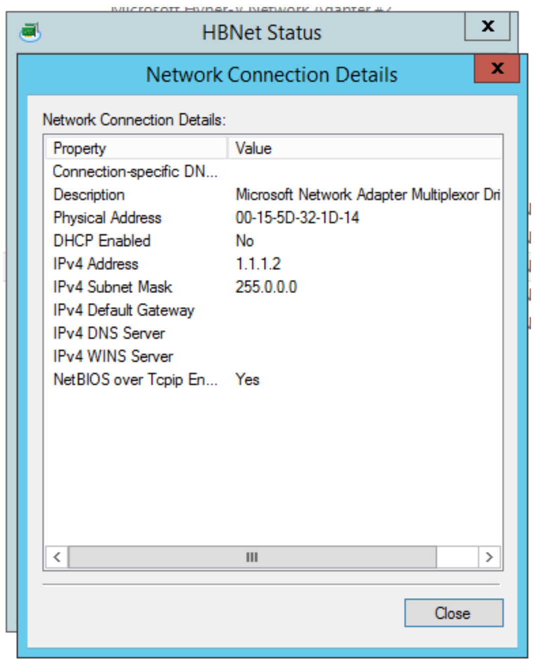

Figure D.16 Node 2 Heartbeat Red Network

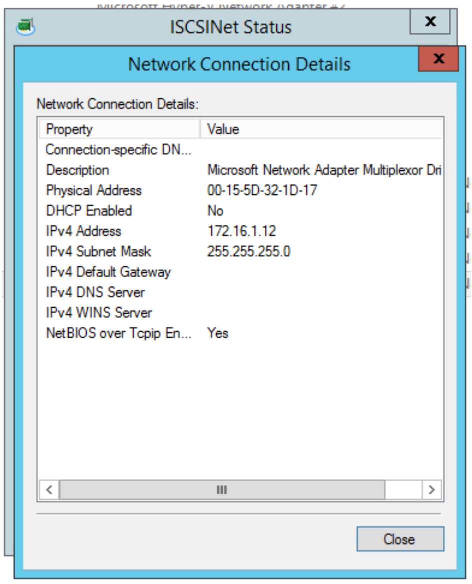

Figure D.17 Node 2 iSCSI Orange Network

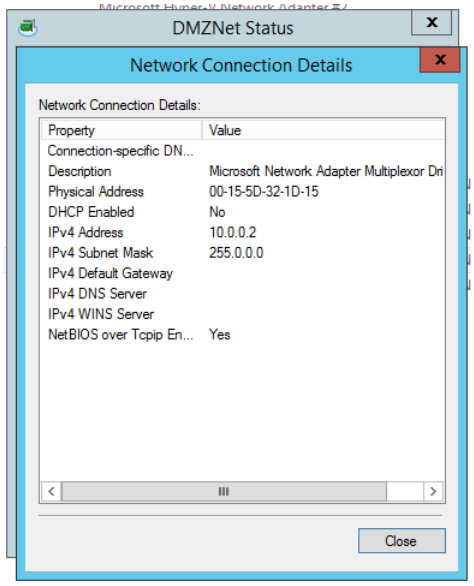

Figure D.18 Node 2 DMZ Black Network

D.2 - Hyper-V Role and Management Tools Installation

Follow the Figure D.19 - D.23 to install Hyper-V Role and Management Tools for both hv-node1 and hv-node1. In the screenshots below, it is shown in hv-node1. After completing all the steps, repeat all steps again for hv-node2.

1. Right Click Windows PowerShell[16] icon and click Run as Administrator;

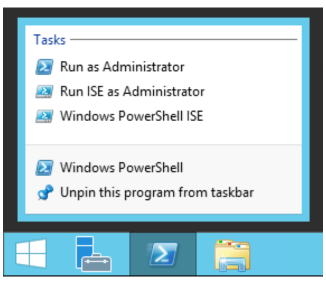

Figure D.19 Windows powershell 1

---

[16] **Windows PowerShell** is an interactive object-oriented command environment with scripting language features that utilizes small programs called cmdlets to simplify configuration, administration, and management of heterogeneous environments in both standalone and networked typologies by utilizing standards-based remoting protocols [What Is PowerShell?, 2016].

Windows Powershell opens. In Figure D.20, type the following command (Enable-WindowsOptionalFeature –Online -FeatureName Microsoft-Hyper-V –All –NoRestart). That will perform the installation of Hyper-V in online mode most do not restart the equipment.

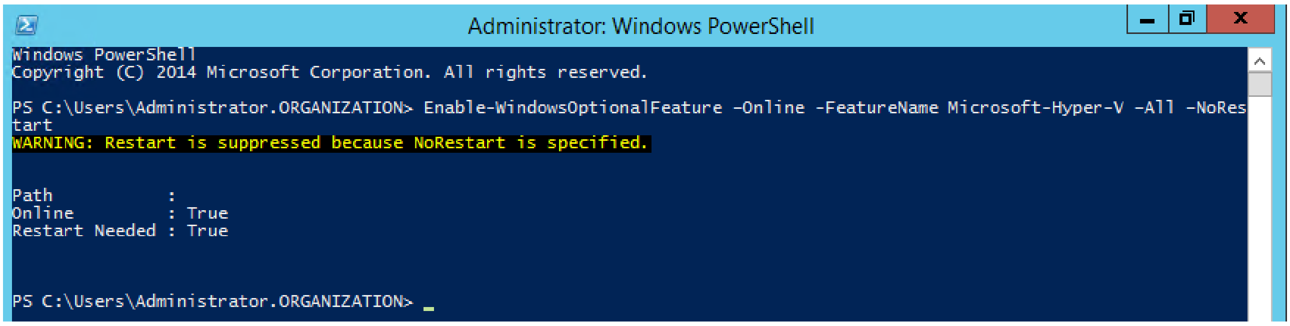

Figure D.20 Windows powershell 2

2. In Figure D.21, type the following command (Install-WindowsFeature RSAT-Hyper-V-Tools –IncludeAllSubFeature). This process will perform the installation of Hyper-V Tools. Server must be restarted after PS command execute;

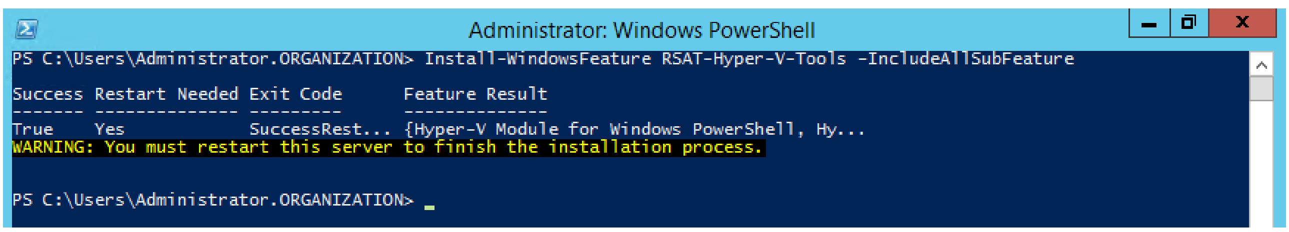

Figure D.21 Windows powershell 3

3. In Figure D.22, after server restarts, verify that Hyper-V Role and Management tools are installed by opening Server Manager. Click Hyper-V Manager to verify installation;

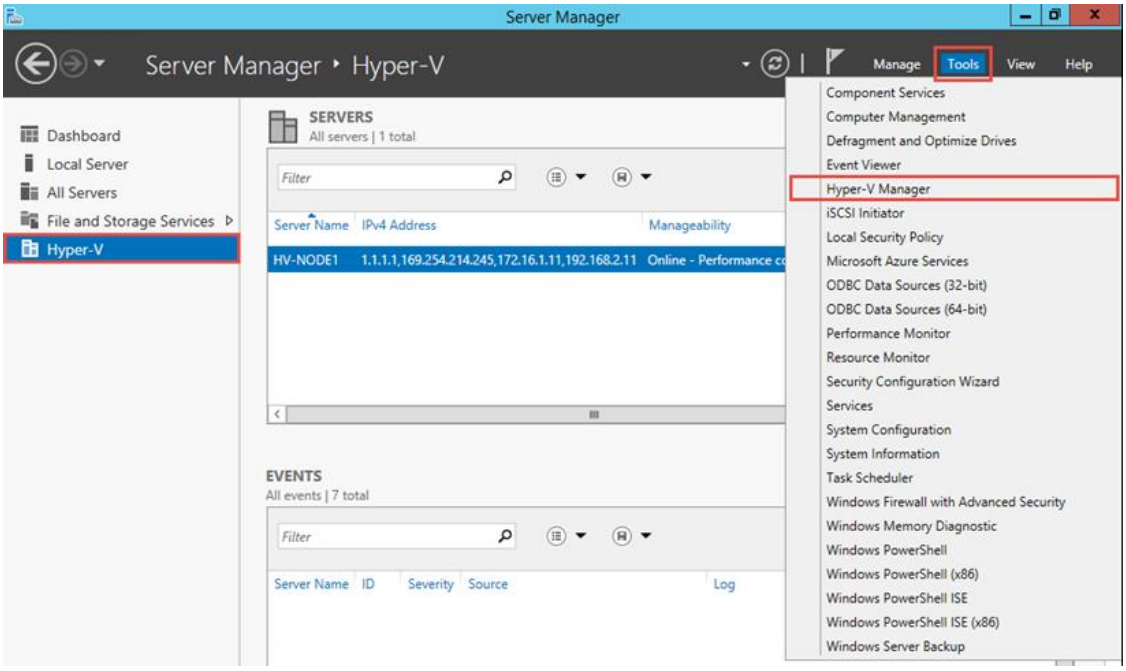

Figure D.22 Server manager

4. In Figure D.23, Hyper-V Manager console is opened and installation is completed successfully.

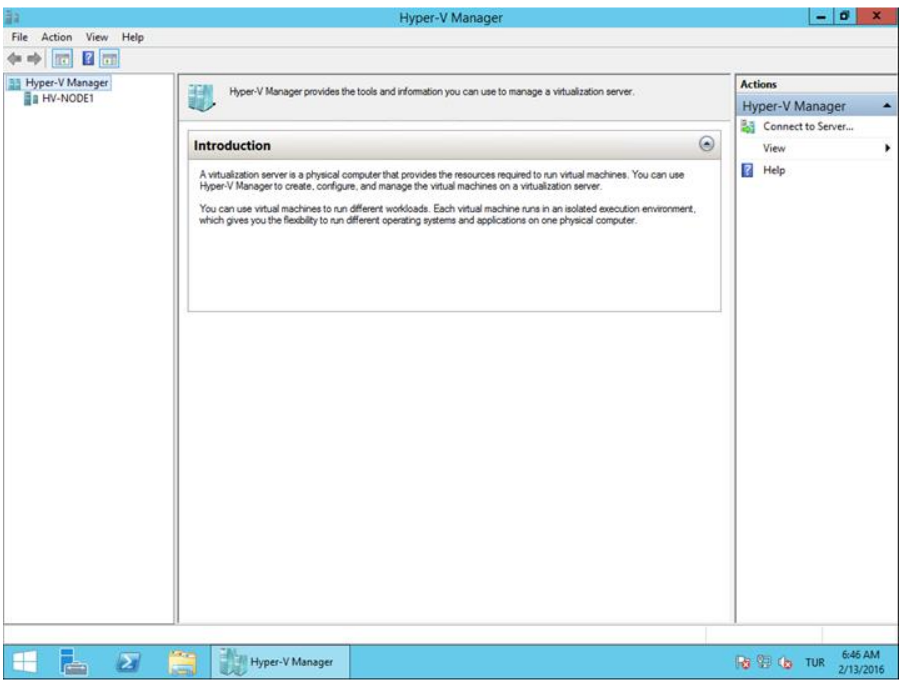

Figure D.23 Hyper-V manager

D.3 - Failover Cluster Component and Management Tools Installation

Follow the Figure D.24-D.34 to install Failover Cluster Component and Management Tools for both hv-node1 and hv-node1. In the screenshots below, it is shown in hv-node1. After complete all the steps, repeat all steps again for hv-node2.

1. In Figure D.24, click Add roles and features;

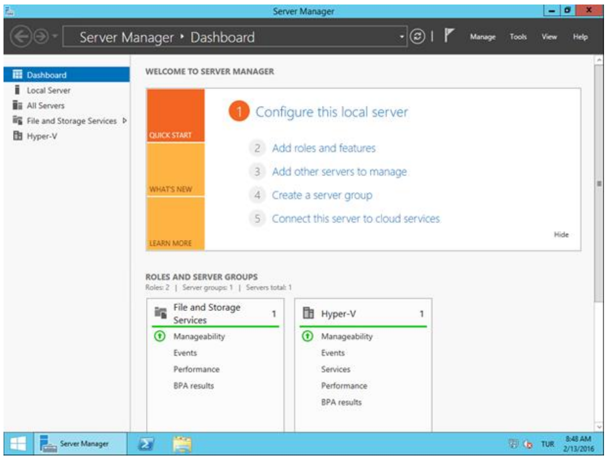

Figure D.24 Roles and features

2. In Figure D.25, click Next;

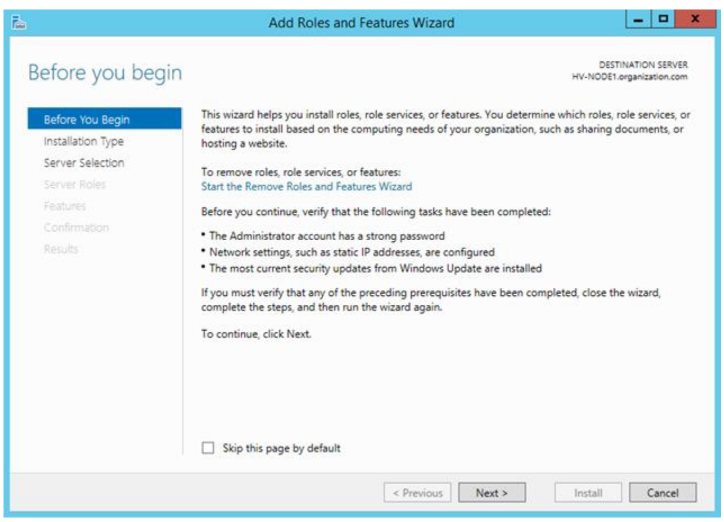

Figure D.25 Add roles and features wizard 1

3. In Figure D.26, click Next;

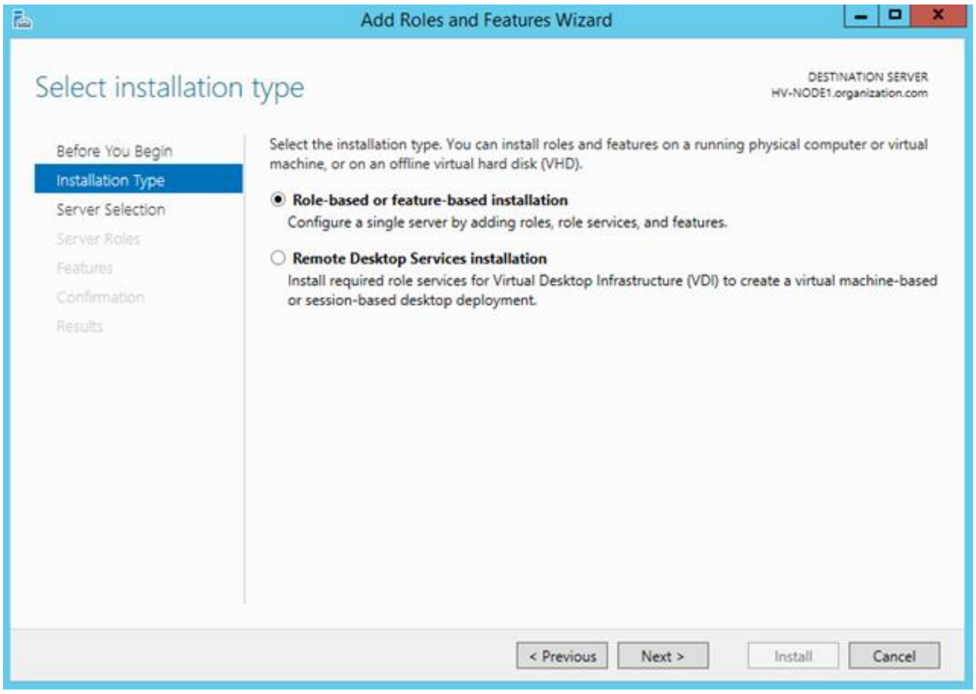

Figure D.26 Add roles and features wizard 2

4. In Figure D.27, click Next;

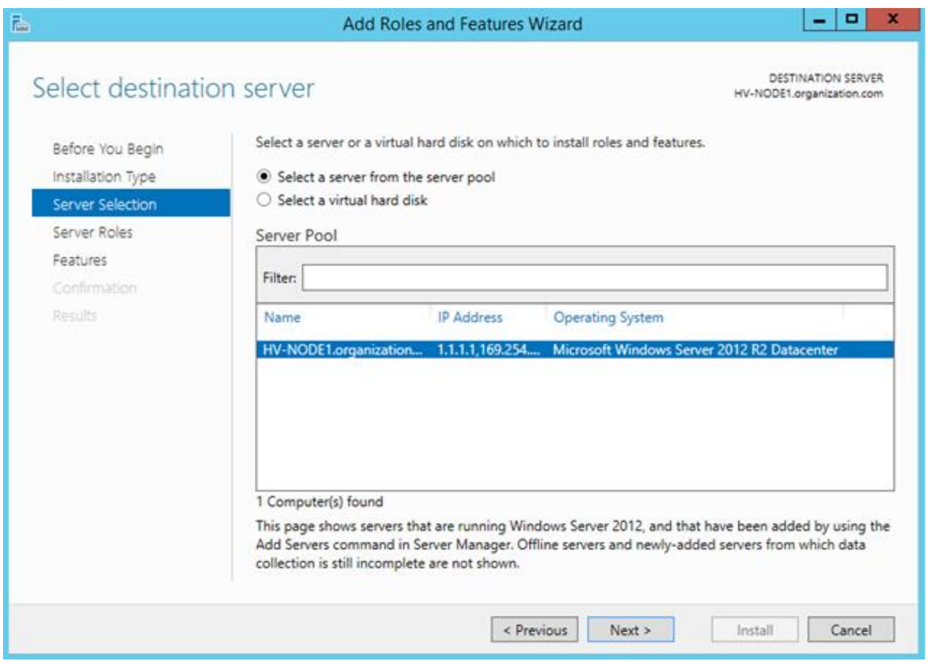

Figure D.27 Add roles and features wizard 3

5. In Figure D.28, click Next;

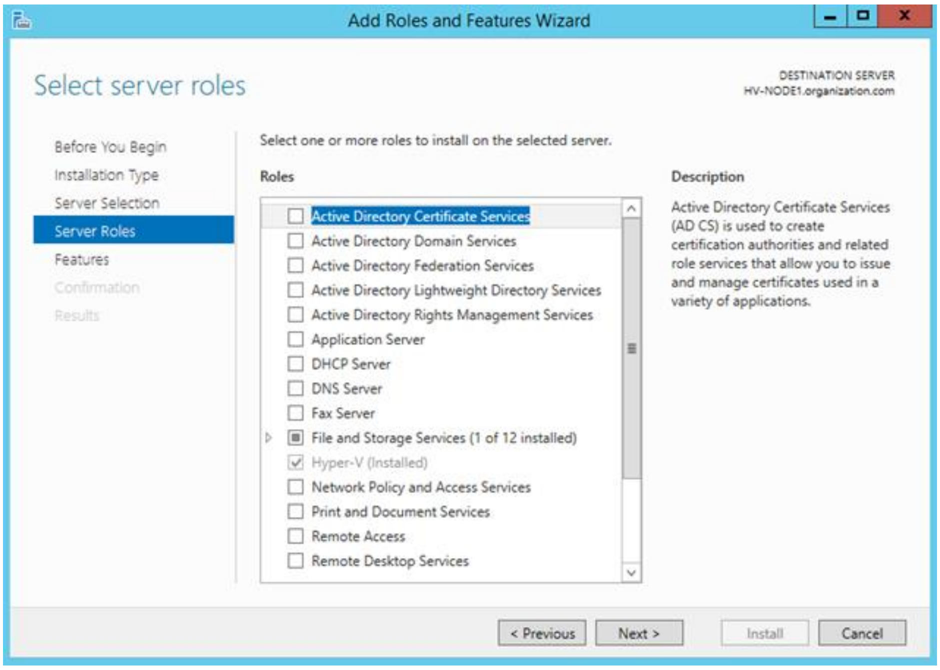

Figure D.28 Add roles and features wizard 4

6. In Figure, D.29, select Failover Clustering;

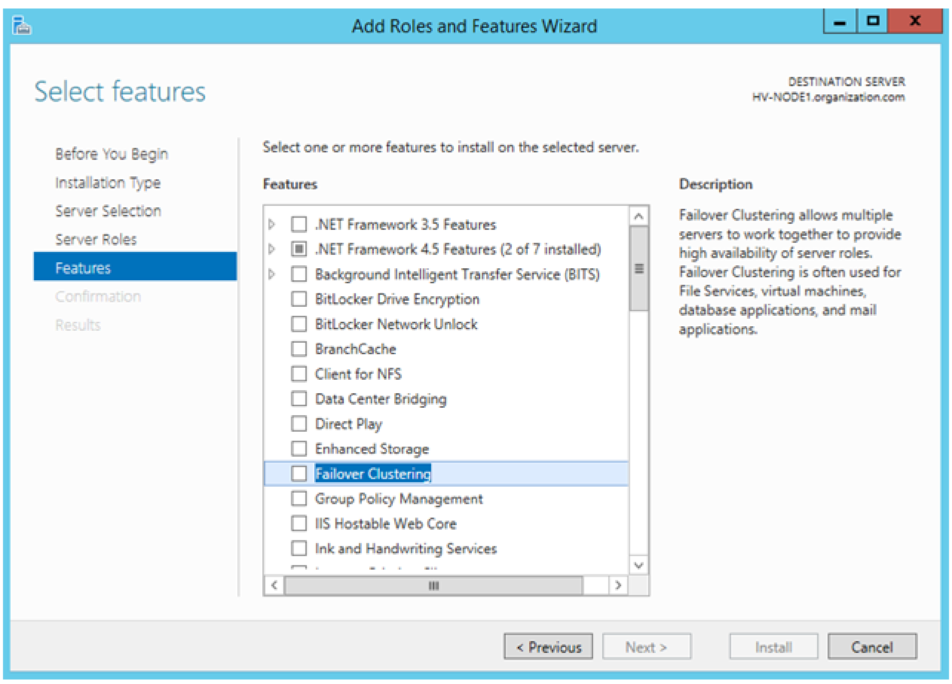

Figure D.29 Add roles and features wizard 5

7. In Figure D.30, click Add features;

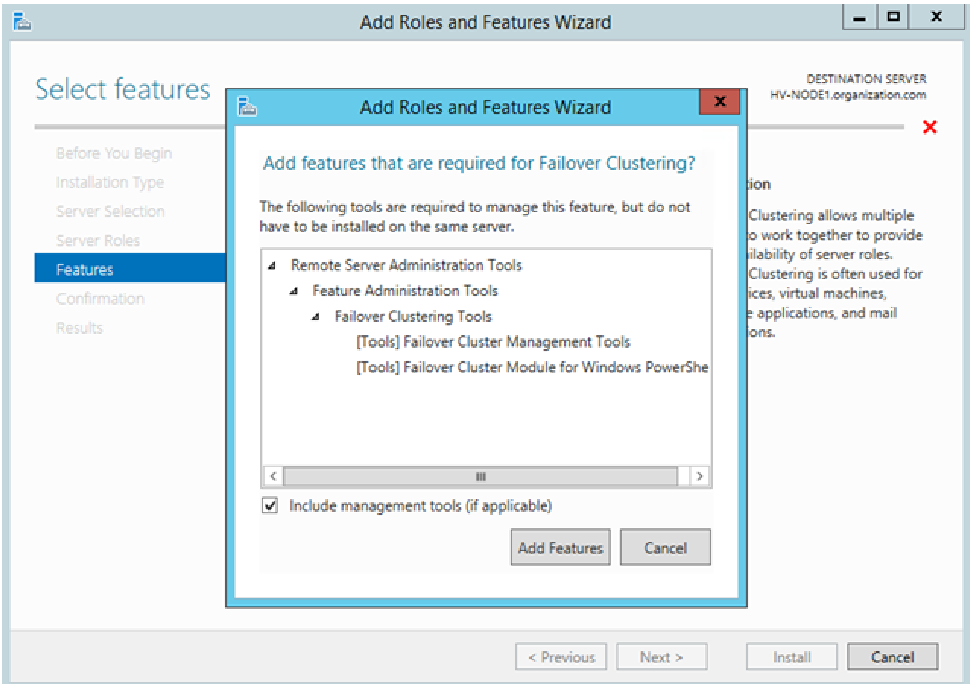

Figure D.30 Add roles and features wizard 6

8. In Figure D.31, click Install;

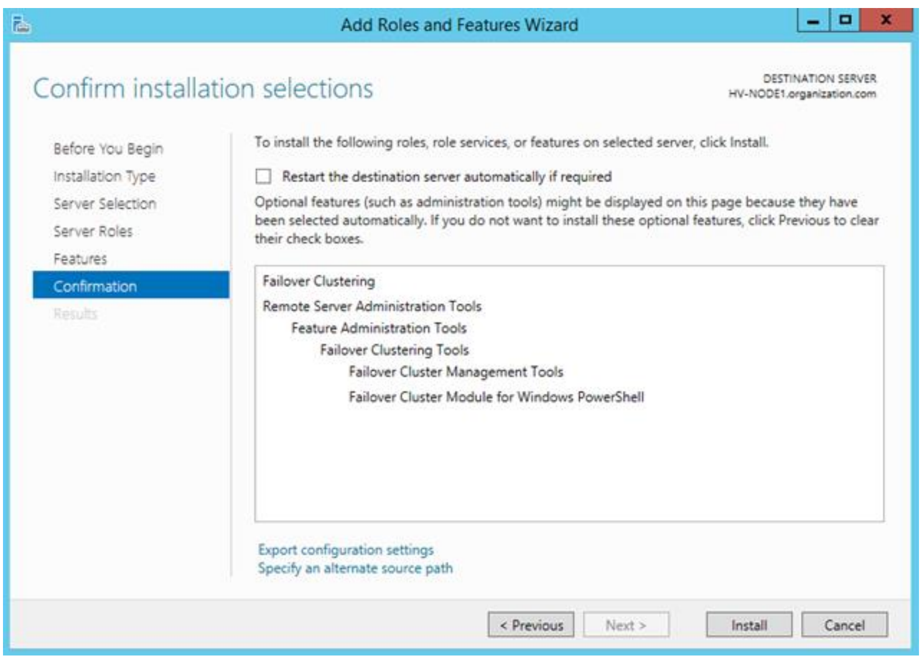

Figure D.31 Add roles and features wizard 7

9. In Figure D.32, Failover Clustering role installation starts;

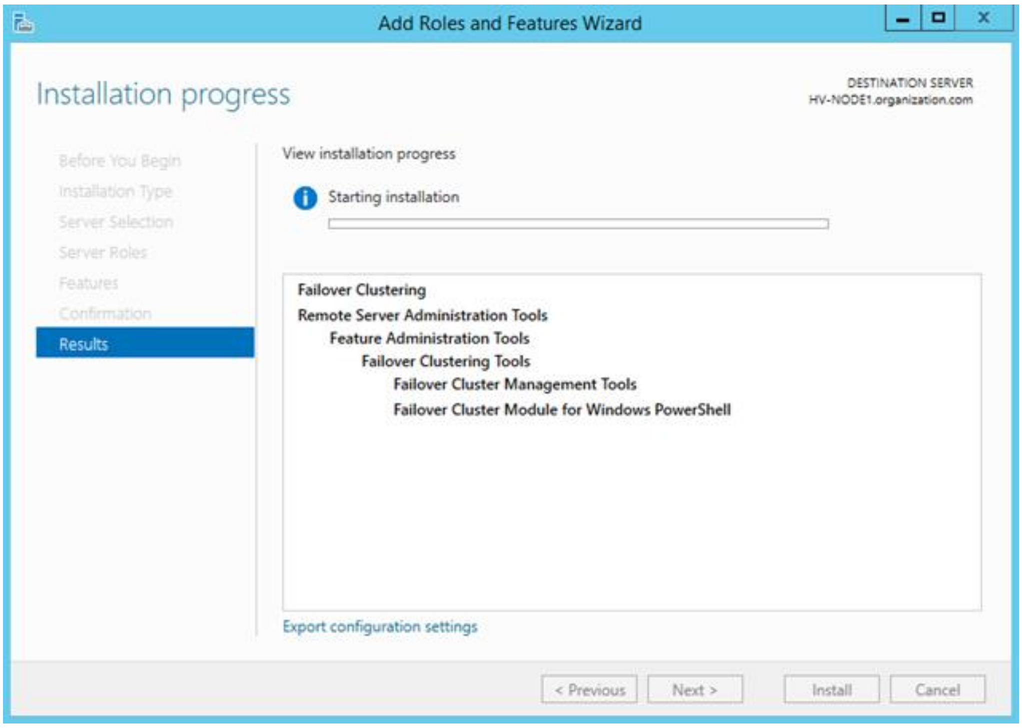

Figure D.32 Add roles and features wizard 8

10. In Figure D.33, after installation completed, verify that failover clustering feature installed by clicking the flag icon;

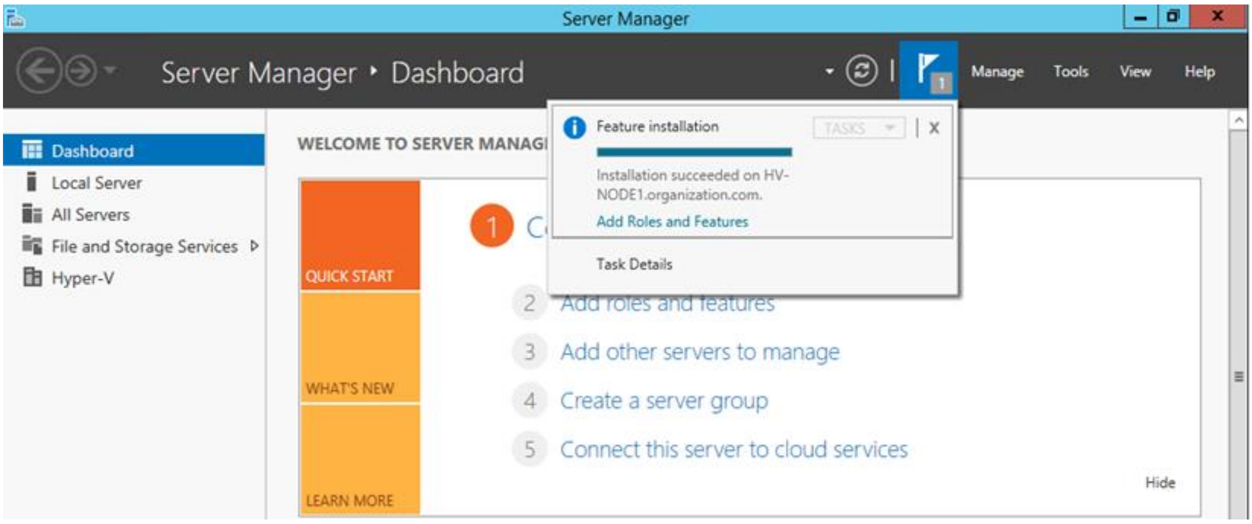

Figure D.33 Failover clustering feature installation control 1

11. In Figure D.34, verify that Failover Cluster Manager is installed by clicking Tools.

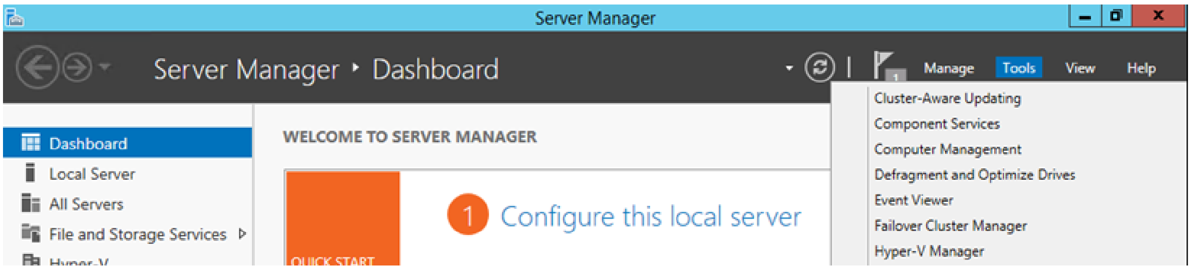

Figure D.34 Failover clustering feature installation control 2

D.4 - Using Hyper-V Virtual Switch and Creating New External Network

"The Hyper-V Virtual Switch is a software-based layer-2 Ethernet network switch that is available in Hyper-V Manager when you install the Hyper-V server role" [Hyper-V Virtual Switch Overview, 2016].

There are three types of virtual switch in Hyper-V as seen in Figure D.35. External, Internal and Private. External creates a connection from the physical adapter and the virtual machine. It allows a virtual machine to access the network through the network adapter. Internal allows communications between the virtualization servers and the virtual machines. Private provides communications only among the virtual machines and allows the virtual machines to talk to each other only [Panek, 2009].

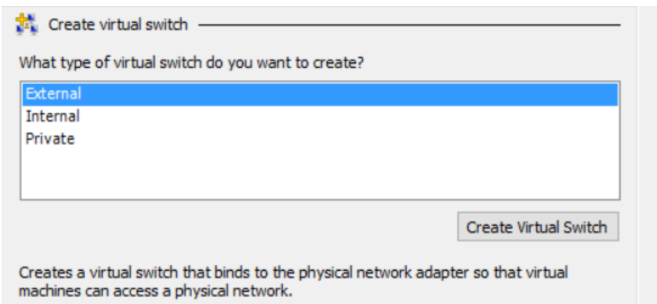

Figure D.35 Create virtual switch

Follow the steps below to create external network for virtual machines to access internet or external any other network in the system depicted as Figures D.36-D.40.

1. Click Virtual Switch Manager from Hyper-V Manager Console;

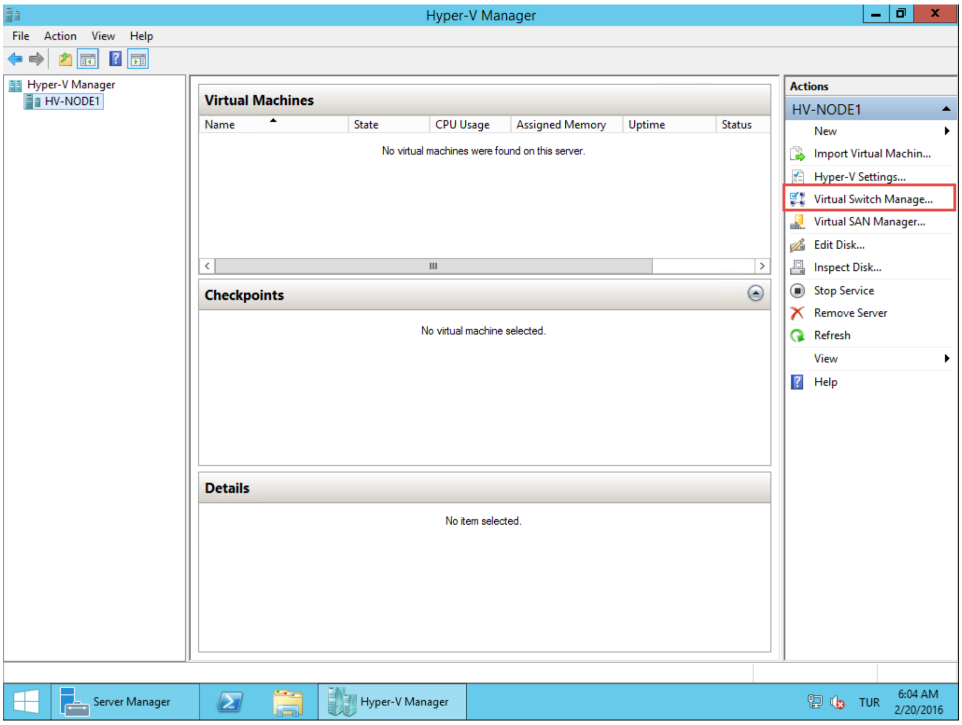

Figure D.36 Virtual switch manager 1

2. Click New Virtual network switch, select External and then click Create Virtual Switch and click Ok;

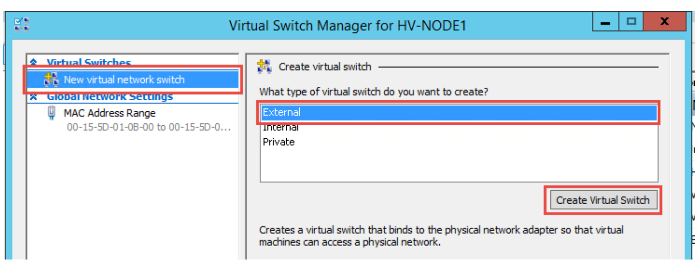

Figure D.37 Virtual switch manager 2

3. Write MNetwork as name, select physical network card that is used as External Switch, click Allow management operating system to share this

network adapter and click Apply. When Allow management … is selected, physical network card is bind to new external switch and all protocols including TCP/IP are removed. New virtual network card is added to hv-node1, if there is old TCP/IP information, it is added to new virtual network card and added as a member of external virtual network;

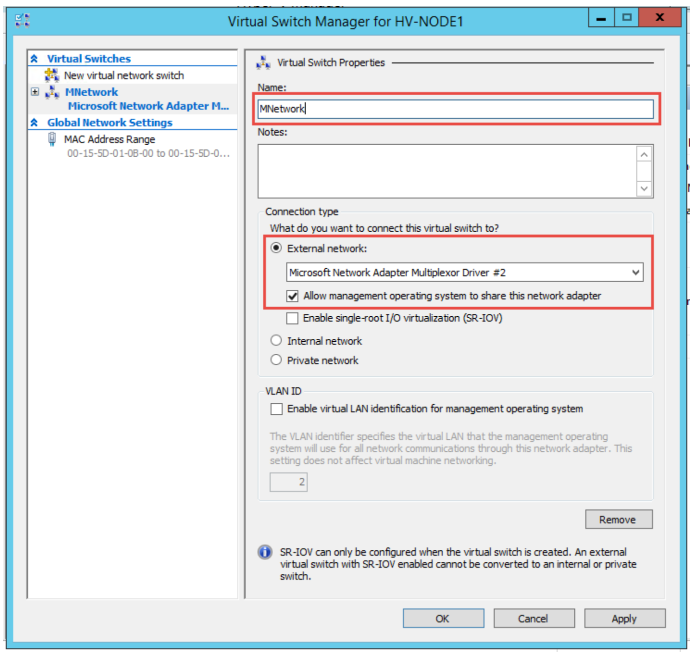

Figure D.38 Virtual switch manager 3

4. It prompts a message regarding network connectivity as seen in Figure D.39. Click Yes;

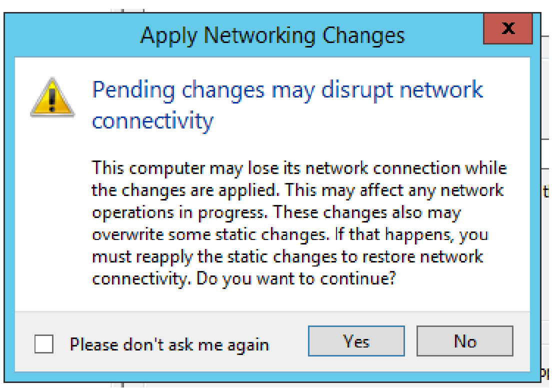

Figure D.39 Apply networking changes

5. After VM1 external virtual network is created, there is not seen TCP/IP information and all the protocols are disabled except Hyper-V Extensible Virtual Switch in Figure D.40.

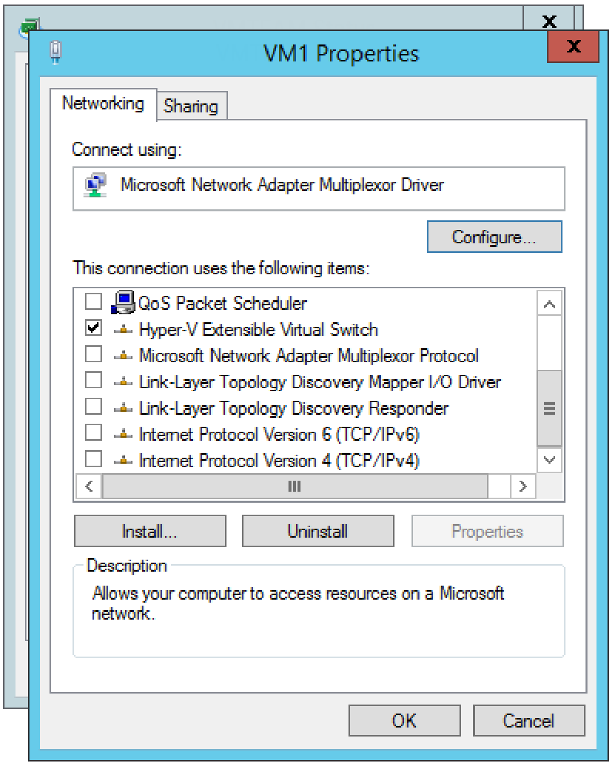

Figure D.40 VM1 properties

D.5 - iSCSI Initiator Settings

Microsoft iSCSI initiator enables to connect a host computer that is running hv-node1 and hv-node2 to an external iSCSI-based storage array which is HPE MSA 1040 2 Port 1G ISCSI DC SFF through an ethernet network adapter. Microsoft iSCSI Initiator in existing network infrastructure is used to enable block-based SAN. SANs provide iSCSI target functionality without investigating in additional hardware, and they enable the use of iSCSI storage devices in enterprise.

Follow the following steps below to configure iSCSI Initiator on both hv-node1 and hv-node2:

1. Right Click Windows icon, and select Control Panel;

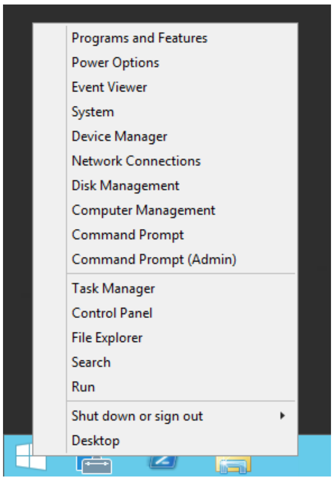

Figure D.41 Control panel

2. Control Panel window opens. Click iSCSI Initiator;

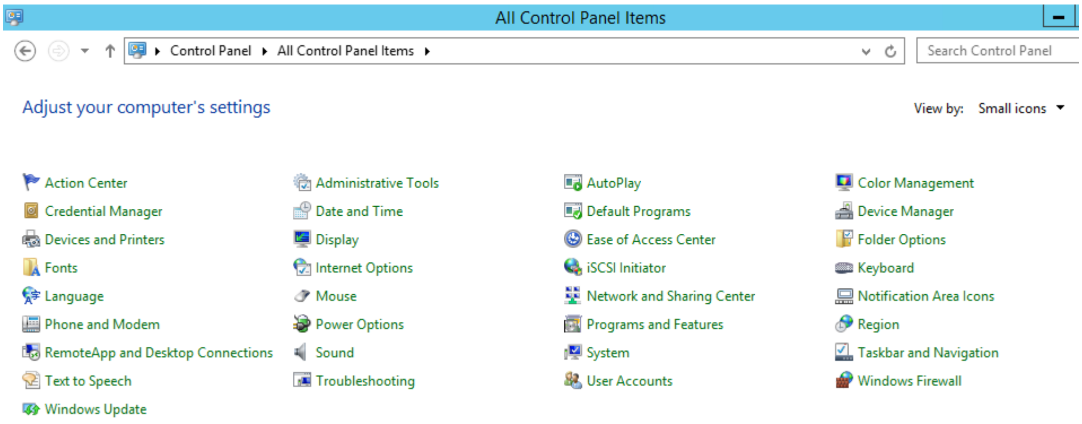

Figure D.42 iSCSI initiator

3. Click Yes only once to start The Microsoft iSCSI service;

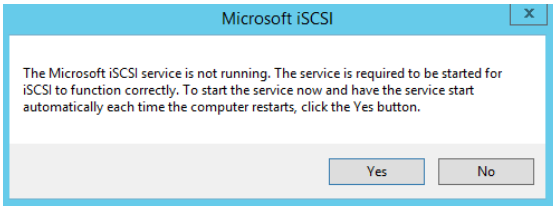

Figure D.43 Microsoft iSCSI service

4. iSCSI Initiator window opens. Click Discovery and click Discover Portal;

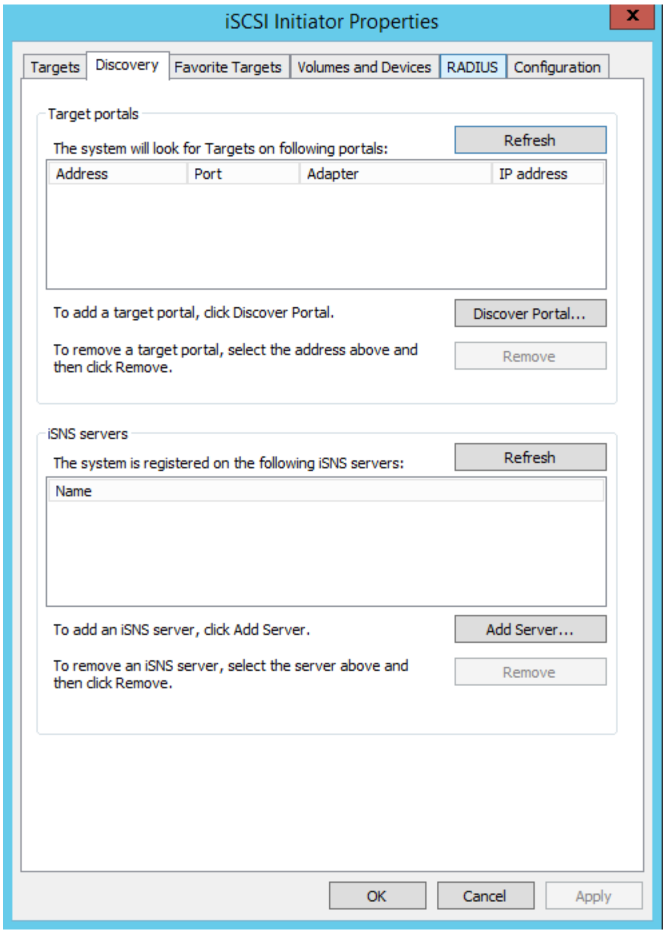

Figure D.44 iSCSI initiator properties

5. Discover Target Portal window opens. Write iSCSI network IP of storage and Port in Figure D.45;

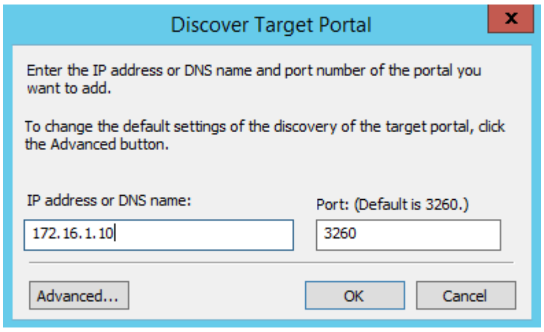

Figure D.45 Discover target portal

6. If iSCSI Initiator connect successfully to target, it shown as Figure D.46 without prompting any warning.

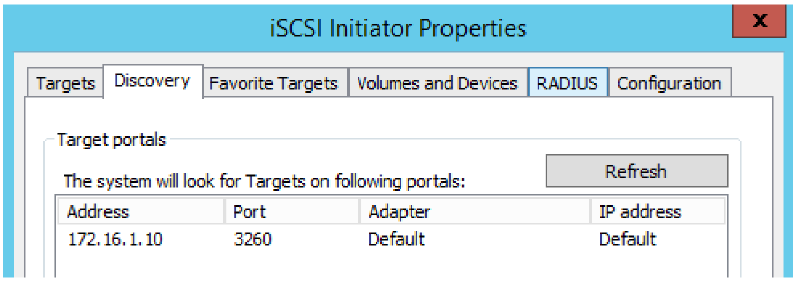

Figure D.46 Discovery

APPENDIX E - Storage Configuration

HPE MSA 1040 has two controllers cards with two 1 Gigabit Ethernet ports on each controller. MSA 1040 front and rear features shown in Figure E.1 and E.2 below.

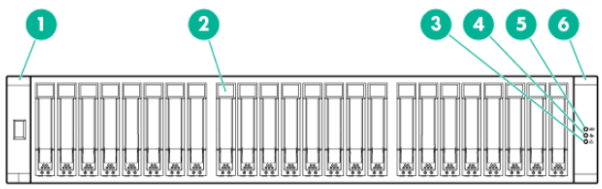

1. Left bezel ear
2. Drive in slot 9
3. Fault LED
4. Heartbeat LED
5. Unit Identification LED
6. Right bezel ear

Figure E.1 Storage front view

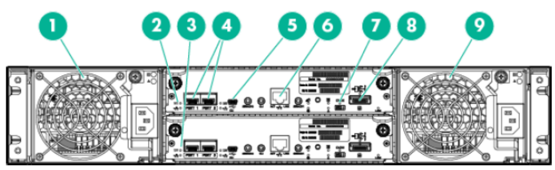

1. Power supply 1
2. Controller module A
3. Controller module B
4. Controller A Host ports
5. CLI port
6. Network management port
7. Cache Status LED
8. SAS expansion port
9. Power supply 2

Figure E.2 Storage rear view

Two controllers are used for redundancy. If one controller fails, the SAN will remain operational on the redundant controller.

For detailed reference information, instructions, and cabling examples, see the following documents:

- HPE MSA 1040 Quick Specs [HPE MSA 1040 SAN Storage, 2016];
- HPE MSA 1040 User Guide [HPE MSA 1040 SAN Storage, 2016];

• HPE MSA 1040 Cable Configuration Guide [HPE MSA 1040 SAN Storage, 2016];

• HPE MSA 1040 SMU Reference Guide [HPE MSA 1040 SAN Storage, 2016].

Direct-connect and switch-connect environment are supported. In this system, switch-connect will be used.

The following topology in Figure E.3 illustrates connecting with redundancy an MSA 1040 iSCSI to two switches and the two switches connecting to two servers. Eight Ethernet CAT6 cables are required. Storage has four IP addresses which are 172.16.1.13, 172.16.1.14, 172.16.1.15 and 172.16.1.16 in order to communicate with hv-node1's IP address which is 172.16.1.11 and hv-node2's IP address which is 172.16.1.12 from iSCSI network. After give an IP address for both the nodes and the storage, verify connectivity by pinging.

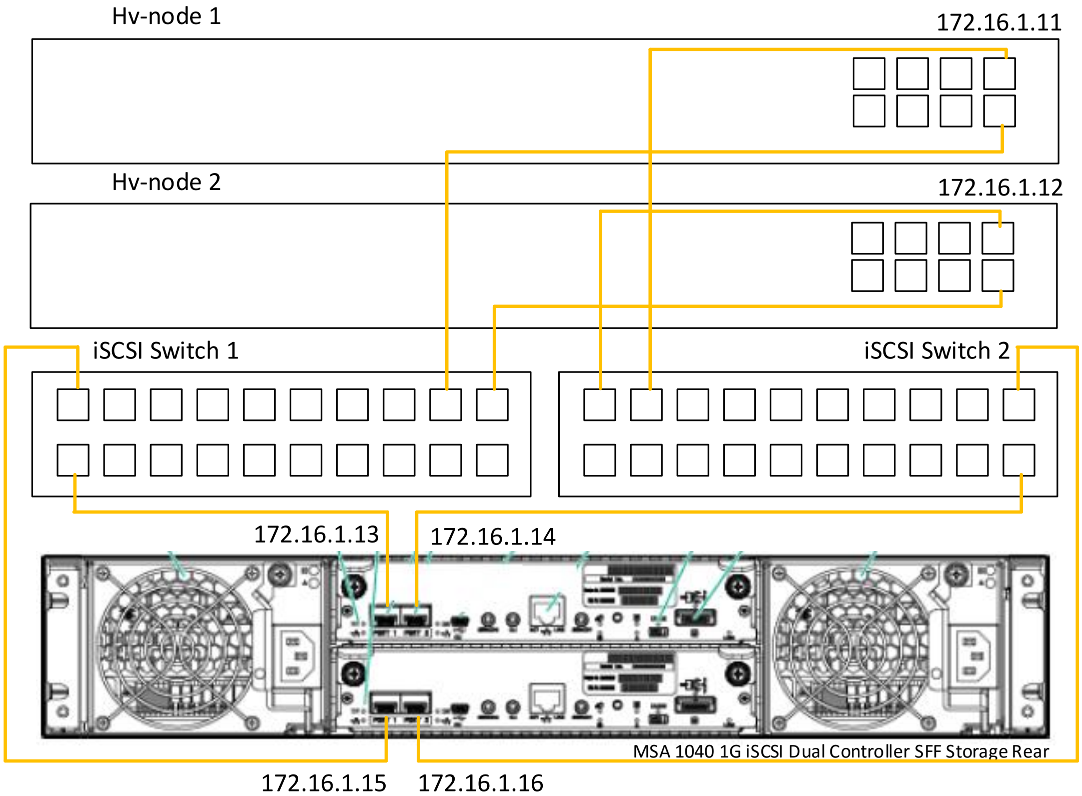

Figure E.3 Storage iSCSI network connection

There are two management ports from two controllers. In Figure E.4, there is a topology to show how to connect storage management ports to LAN switch to communicate with hv-node1 and hv-node2.

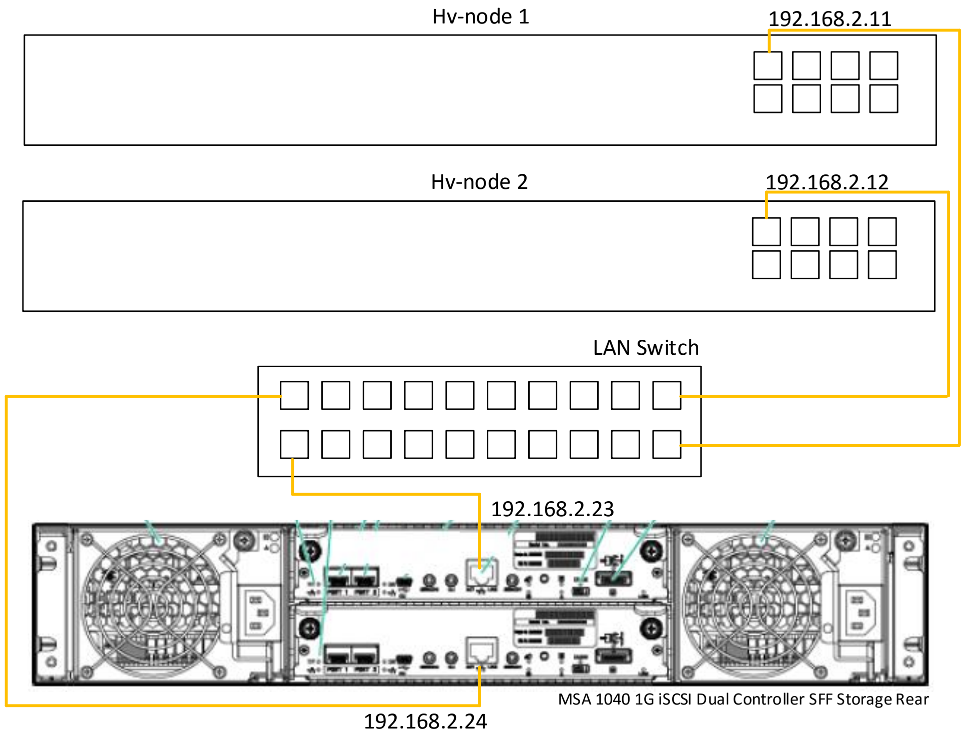

Figure E.4 Storage management network connection

After configure management connection successfully, the web based console is accessed via Internet browser. After logged in, to create a VDisk encompassing the entire drive space available, click Provisioning and then Create Vdisk in Figure E.5.

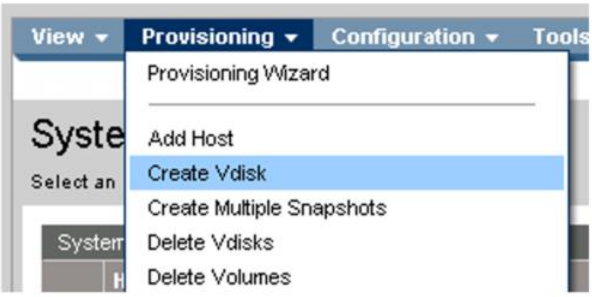

Figure E.5 Create vdisk

Raid 5 on 13x900 GB drives are used. After create Raid 5, five volumes are created. One volume which is Quorum for the configuration of failover cluster and four volumes for the storage of virtual machines. The Quorum volume will be 1GB and 1.8 TB for each four other volumes. In order to create volume for Quorum click on the Vdisk created and after that click Provisioning – Create Volume in Figure E.6.

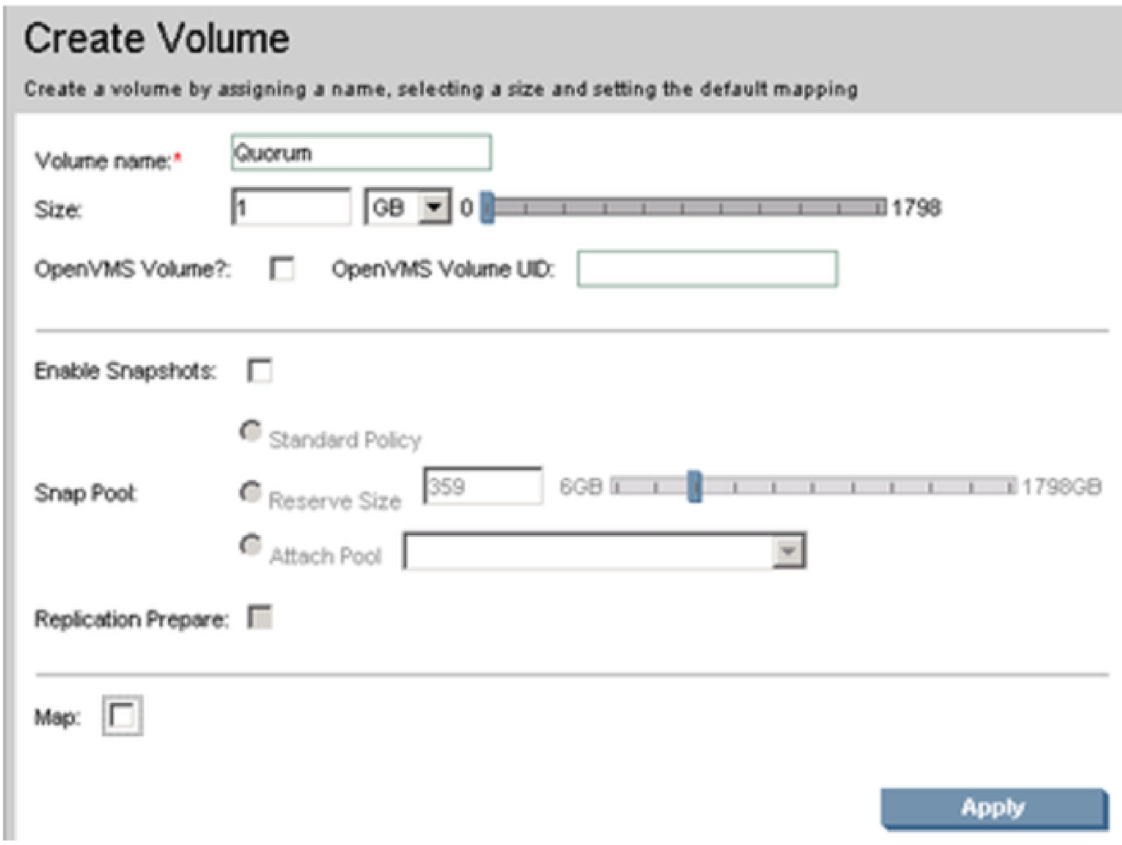

Figure E.6 Quorum volume creation

Repeat the same step in Figure E.6 for other four volumes. Four volumes' size are 1.8 TB and names are LUN1, LUN2, LUN3 and LUN4.

APPENDIX F - Hyper-V Nodes iSCSI Initiator Advanced Configuration

Follow the steps below to configure iSCSI Initiator on hv-node1.

1. Click iSCSI Initiator in Figure F.1;

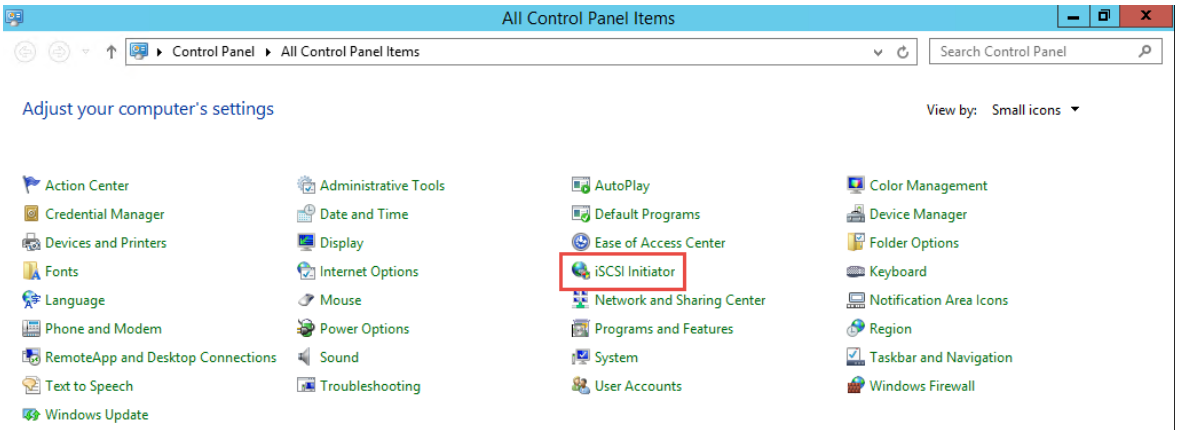

Figure F.1 iSCSI initiator 1

2. Click Targets and click Refresh in Figure F.2 There will be information at the Discovered targets about storage below and click Connect;

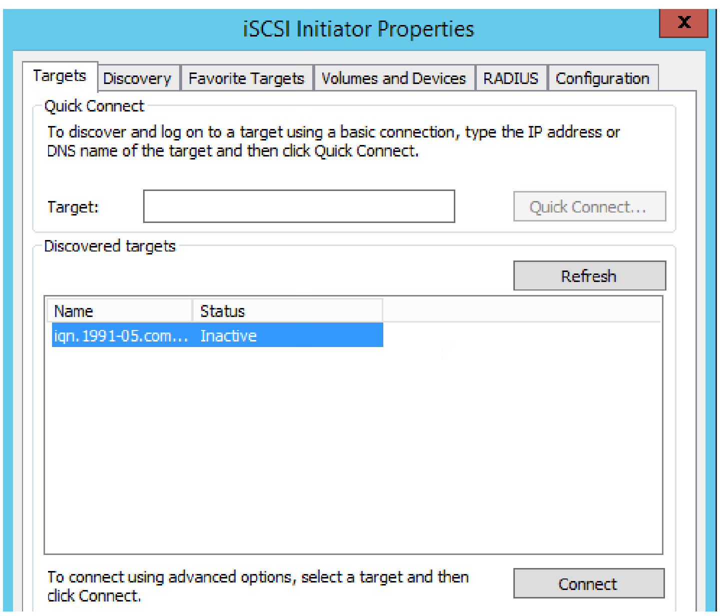

Figure F.2 iSCSI initiator 2

3. After click Connect, there will be popup message regarding connection every time this computer restarts in Figure F.3 Click Ok to accept it;

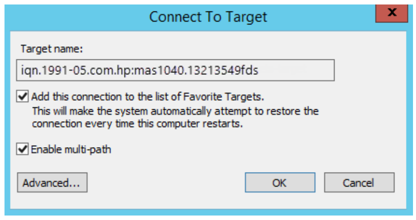

Figure F.3 Connect to target

4. Verify that it is connected in Figure F.4;

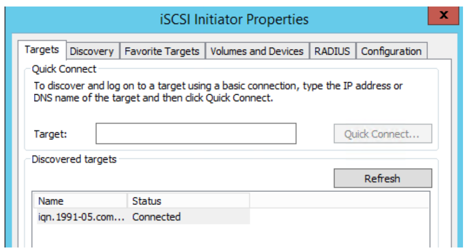

Figure F.4 Target connection

5. In order to see the disks from storage click Volumes and Devices and click Auto Configure and Click OK;

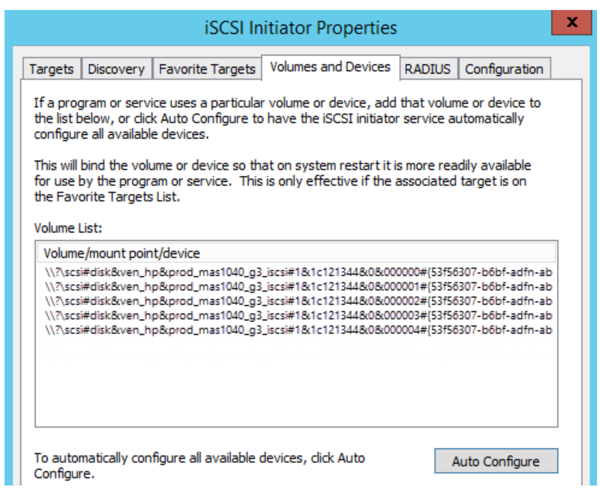

Figure F.5 Volumes and devices

6. Open Disk Management and verify that all the disks are mapped from storage. It is normal that disk will be showed as offline and unknown in Figure F.6;

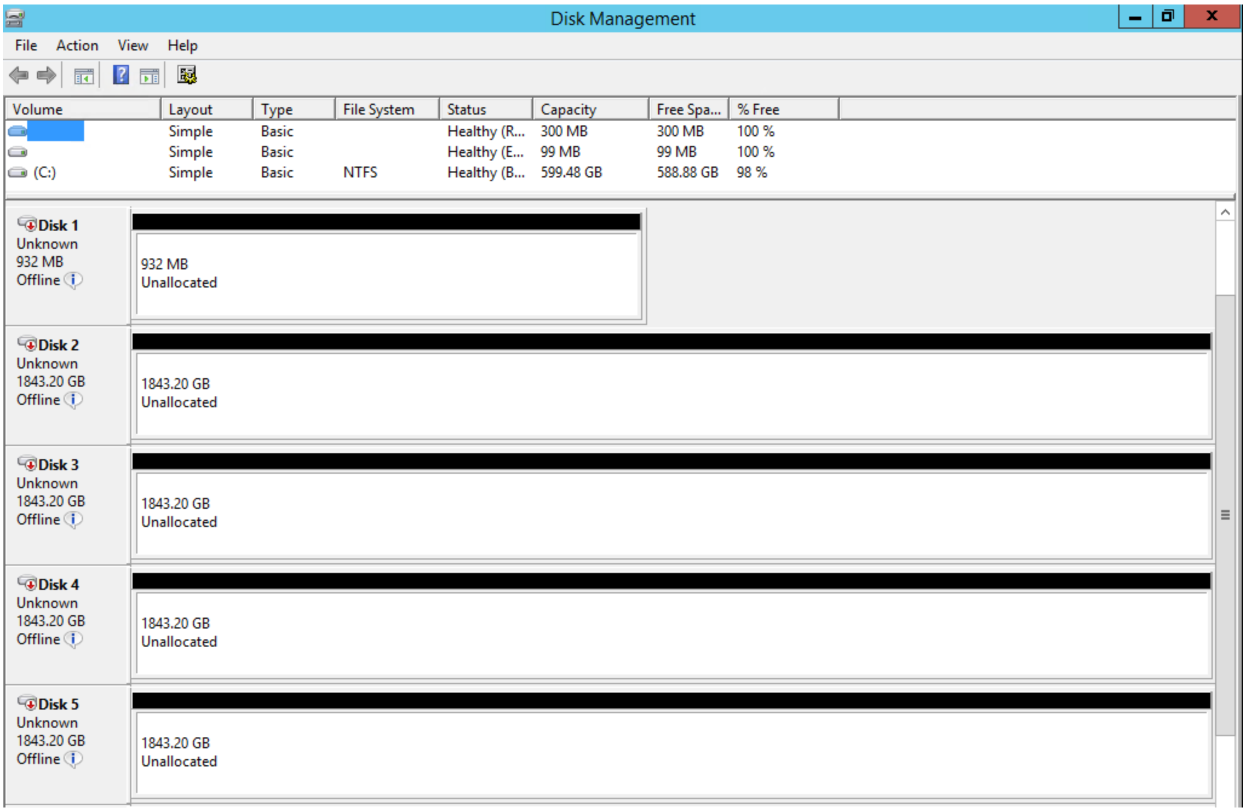

Figure F.6 Disk management

7. Repeat the steps between 1 and 6 for hv-node2 and verify that all disks are shown from storage for hv-node2.

APPENDIX G - Hyper-V Failover Cluster Installation and Configuration

Follow the steps below on hv-node1.

1. Open Disk Management. Right click Disk 1 and Click Online in Figure G.1. Repeat step for Disk2, Disk3, Disk4 and Disk5;

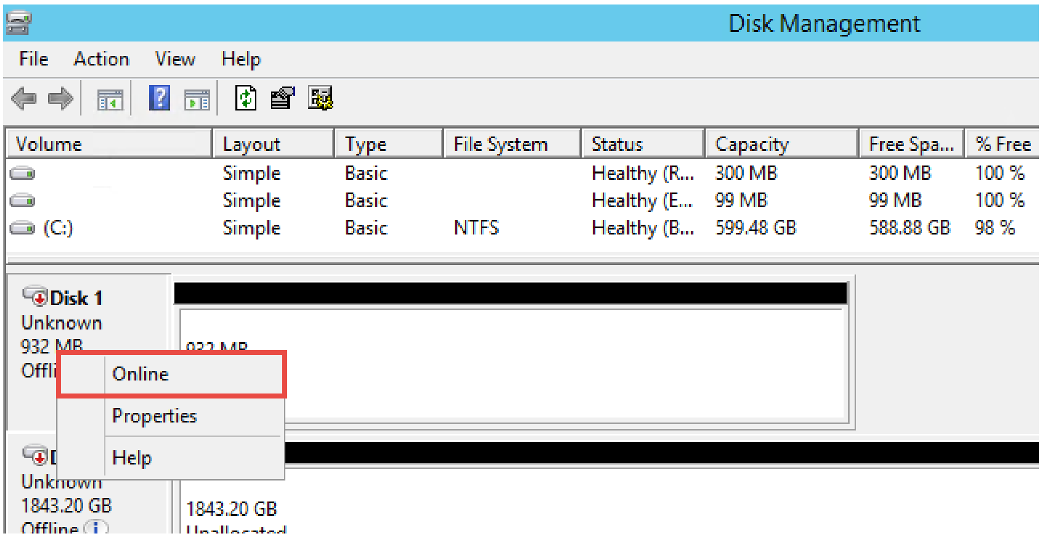

Figure G.1 Online disk

2. Right click Disk 1 and click Initialize Disk in Figure G.2;

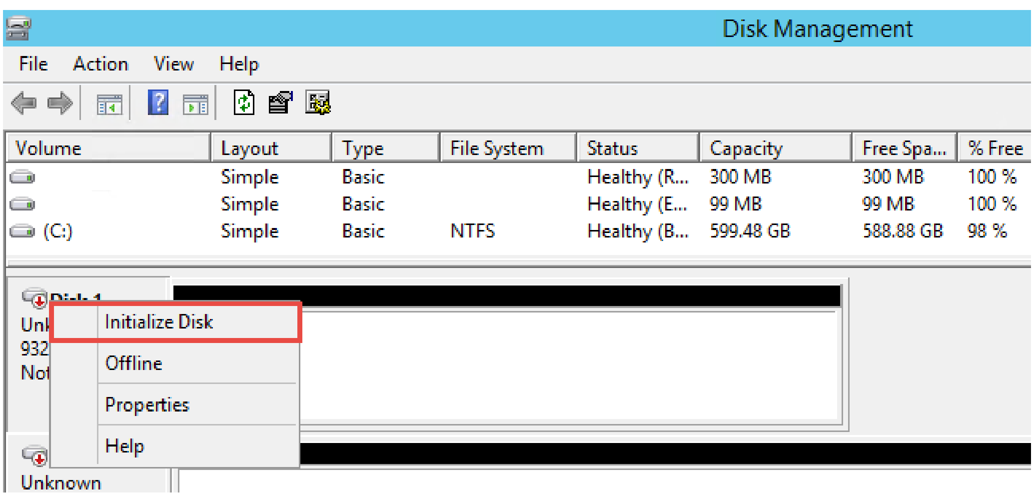

Figure G.2 Initialize disk

3. Initialize Disk window opens. Select all Disks, GPT (GUID Partition Table) and click OK in Figure G.3;

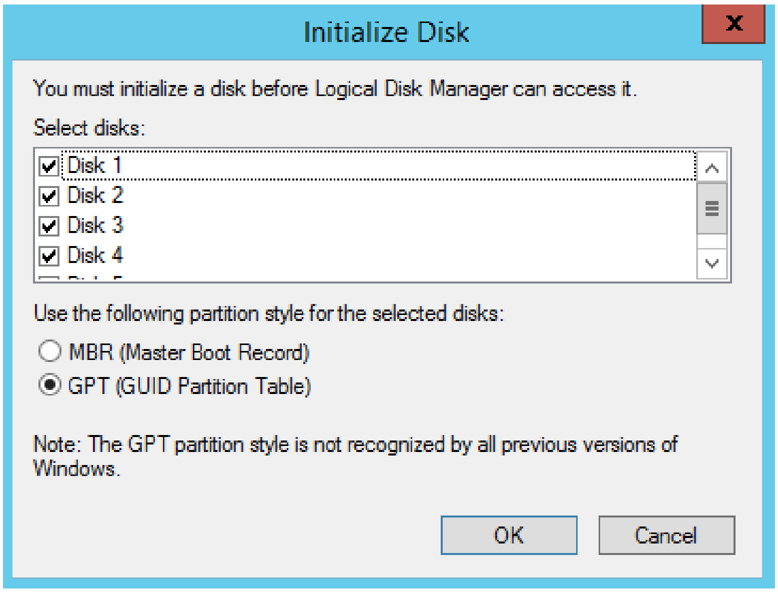

Figure G.3 Initialize all disks

4. Right click Disk 1 and click New Simple Volume in Figure G.4;

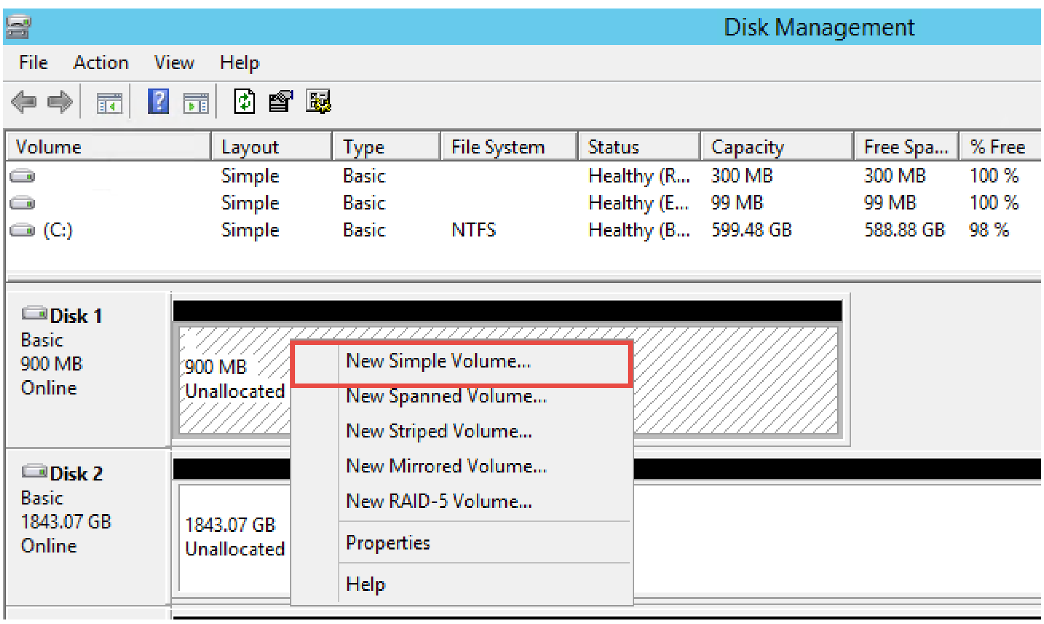

Figure G.4 New simple volume

5. Click Next two times and assign letter and write Volume label as QUORUM for Disk1 click OK in Figure G.5 Then, click Finish;

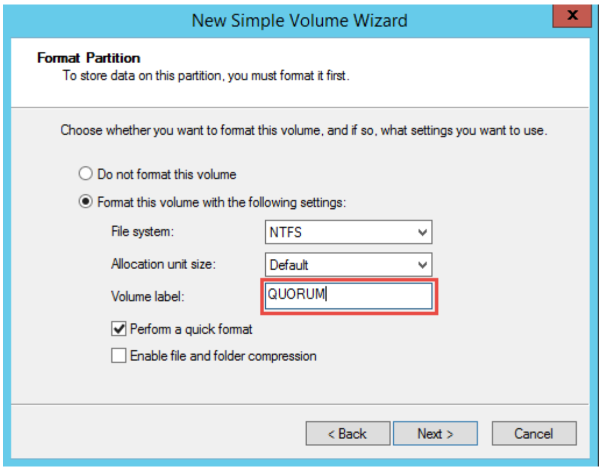

Figure G.5 Volume label

6. Repeat steps between 4 and 5 for Disk2, Disk3, Disk4, Disk5 and named as LUN1, LUN2, LUN3 and LUN4;
7. Disks are shown below on hv-node1 in Figure G.6;

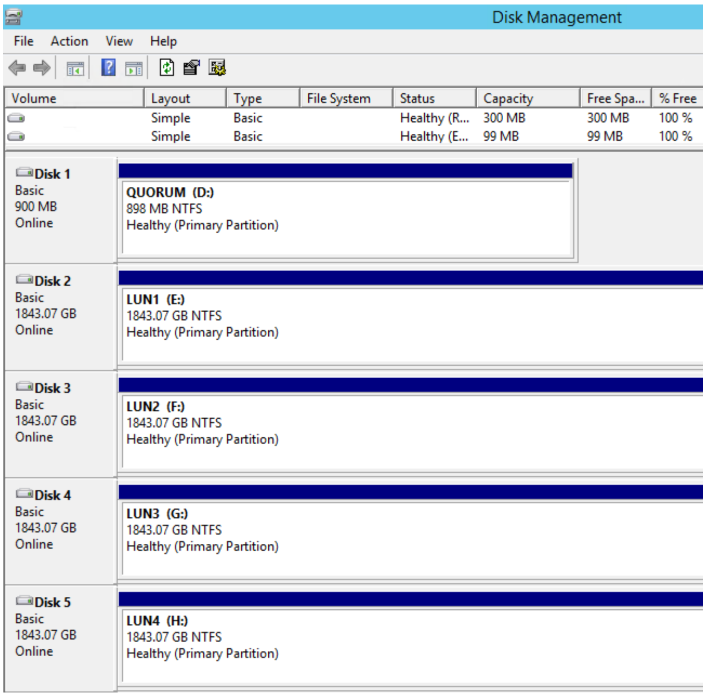

Figure G.6 All disks

8. Although disks are online in hv-node1, disks shown as offline in hv-node2 in Figure G.7.

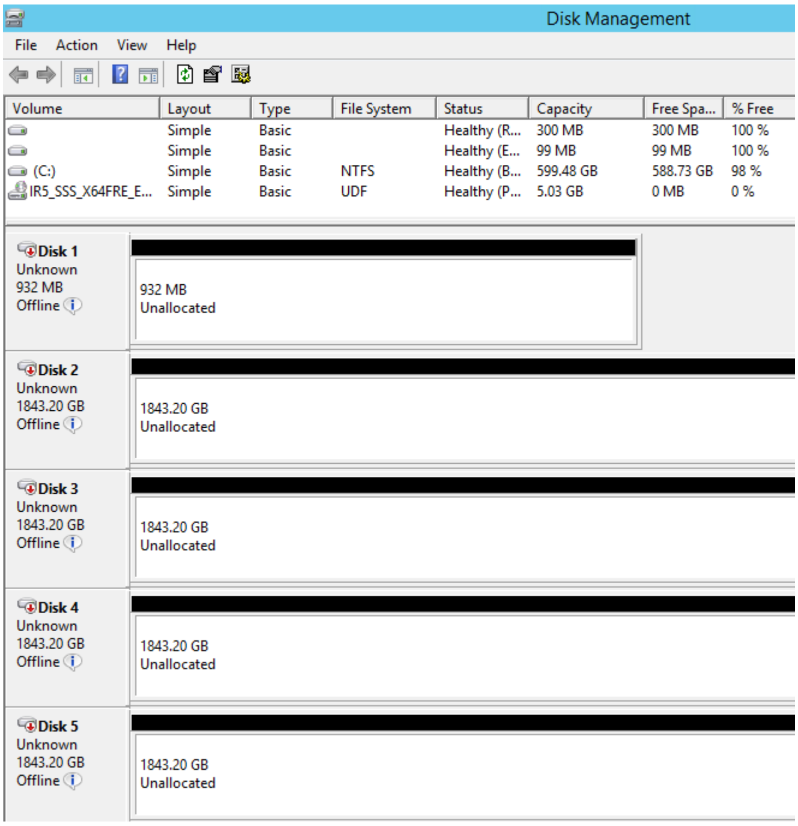

Figure G.7 Disks comparison for hv-node1 and hv-node2

G.1 - Hyper-V Failover Cluster Installation

Follow the steps below on hv-node1.

1. Open Failover Cluster Manager in Figure G.8;

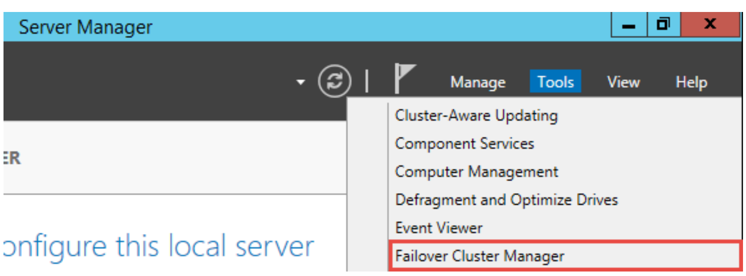

Figure G.8 Failover cluster manager

2. Right click Failover Cluster Manager and Click Create Cluster in Figure G.9;

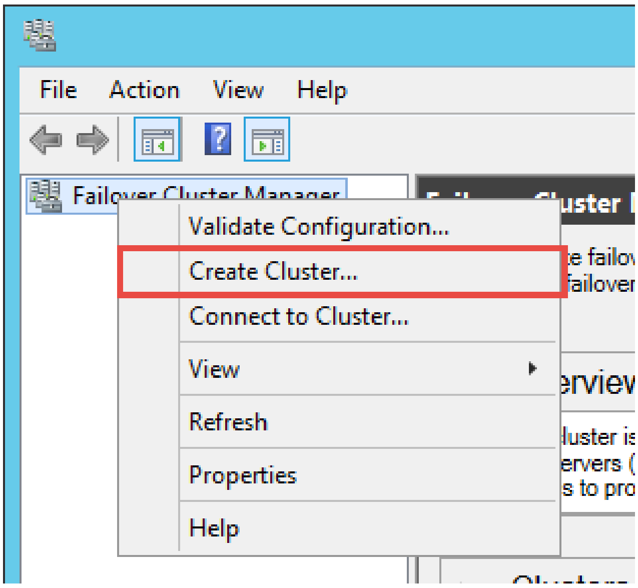

Figure G.9 Create cluster

3. Click Next in Figure G.10;

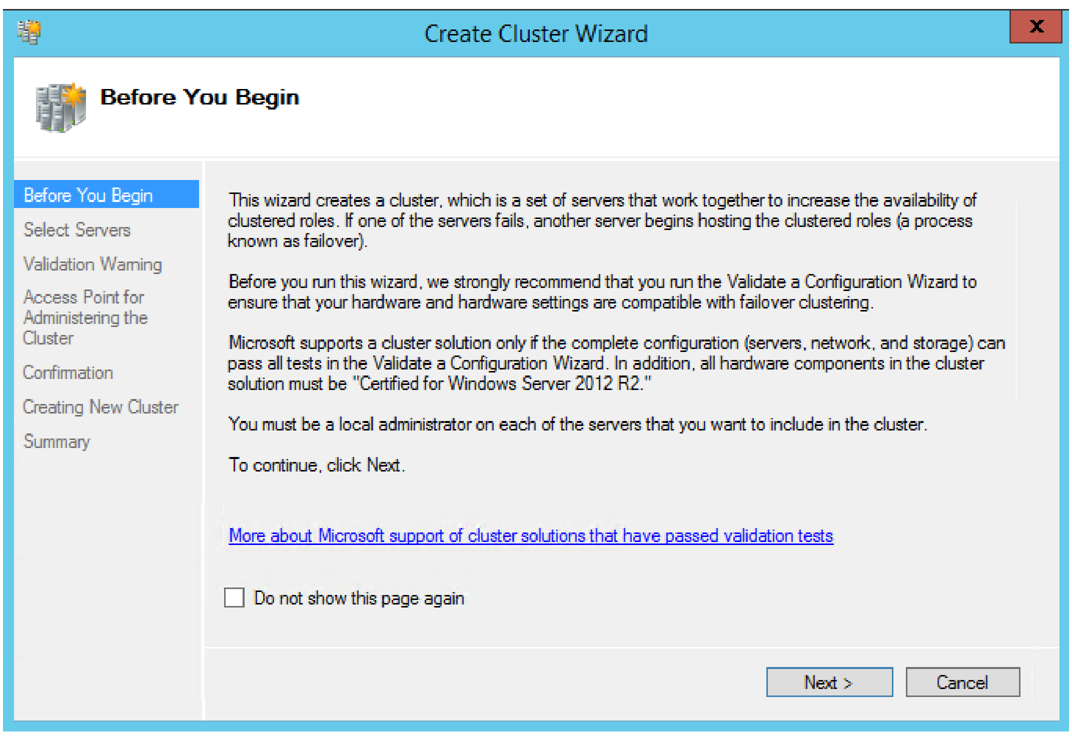

Figure G.10 Create cluster wizard

4. Click Browse, write host names and click Check Names and Click OK in Figure G.11;

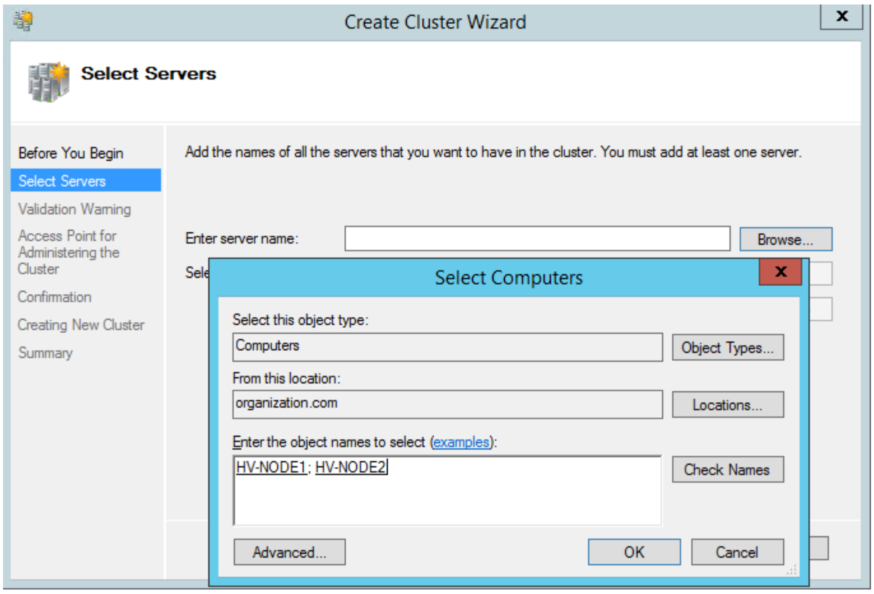

Figure G.11 Select computers

5. Verify that hosts are listed and Click Next in Figure G.12;

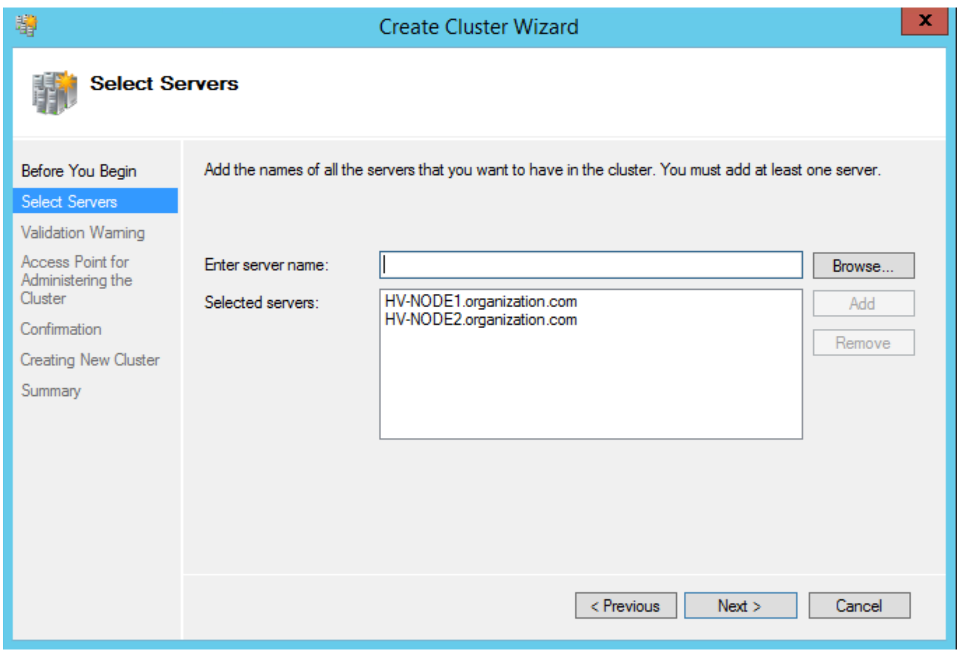

Figure G.12 Selected servers

6. Click Next to run configuration validation tests in Figure G.13;

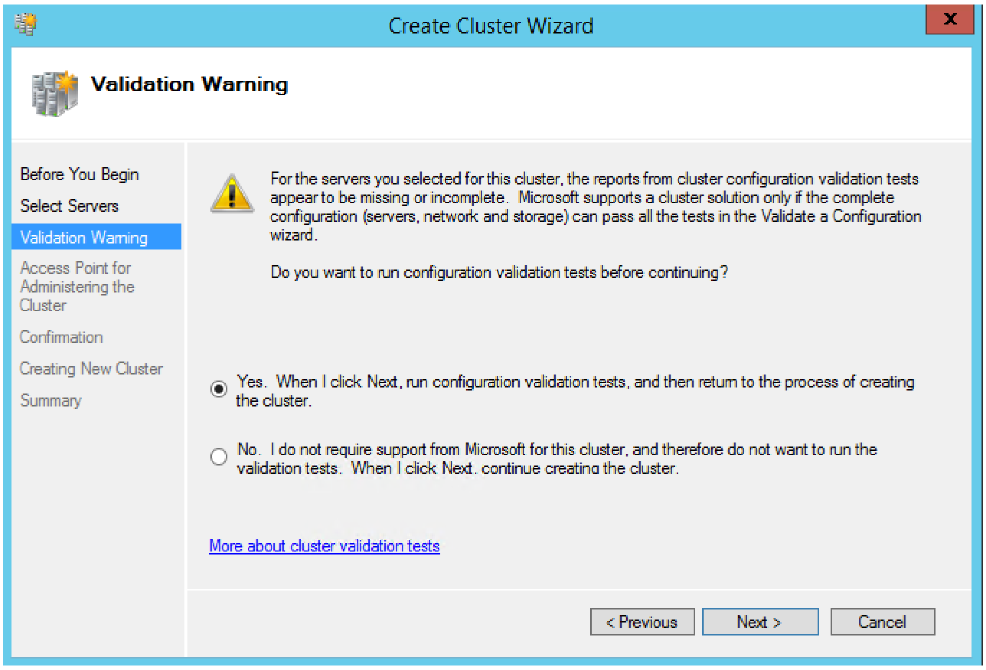

Figure G.13 Run configuration validation tests

7. Click Next in Figure G.14;

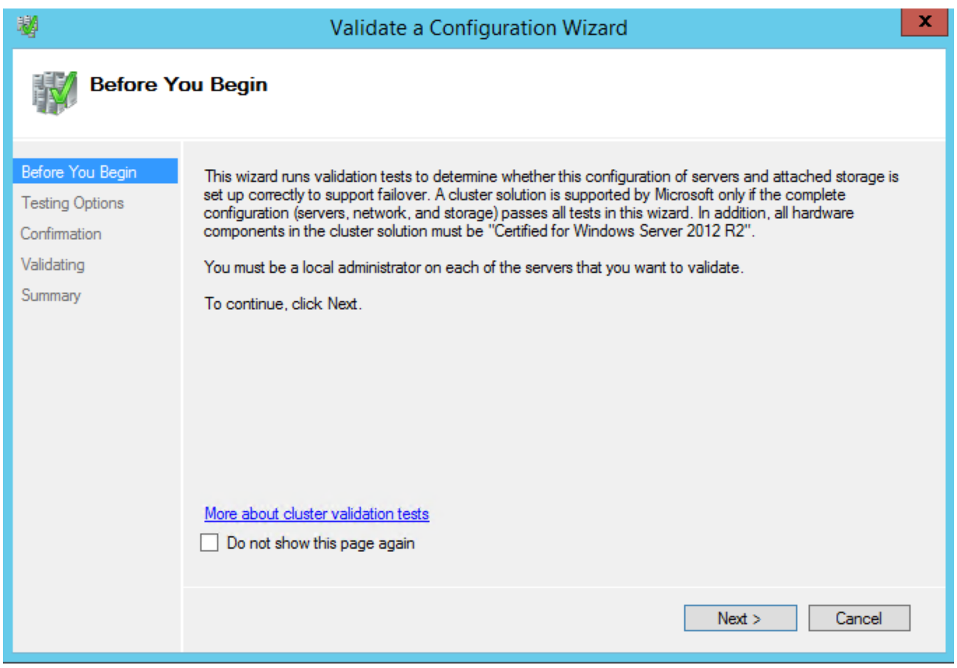

Figure G.14 Validate a configuration wizard

8. Click Run all tests in Figure G.15;

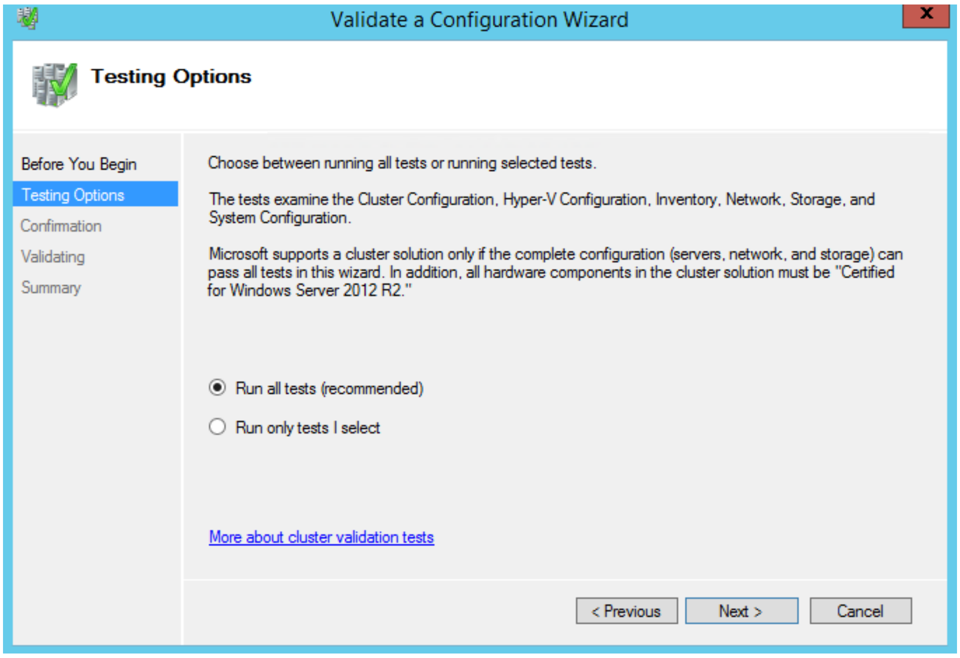

Figure G.15 Run all tests

9. Click Next in Figure G.16;

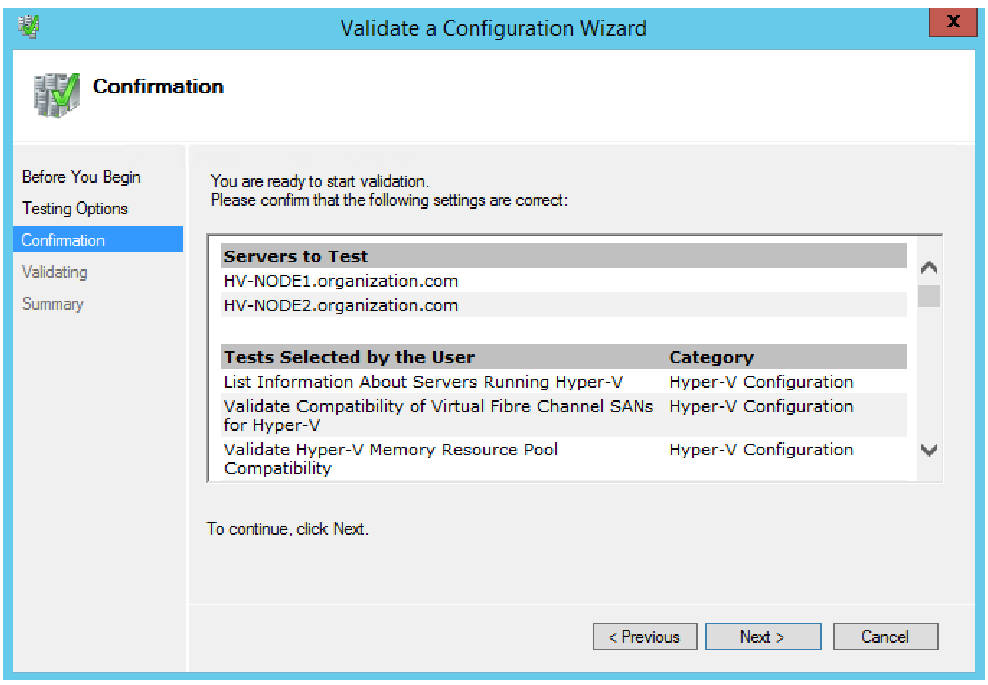

Figure G.16 Validation

10. Verify test results and if there is no problem click Finish in Figure G.17;

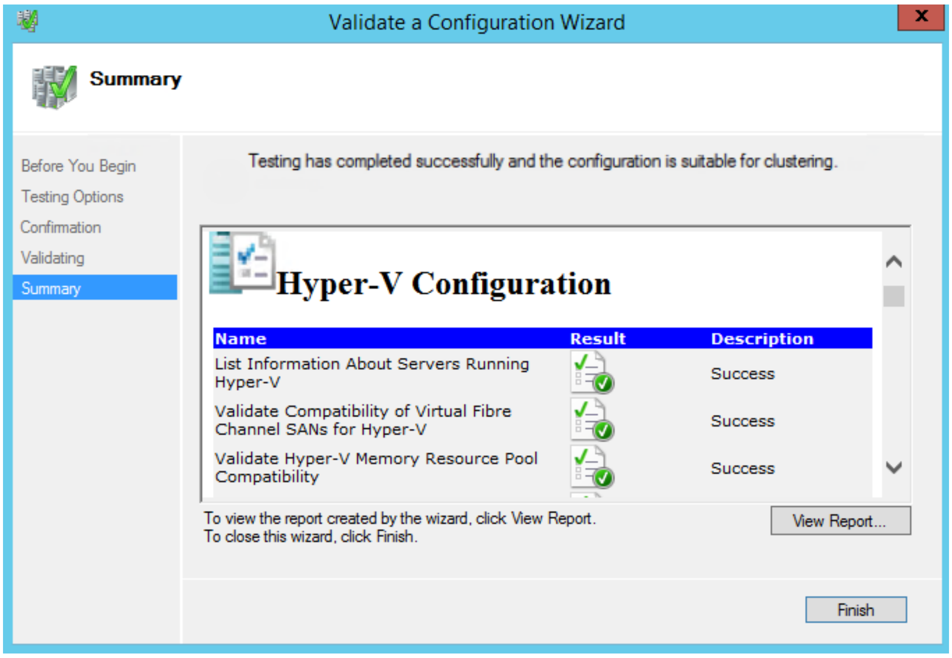

Figure G.17 Test completion

11. Write Cluster Name and give an IP address in Figure G.18;

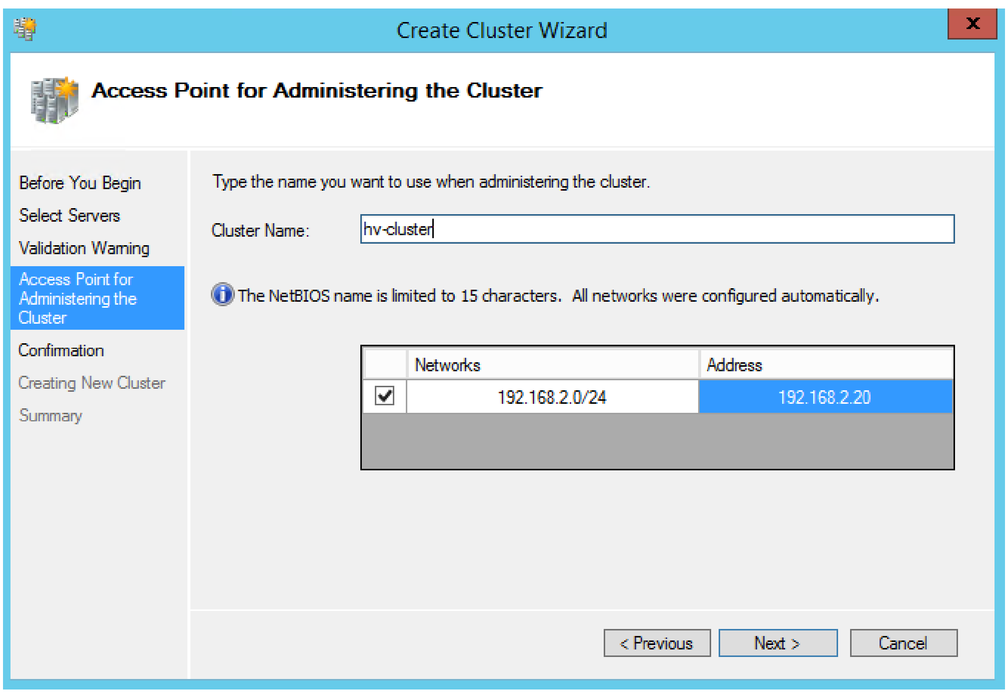

Figure G.18 Cluster name

12. Click Next in Figure G.19;

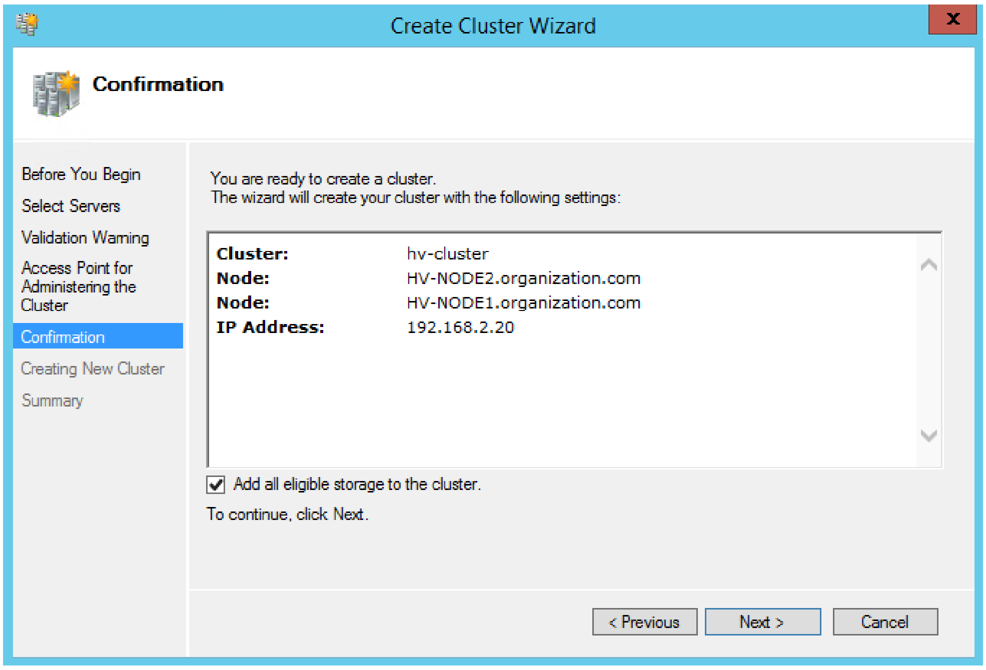

Figure G.19 Create cluster wizard

13. It is completed. Click Finish in Figure G.20;

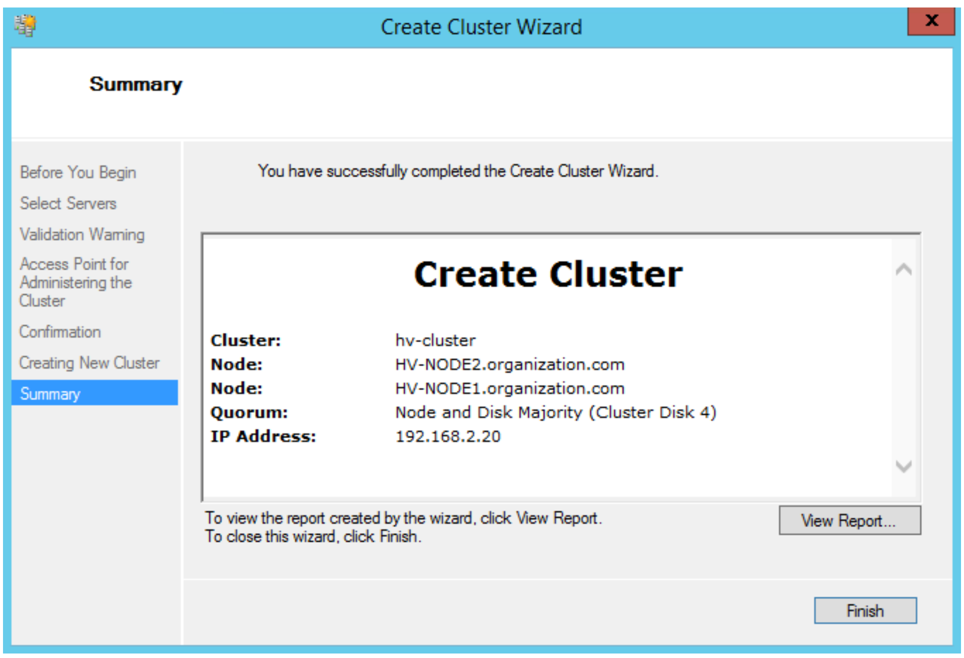

Figure G.20 Create cluster completion

14. Failover Cluster is installed in Figure G.21;

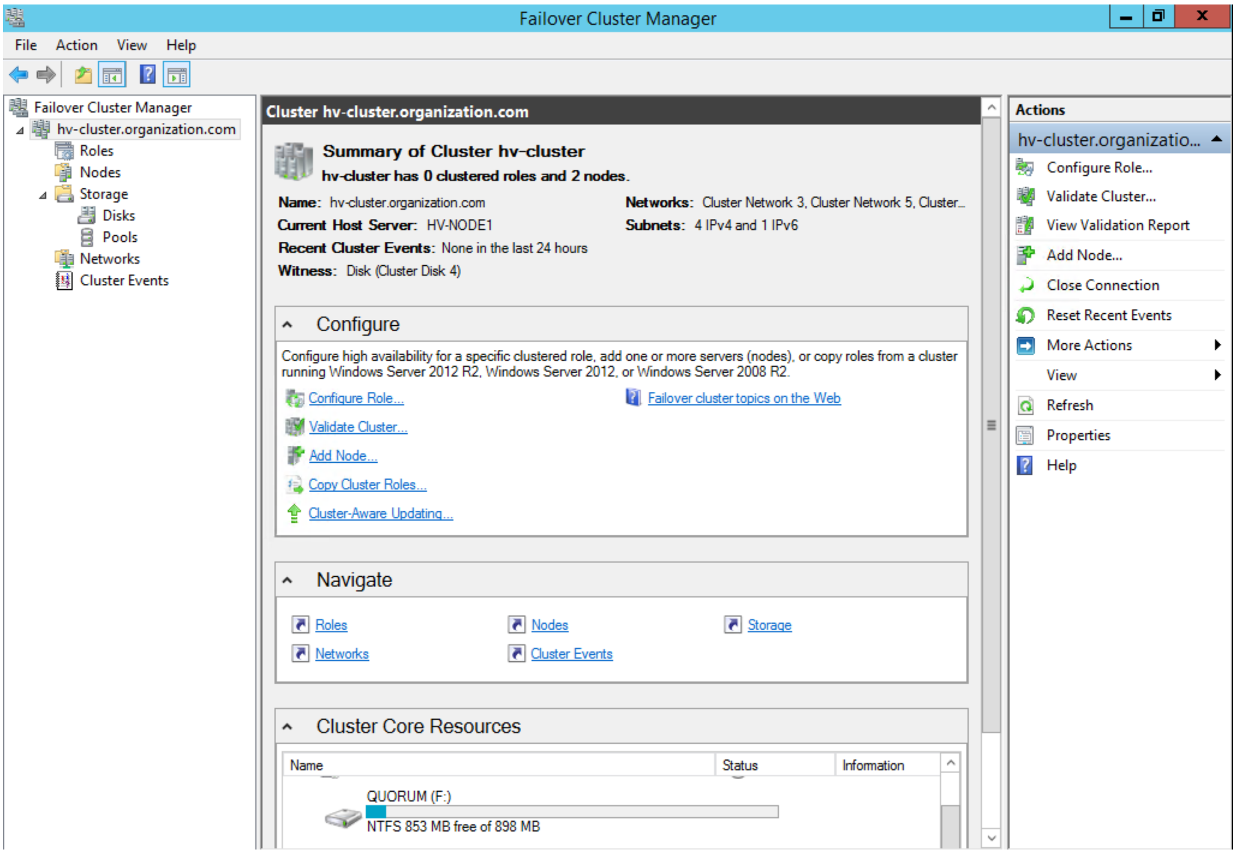

Figure G.21 Failover cluster manager

15. Nodes are shown in Figure G.22;

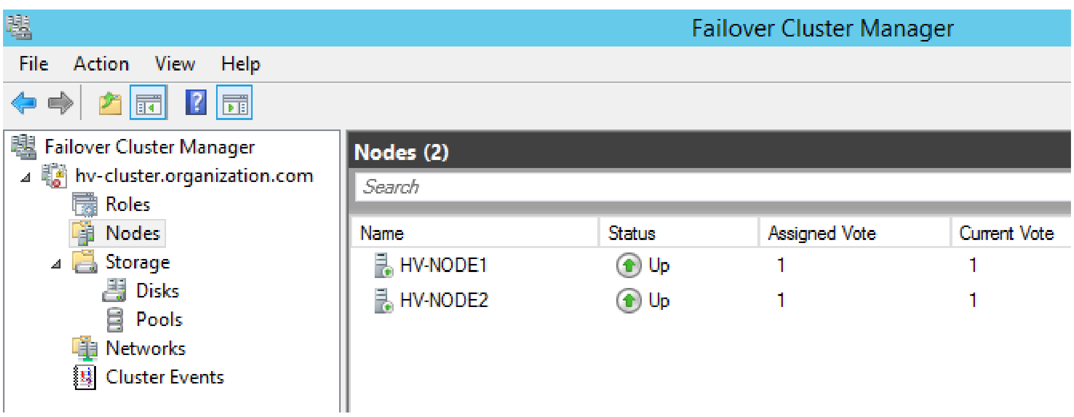

Figure G.22 Nodes

16. Networks are shown in Figure G.23;

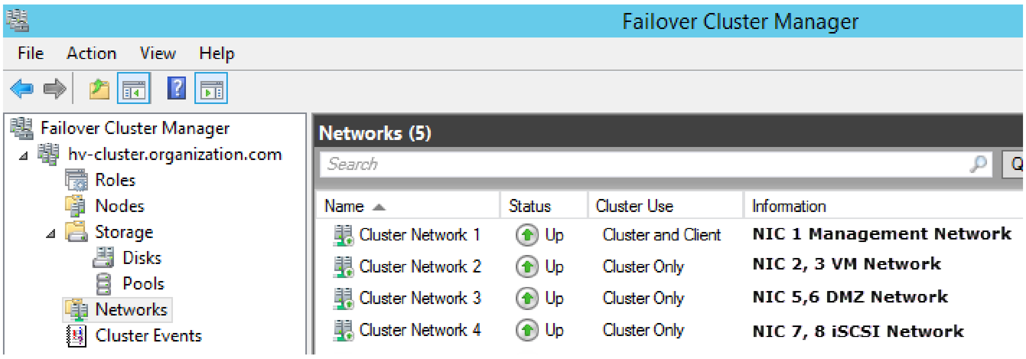

Figure G.23 Networks

17. Disks are shown in Figure G.24.

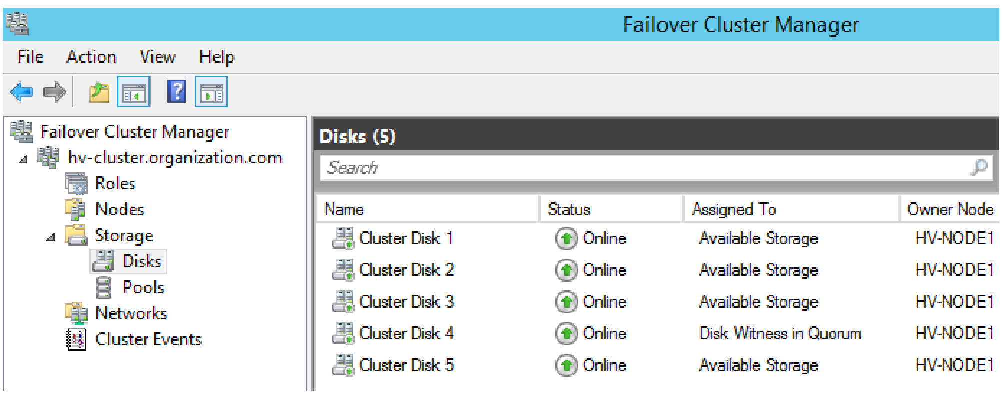

Figure G.24 Disks

G.2 - Cluster Shared Volumes Configuration

"Cluster Shared Volumes (CSV) works in a failover cluster. Without CSV, a failover cluster allows a given disk (LUN) to be accessed by only one node at a time. Given this constraint, each Hyper-V virtual machine in the failover cluster requires its own set of LUNs in order to be migrated or fail over independently of other virtual machines. In this type of deployment, the number of LUNs must increase with the addition of each virtual machine, which makes management of LUNs and clustered virtual machines more complex. In contrast, on a failover cluster that uses CSV, multiple virtual machines that are distributed across multiple cluster nodes can all access their Virtual Hard Disk (VHD) files at the same time, even if the VHD files are on a single disk (LUN) in the storage. The clustered virtual machines can all fail over independently of one another.” [Microsoft Technet(c), 2016]

In Figure G.25, it is showed hv-node1 is shown as the owner of shared disks that are not clustered yet.

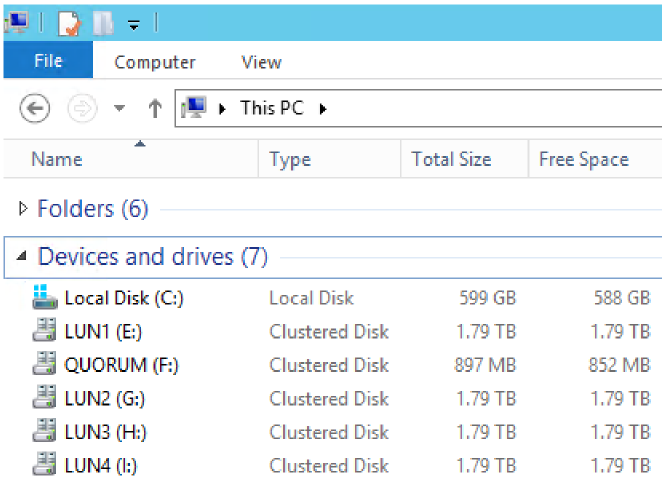

Figure G.25 Clustered disks

In Figure G.26, select all disks except Quorum and right click. Click Add to Cluster Shared Volumes.

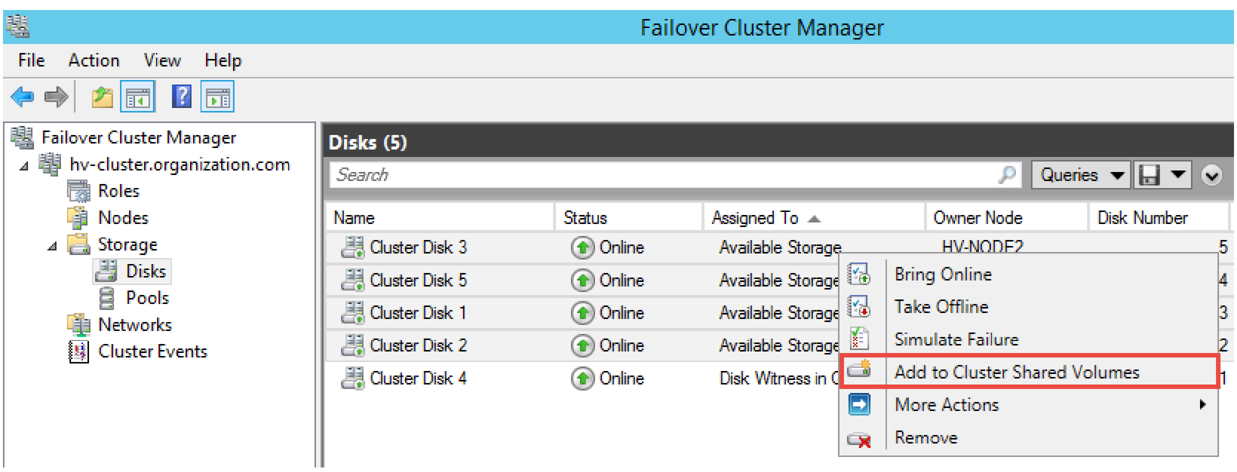

Figure G.26 Add to cluster shared volumes

In Figure G.27, all disks assigned to Cluster Shared Volume.

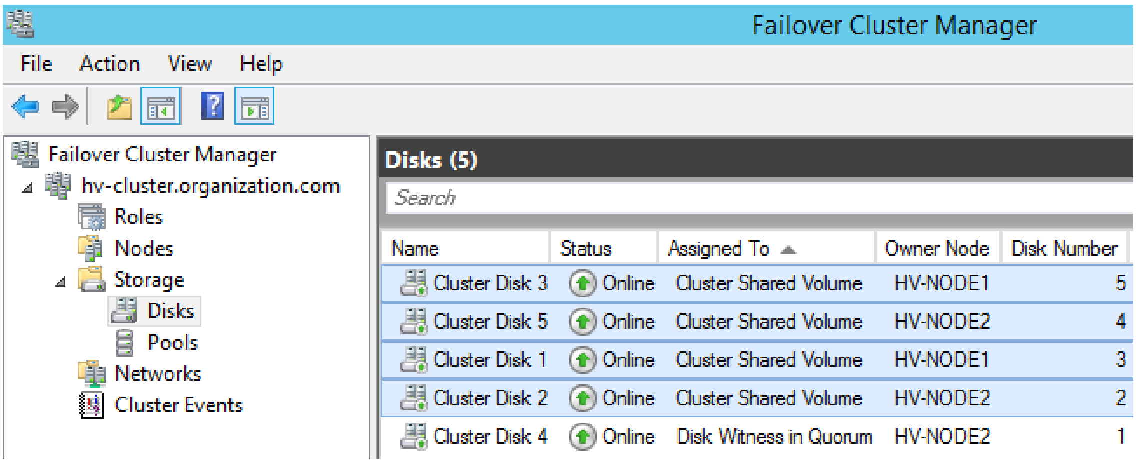

Figure G.27 Clustered shared volumes

In Figure G.28 and Figure G.29, drive letters are not seen any more for hv-node1 and hv-node2. Quorum disk is belonging to hv-node2. Cluster shared volume are not created on Quorum disk.

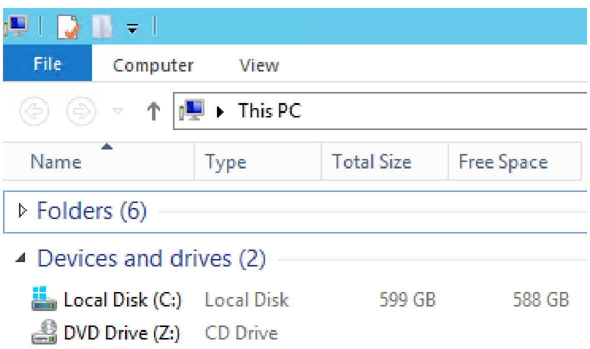

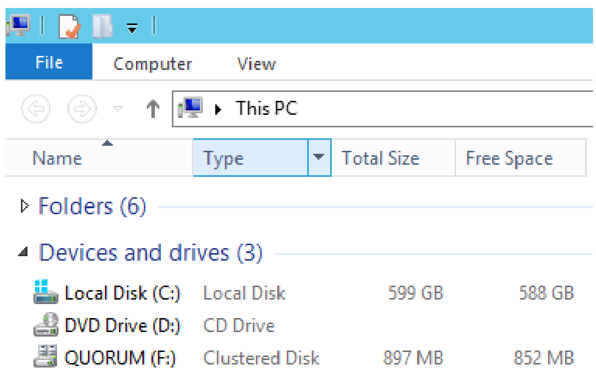

Figure G.28 Hv-node1 disks

Figure G.29 Hv-node2 disks

In Figure G.30 and Figure G.31 Disk Management for hv-node1 and hv-node2 are given.

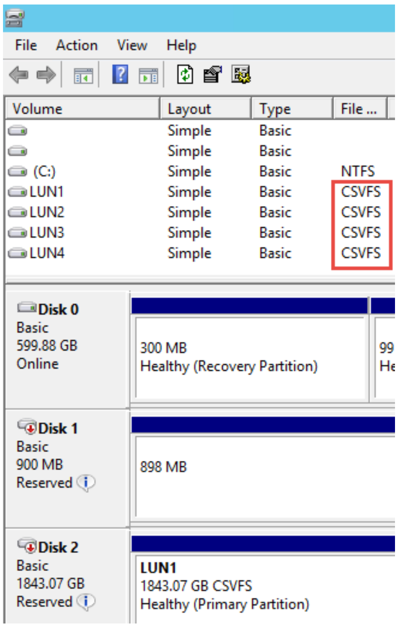

Figure G.30 Hv-node1 csvfs

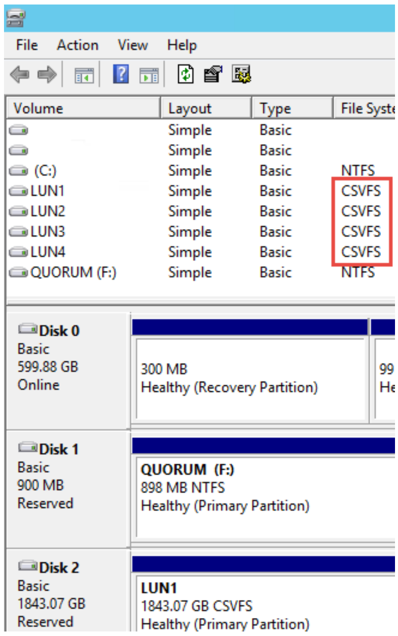

Figure G.31 Hv-node2 csvfs

In Figure G.32, Member of cluster hosts will use Volume1, Volume2, Volume3 and Volume4 to storage virtual machines in hv-node1 and hv-node2. All Volumes are mapped to clustered disks.

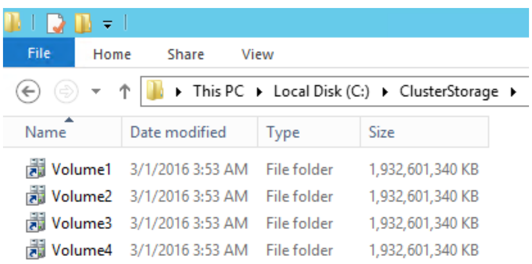

Figure G.32 Shared clustered shared volumes

G.3 - Network Configuration

One of the good feature of Hyper-V is Live Migration. “Hyper-V live migration moves running virtual machines from one physical server to another with no impact on virtual machine availability to users” [Microsoft Technet(d), 2016].

In this topology, iSCSI network shouldn't be used for Live Migration in order not to impact data performance between nodes and storage. Follow the steps below to configure Live Migration network:

1. Open Failover Cluster Manager. Click Networks and click Live Migration Settings in Figure G.33;

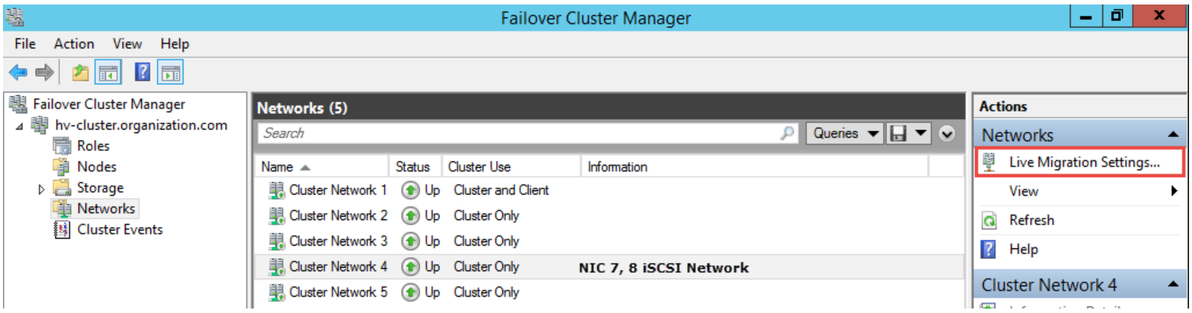

Figure G.33 Live migration settings 1

2. In Figure G.34, Live Migration Settings opens. Unselect Cluster Network 3 (DMZ) and Cluster Network 4 (iSCSI) and Click OK. Live migration operation will use firstly Cluster Network 1(Management) and if it is busy, it will use Cluster Network 2(VM).

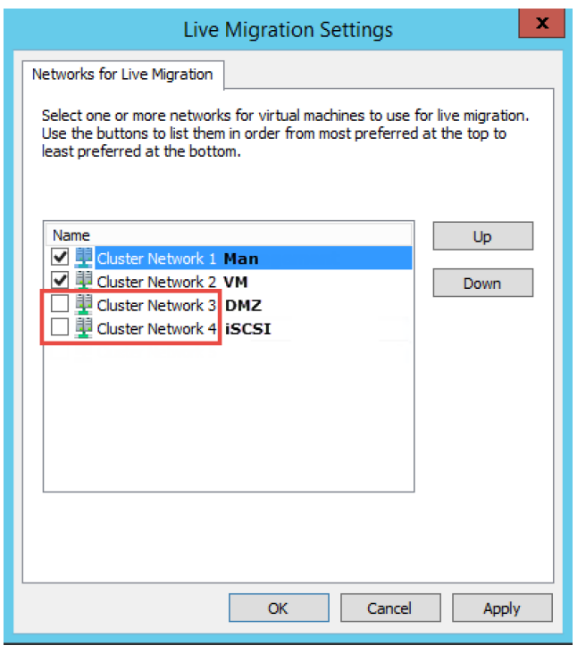

Figure G.34 Live migration settings