Received 20 March 2023 | Revised 15 May 2023 | Accepted 13 June 2023
Vol. 04, No. 01, 22 pages

RESEARCH ARTICLE TUJES Tobruk University Journal of Engineering Sciences

Implementing a Disaster Recovery Solution for Datacenters Using VMware Site Recovery Manager

Khaled Elgdamsi[1], Mohamed Embarak[2]

[1,2] Electrical and Electronic Department, Faculty of Engineering, University of Tripoli, Libya
k.elgdamsi@uot.edu.ly

Abstract. Many organizations realize the importance of implementing a robust disaster recovery (DR) solution for reasons including, but not limited to, business continuity, compliance with industry regulations, protection against disasters, ransomware, and security breaches. Traditional DR solutions are complex and may not able to provide the required levels of protection that organizations need. VMware disaster recovery solution known as site recovery manager offers policy-based management, minimizes downtime in case of disasters via automated orchestration, and enables no disruptive testing of your disaster recovery plans. In this paper, a virtual environment that simulates enterprise works environment such as banks, companies and other organizations will be created on a HPE server using a VMware platform to study disaster recovery solution provided by VMware and assessment of several cases of applications and systems that operate on this environment in terms of different specifications and validate the recovery solution from disasters when a disaster strikes and minimize downtime duration.

Keywords: Disaster Recovery (DR), VMware, Virtual Machine (VM), Downtime, Site Recovery Manager (SRM)

1 Introduction

Many IT organizations deploy servers that are only running at a fraction of their capacity, often because they are dedicating their physical server to a specific application. This is usually an inefficient mechanism because there is excess capacity not being consumed which leads to higher operating costs. In efforts to drive higher capacity utilization and reduce costs, virtualization was created. Virtualization uses software to create an abstraction layer over the physical hardware. It creates a virtual computer system, known as virtual machines (VMs). This allows organizations to run multiple virtual computers, operating systems, and applications on a single physical server (essentially partitioning it into multiple virtual servers). One of the main advantages of virtualization is a more efficient use of the physical computer hardware; this, in turn, provides a greater return on a company's investment [1].

Disasters can inflict many types of damage with varying levels of severity, depending on the scenario. A brief network outage could result in frustrated customers and some loss of business to an e-commerce system. A hurricane or tornado could destroy an entire manufacturing facility, data center or office. The monetary costs can be significant, the Uptime Institute's Annual Outage Analysis 2021 report estimated that 40% of outages or service interruptions in businesses cost between $100,000 and $1 million, while about 17% cost more than $1 million. A data breach can be more expensive; the average cost in 2020 was $3.86 million, according to the 2020 Cost of a Data Breach Report by IBM and the Ponemon Institute.

Thinking about disasters before they happen and creating a plan for how to respond can provide many benefits. It raises awareness about potential disruptions and helps an organization to prioritize its mission-critical functions. Disaster recovery is an organization's method of regaining access and functionality to its IT infrastructure after events like a natural disaster or cyber-attack. A variety of disaster recovery (DR) methods can be part of a disaster recovery plan. DR is one aspect of business continuity [2]. It relies upon the replication of data and computer processing in an off-premises location not affected by the disaster. When servers go down because of a natural disaster, equipment failure or cyber-attack, a business needs to recover lost data from a second location where the data is backed up. Ideally, an organization can transfer its computer processing to that remote location as well in order to continue operations.

SRM can be a cost-effective disaster recovery solution compared to all the recovery solutions available. It will also cover Recovery Point Objective and Recovery Time Objective. The SRM works on two different replication methodologies that is vSphere replication and Array based replications. Site Recovery Manager Server operates as an extension to the vCenter Server at a site and is compatible with other VMware solutions, besides with third-party software. You can run other VMware solutions such as vCenter Update Manager, vCenter Server Heartbeat, VMware Fault Tolerance, vSphere Storage vMotion, and vSphere Storage DRS in deployments that you protect using Site Recovery Manager [3].

VMware site recovery manager is a business continuity and disaster recovery solution that helps to plan, test, and run the recovery of virtual machines between a protected vCenter Server and recovery vCenter server in DR site. SRM brings the powerful orchestration capabilities, simplifies the deployment in case of a DR event, and increases the reliability of the DR solution with non-disruptive testing. For large businesses, providing a high degree of availability is as important as providing an overall great hosting service. Disaster is always being unpredictable, the destruction caused by it is always worse than expected. Sometimes it ends up with the loose of information, data, and records. Disaster can also make services inaccessible for very long time if disaster recovery was not planned properly. This paper focuses on protecting a vSphere virtual datacenters using Site Recovery Manager (SRM) [4].

2 Literature review

While a major number of papers focus on security and privacy with regard to Cloud Computing, the number of papers focusing on disaster recovery (DR) is relatively slight. Existing literature on DR process is presented as follows. Pokharel et al [5] explain the drawback of existing DR practices, and suggest that factors such as high cost, loss of energy, undesirable downtime, underutilization of infrastructure and low reliability and availability prevent successful DR process delivery. Subashini and Kavitha [6] explain the significance of Cloud security including disaster recovery. They acknowledge the importance of DR process and only provide the overview and the related literature. There are no any empirical studies and thus their proposal does not have a strong support and recommendation. Snedaker [7] focuses on the organization's strategies, steps and policies to achieve DR and suggest ways for organizations to restore their data over a consolidated process built in place. Sengupta and Annerrvaz [8] present their multi-site DR data distribution, including their system architecture, theories, data center details and costs involved in the DR process. The recent studies explains that Traditional DR solutions often fail to meet business requirements because they are too expensive, complex, and unreliable. Organizations using Site Recovery Manager ensure highly predictable RTOs at a much lower cost and level of complexity.

Our novel contribution through this paper is to validate a new empirical approach for disaster recovery, using a software known as hypervisor, by identifying its advantages and types, especially VMware hypervisor, which will be used in this paper , with central monitoring and resource management software known as vCenter, which enables monitoring and management of several ESXi hosts from a single control interface.

2 What is virtualization?

Virtualization is the creation of a virtual version of something, such as an operating system (OS), a server, a storage device or network resources [9]. Virtualization uses software that simulates hardware functionality to create a virtual system. This practice allows IT organizations to operate multiple operating systems, more than one virtual system and various applications on a single server. The benefits of virtualization include greater efficiencies and economies of scale [10].

OS virtualization is the use of software to allow a piece of hardware to run multiple operating system images at the same time. The technology got its start on mainframes decades ago, allowing administrators to avoid wasting expensive processing power. In more practical terms, imagine that there are 3 physical servers with individual dedicated purposes. One is a mail server, another is a web server, and the last one runs internal legacy applications as shown in figure (1).

Each server is being used at about 30% capacity (just a fraction of their running potential) but since the legacy apps remain important to the internal operations, it has to keep them and the third server that hosts them.

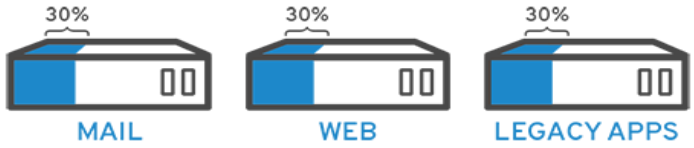

Fig. 1. Traditional Architecture

Traditionally, it was often easier and more reliable to run individual tasks on individual servers: 1 server, 1 operating system, 1 task. It was not easy to give 1 server multiple brains. Nevertheless, with virtualization, it can split the mail server into two unique ones that can handle independent tasks so the legacy apps can be migrated as shown in figure (2). It is the same hardware; just using more of it more efficiently.

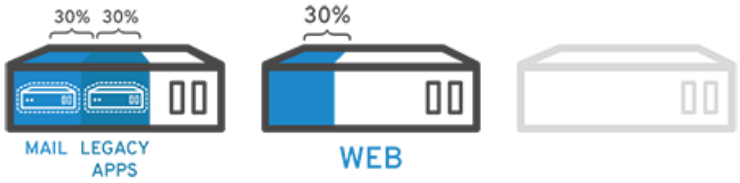

Fig. 2. Virtual Architecture

Keeping security in mind, it could split the first server again so it could handle another task increasing its use from 30%, to 60%, to 90%. Once you do that, the now empty servers could be reused for other tasks or retired altogether to reduce cooling and maintenance costs [11].

3 Virtualization Mechanism

Software called hypervisors separates the physical resources from the virtual environments. Hypervisors can sit on top of an operating system (like on a laptop) or be installed directly onto hardware (like a server), which is how most enterprises virtualize. Hypervisors take your physical resources and divide them up so that virtual environments can use them. Figure (3) shows how hypervisors convert traditional server to virtual servers.

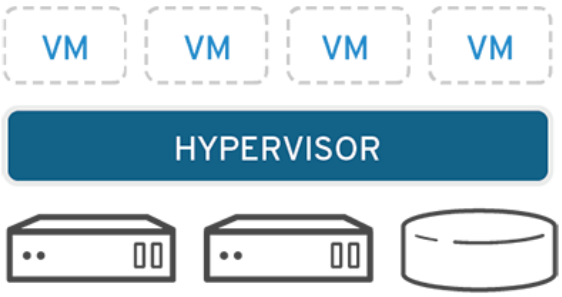

Fig. 3. Hypervisor Software

Resources are partitioned as needed from the physical environment to the many virtual environments. Users interact with and run computations within the virtual environment (typically called a guest machine or virtual machine). When the virtual environment is running and a user or program issues an instruction that requires additional resources from the physical environment, the hypervisor relays the request to the physical system and caches the changes (which all happens at close to native speed) [10].

Types of Hypervisors

There are two main hypervisor types, referred to as type 1 (or bare metal) and type 2 (or hosted). A type 1 hypervisor acts like a lightweight operating system and runs directly on the host's hardware, while a type 2 hypervisor runs as a software layer on an operating system, like other computer programs.

The most commonly deployed type of hypervisor is the type 1 or bare-metal hypervisor, where virtualization software is installed directly on the hardware where the operating system is normally installed. Because bare-metal hypervisors are isolated from the attack-prone operating system, they are extremely secure. In addition, they generally perform better and more efficiently than hosted hypervisors. For these reasons, most enterprise companies choose bare-metal hypervisors for data center computing needs [9]. There are several types for type 1 and type 2 hypervisors as shown below:

Type 1 hypervisors:

- VMware ESXi: These hypervisors offer advanced features and scalability, and it is considered the leader in the Type-1 hypervisors. In this project, we will be working on an environment that runs on ESXi hypervisor.
- Microsoft Hyper-V.
- Citrix XenServer.
- Oracle VM.

Type 2 hypervisor:

- VMware Workstation/Fusion/Player.
- Microsoft Virtual PC.
- Oracle VM VirtualBox.
- Red Hat Enterprise Virtualization.

4 VMware ESXi Hypervisor

VMware ESXi, also called VMware ESXi Server, is a bare-metal hypervisor developed by VMware for vSphere. ESXi is one of the primary components in the VMware infrastructure software suite. ESXi is a type 1 hypervisor, meaning it runs directly on system hardware without the need for an OS. Hypervisors help run multiple VMs efficiently on a physical server as shown in figure (4).

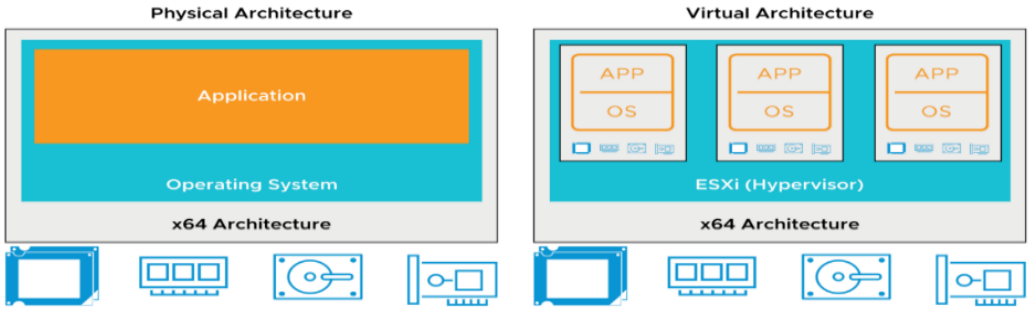

Fig. 4. VMware ESXi Hypervisor

VMware vSphere is VMware virtualization platform, which transforms data centers into aggregated computing infrastructures that include CPU, storage, and networking resources. vSphere manages these infrastructures as a unified operating environment, and provides you with the tools to administer the data centers that participate in that environment. The two core components of vSphere are ESXi and vCenter Server as shown in figure (5) [12].

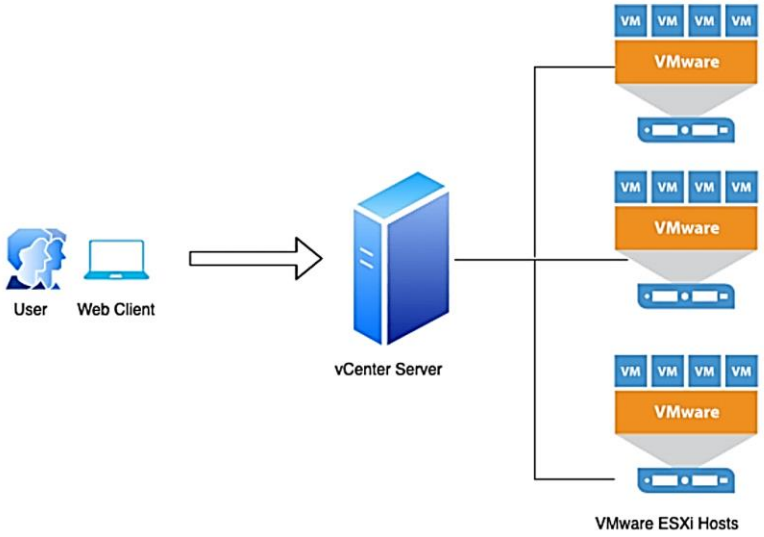

Fig. 5. VMware vSphere Components

ESXi is the virtualization platform where you create and run virtual machines and virtual appliances. vCenter Server is the service through which you manage multiple hosts connected in a network and pool host resources

5 VMware vCenter Server

VMware vCenter Server is the centralized monitoring and resource management software for VMware vSphere virtual infrastructure. It performs a number of tasks, including resource provisioning and allocation, performance monitoring, workflow automation and user privilege management. It enables a vSphere administrator to manage multiple ESXi servers and virtual machines (VMs) through a single console [13].

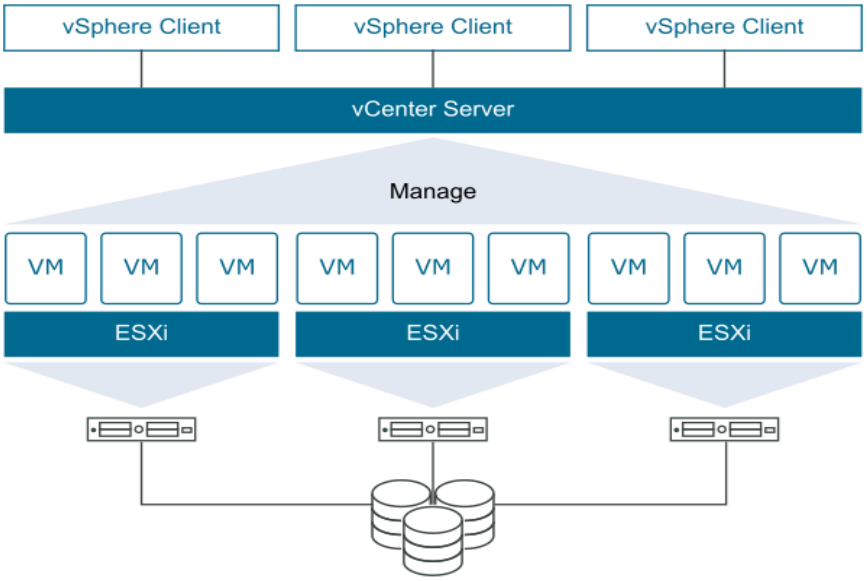

Fig. 6. VMware vCenter Server

The vCenter Server architecture consists of three main components as shown in figure (6) above: vSphere Web Client, vCenter Server Database and vCenter Single Sign-On. The vSphere Web Client is a web application that acts as the vCenter Server user interface. It enables the administrator to manage installation and handle inventory objects in a vSphere deployment and provides console access to VMs. The vCenter Server Database stores and manages server data, from inventory items to resource pools. vCenter Single Sign-On (SSO) is an authentication broker and security token that enables the user to use one login to access the entire vSphere infrastructure without further authentication. Important vCenter Server features include the following:

- **Multi-hypervisor management**. VMware vCenter Server offers integrated management for VMware and Microsoft Hyper-V hosts.
- **VMware Host Profiles**. This tool automates ESXi host configuration. A vSphere administrator can use Host Profiles to create a standard configuration, which serves as a sort of blueprint for all other hosts, and automate compliance to this configuration across all hosts or clusters.
- **Automatic VM restart**. VMware vCenter Server uses vSphere HA to pool VMs and their hosts into a cluster. In the event of a server failure, vSphere HA will automatically restart these VMs on other hosts within the cluster.
- **vCenter Server Linked Mode**. This feature provides an administrator with a single view of their vSphere deployment. An administrator can also use Linked Mode to connect multiple vCenter Server systems and grant them permission to share information. Linked Mode automatically replicates all new resources created by the administrator, including roles, policies and permissions, across the linked vCenter Server systems [14].

6 Disaster Recovery

Disaster recovery (DR) is an organization's method of regaining access and functionality to its information technology infrastructure after events like a natural disaster, or cyber-attack. A variety of disaster recovery methods can be part of a disaster recovery plan. DR is one aspect of business continuity [15].
Disaster recovery relies upon the replication of data and computer processing in an off-premises location not affected by the disaster. When servers go down because of a natural disaster, equipment failure or cyber-attack, a business needs to recover lost data from a second location where the data backed up. Ideally, an

organization can transfer its computer processing to that remote location as well in order to continue operations [16]. There are numerous types of disaster recovery methods, and organizations can choose from a variety of disaster recovery methods, or combine several:

- **Back-up**: This is the simplest type of disaster recovery and entails storing data off site or on a removable drive. However, just backing up data provides only minimal business continuity help, as the IT infrastructure itself is not backed up.
- **Disaster Recovery as a Service (DRaaS)**: In the event of a disaster or ransomware attack, a DRaaS provider moves an organization's computer processing to its own cloud infrastructure, allowing a business to continue operations seamlessly from the vendor's location, even if an organization's servers are down.
- **Back Up as a Service**: Similar to backing up data at a remote location, with Back Up as a Service, a third-party provider backs up an organization's data, but not its IT infrastructure.
- **Datacenter disaster recovery**: The physical elements of a data center can protect data and contribute to faster disaster recovery in certain types of disasters. For instance, fire suppression tools will help data and computer equipment survive a fire. A backup power source will help businesses sail through power outages without grinding operations to a halt. Of course, none of these physical disaster recovery tools will help in the event of a cyber-attack.
- **Cold Site**: In this type of disaster recovery, an organization sets up a basic infrastructure in a second, rarely used facility that provides a place for employees to work after a natural disaster or fire.
- **Hot Site**: For the facility, that has more than one branch operational, and one of the branches will be chosen as the disaster recovery site.

The two most important benefits of having a disaster plan in place, including effective DR software, are:

- Cost savings: Planning for potential disruptive events can save businesses hundreds of thousands of dollars and even mean the difference between a company surviving a natural disaster or folding.
- Faster recovery: Depending on the disaster recovery strategy and the types of disaster recovery tools used, businesses can get up and running much faster after a disaster, or even continue operations as if nothing had happened [15][16] [17].

7.1 VMware Site Recovery Manager

VMware Site Recovery Manager (SRM) is an extension to VMware vCenter that provides disaster recovery and site migration to ensure a business continuity. SRM helps to plan, test, and run the recovery of virtual machines between a protected vCenter server site and a recover vCenter server site.

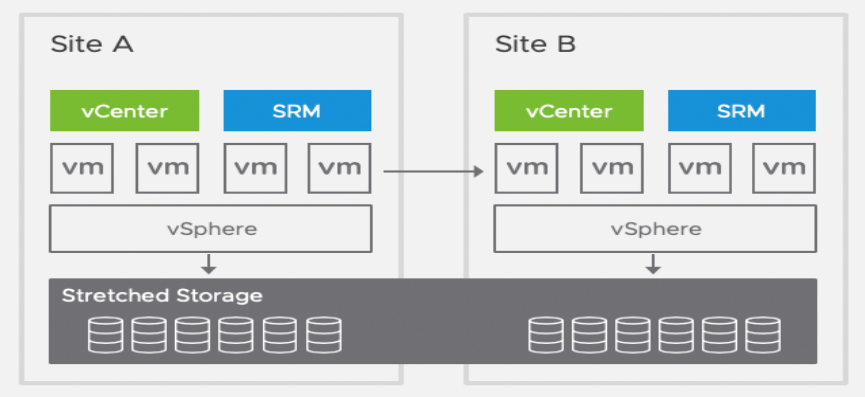

Fig. 7. VMware Site Recovery

SRM able to be used toward implement different types of recovery from the protected site to the recovery site. It orchestrates the recovery process with the replication mechanisms, to minimize data loss and system downtime.

SRM servers coordinate the operations of the VMware vCenter server at two sites as shown in figure (7). This is so that as virtual machines at the protected site are shut down, copies of these virtual machines at the recovery site startup. By using the data replicated from the protected site, these virtual machines assume responsibility for providing the same services. as shown in figure (8).

A recovery plan specifies the order in which virtual machines start up on the recovery site, also specifies network parameters, such as IP addresses, and can contain user -specified scripts that SRM can run to perform custom recovery actions on virtual machines. SRM gives the possibility to test recovery plans by using a temporary copy of the replicated data in a way that does not disrupt ongoing operations at either site [18].

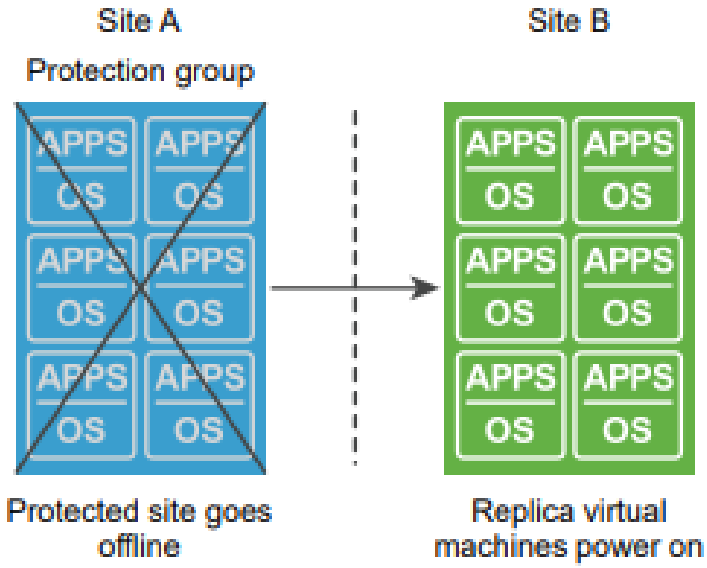

Fig. 8. Site Recovery Manager Protected and Recovery Sites

7.2 VMware vSphere Replication

VMware vSphere Replication is an extension to VMware vCenter server that provides a hypervisor-based virtual machine replication and recovery as shown in figure (9). vSphere Replication is an alternative to storage-based replication. It protects virtual machines from partial or complete site failures by replicating the Virtual Machines Disk file (VMDK) between two sites. VMDK file is the actual raw disk file that is created for each virtual hard drive [19].

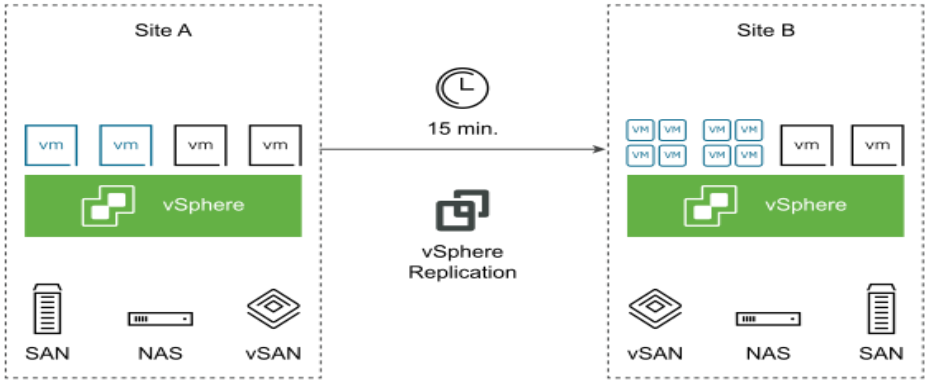

Fig. 9. VMware vSphere Replication

7.3 Site Recovery Manager with vSphere Replication

Site Recovery Manager works in conjunction with VMware vSphere Replication to automate the process of migrating, recovering, testing, re-protecting, and failing-back virtual machine workloads. Figure (10) shows VMware disaster recovery solution which consists of SRM and vSphere replication

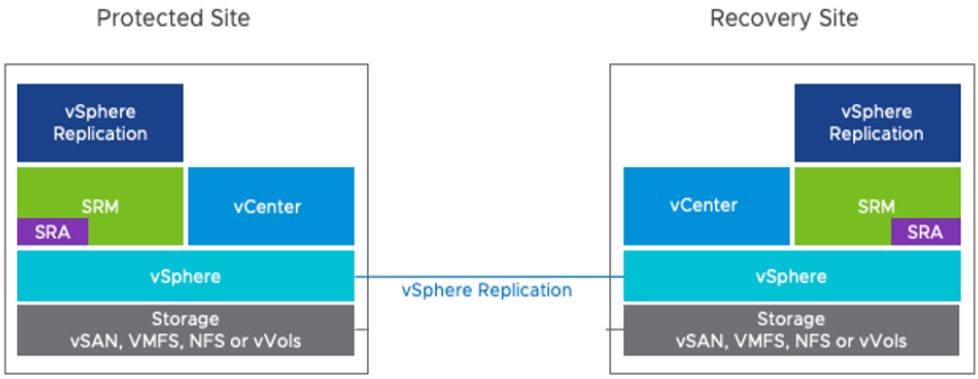

Fig. 10. SRM with vSphere Replication

Migration of protected inventory and services from one site to the other is controlled by a recovery plan that specifies the order in which virtual machines are shut down and started up, the resource pools to which they are allocated, and the networks they can access. Site Recovery Manager enables the testing of recovery plans, using a temporary copy of the replicated data, and isolated networks in a way that does not disrupt ongoing operations at either site. Multiple recovery plans can be configured to migrate individual applications and entire sites providing finer control over what virtual machines are failed over and failed back. This also enables flexible testing schedules.

Site Recovery Manager is deployed in a paired configuration, for example, a protected site and a recovery site. The SRM is deployed as an appliance at both sites. It supports multiple versions of vCenter at either site. There must be one or more ESXi hosts running version 6.5 or higher at each site. Figure (11) shows site recovery manager architecture with vSphere replication [18]. Infrastructure services like Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP), and Network Time Protocol (NTP) are required in both sites or at least at the recovery site. SRM supports protection for up to 5,000 virtual machines and is able to simultaneously run up to 10 recovery plans containing up to 500 virtual machines.

SRM is managed using an HTML5 based User Interface (UI). During the installation of SRM, a plugin labeled "Site Recovery Manager" is installed in the vSphere Web Client and an icon labeled "Site Recovery" is displayed as shown in figure (12) [20].

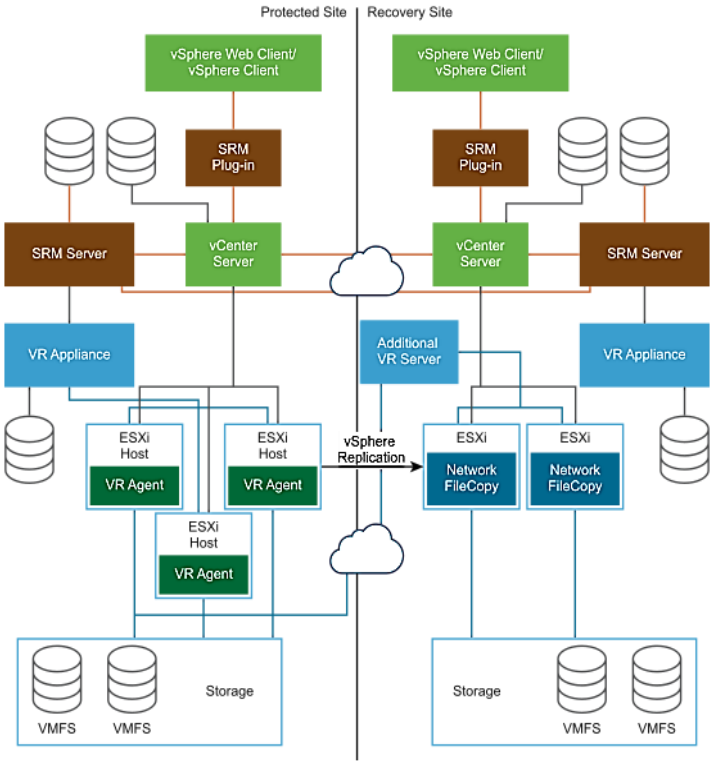

Fig. 11. SRM Architecture with vSphere Replication

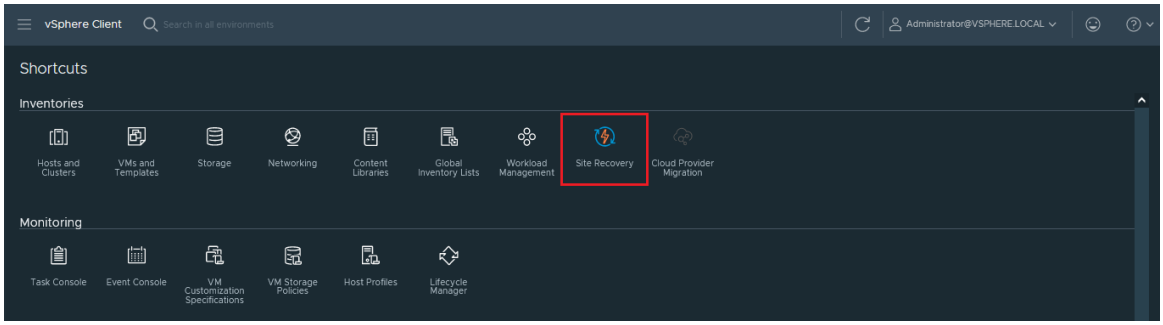

Fig. 12. Site Recovery icon in vCenter Serve

Disaster recovery or an unplanned failover is what SRM was specifically designed to accomplish. This is the most critical but least frequently used for SRM. Unexpected site failures don't happen often but when they do fast recovery is critical to business. Site Recovery Manager can help in this situation by automating and orchestrating the recovery of critical business systems for partial or full site failures ensuring the fastest Recovery Time Objective (RTO) [4].

The most common way SRM is used on a regular basis is for movement of virtual machines and applications between sites. This can be for datacenter relocation, global load balancing or planned site maintenance [20]. SRM can be used in a number of different failover scenarios depending on the requirements, constraints and objectives:

- Active-Passive: In the traditional active-passive scenario there is a production site running applications and services and a secondary or recovery site that is idle until needed for recovery. This topology is common and though it provides dedicated recovery resources it means paying for a site, servers and storage that are not utilized much of the time.
- Active-Active: SRM can be used in a configuration where low-priority workloads such as test and development run at the recovery site and are powered off as part of the recovery plan. This allows for the utilization of recovery site resources as well as sufficient capacity for critical systems in case of a disaster.
- Bi-Directional: In situations where production applications are operating at both sites SRM supports protecting virtual machines in both directions, virtual machines at site A protected at site B and virtual machines at site B protected at site A [20].

The first step to deploy and configure SRM is site pairing. The most common configuration is pairing two sites, though as was outlined in the previous section on topologies.

Next step is inventory mappings. There are multiple types of inventory mappings in SRM: resource mappings, folder mappings, and network mappings. These mappings provide default settings for recovered virtual machines. For example, a mapping can be configured between a network port group named Production-100 at the protected site and a network port group named Production-200 at the recovery site. As a result of this mapping, virtual machines connected to Production-100 at the protected site will, by default, automatically be connected to Production-200 at the recovery site.

For each protected virtual machine Site Recovery Manager creates a placeholder virtual machine at the recovery site. Placeholder virtual machines are contained in a data store and registered with the vCenter Server at the recovery site [18].

SRM offers a choice of replication technologies. Virtual machines can be replicated with vSphere Replication or with other replication technologies such as array-based replication and virtual volumes replication [4].

To use vSphere replication requires deployment and configuration of the vSphere Replication appliance. This is done independently of Site Recovery Manager. vSphere replication is able to utilize any storage supported by vSphere so there is no requirement for storage arrays, similar or otherwise, at either site [20].

One of the most important components in SRM is the protection groups. Protection groups are a way of grouping virtual machines that will be recovered together. In many cases, a protection group will consist of the virtual machines that support a service or application such as email or an accounting system. For example, an application might consist of a two-server database cluster, three application servers, and four web servers. In most cases, it would not be beneficial to failover part of this application, only two or three of the virtual machines in the example, so all nine virtual machines would be included in a single protection group. Figure (13) shows examples of protection groups.

Fig. 13. Protection groups

A protection group contains virtual machines whose data has been replicated by vSphere replication and other replication technologies. Before a protection group can be created, replication must be configured. A protection group cannot contain virtual machines replicated by more than one replication solution and a virtual machine can only belong to a single protection group.

The other main component in SRM is recovery plans. Recovery plans in SRM are like an automated run book, controlling all the steps in the recovery process. It is the level at which actions like failover, planned migration, testing and re-protect are conducted. A recovery plan contains one or more protection groups and a protection group can be included in more than one recovery plan. This provides for the flexibility to test or recover an application by itself and also test or recover a group of applications or the entire site.

In figure, 14 below there are three protection groups: Web App, Email and SharePoint. And there are three recovery plans: The Web App recovery plan containing the Web App protection group, the Email recovery plan containing the Email protection group, and the Whole Site recovery plan containing all three protection groups [20].

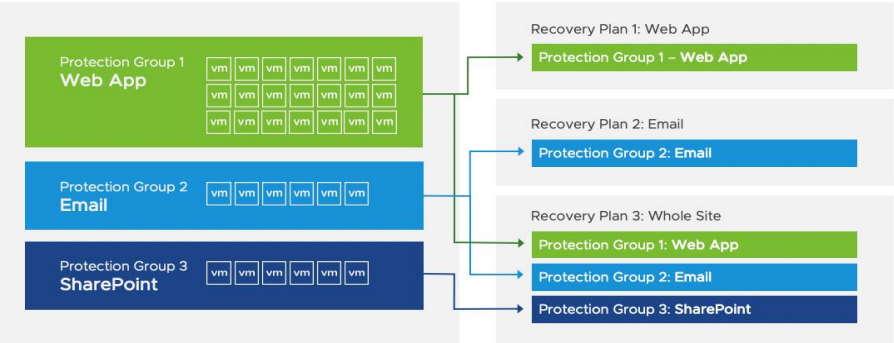

Fig. 14. Recovery Plan

The most commonly modified virtual machine recovery property is IP customization. The majority of organizations have different IP address ranges at the protected and recovery sites. When a virtual machine is failed over, SRM can automatically change the network configuration (IP address, default gateway, etc.) of the virtual network interface card in the virtual machine. This functionality is available in both failover and failback operations.

There are multiple IP customization modes in SRM. For example, it is possible to create an IP customization rule that maps one range of IP addresses to another. In figure 15 below, an administrator has mapped 10.10.10.0/24 to 198.168.100/24.

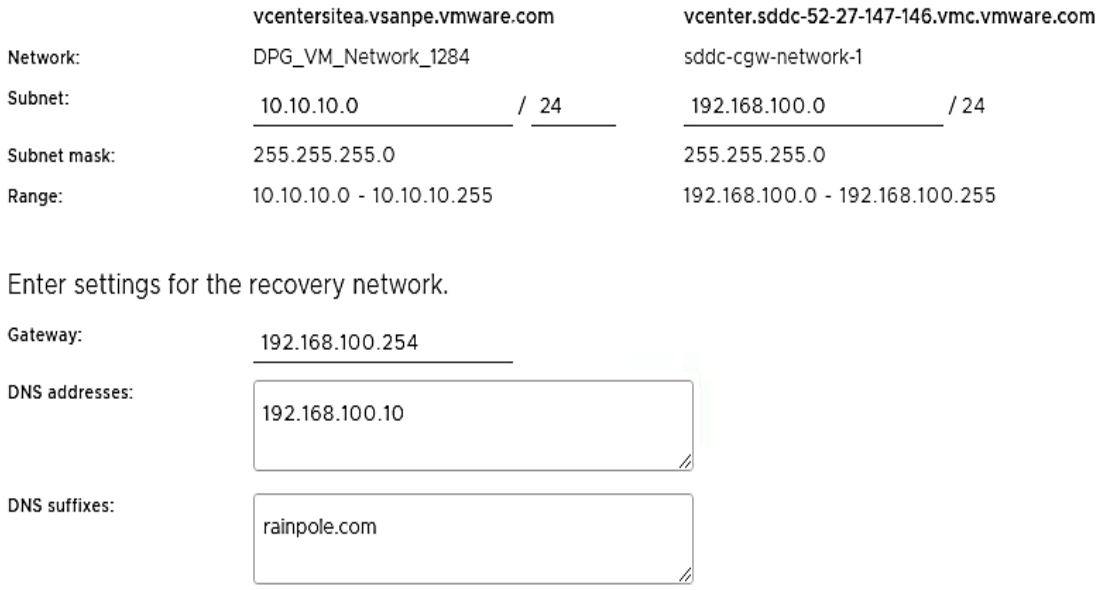

Fig. 15. IP Customization

After creating a recovery plan, it is beneficial to test the recovery plan to verify it works as expected. SRM features a non-disruptive testing mechanism to facilitate testing at any time. It is common for an organization to test a recovery plan multiple times after creation to resolve any issues encountered the first time the recovery plan was tested.

When run is complete, a recovery plan must be cleaned up. This operation powers off virtual machines and removes snapshots associated with the run. Once the cleanup workflow is finished, the recovery plan is ready for testing or running [18].

Running a recovery plan differs from testing a recovery plan. Testing a recovery plan does not disrupt virtual machines at the protected site. When running a recovery plan, SRM will attempt to shut down virtual machines at the protected site, or cross-vCenter vMotion them, if conducting a planned migration and utilizing stretched storage, before the recovery process begins at the recovery site. Recovery plans are run when a disaster has occurred and failover is required or when a planned migration is desired [4].

A recovery plan cannot be immediately failed back from the recovery site to the original protected site. The recovery plan must first undergo a re-protect workflow. This operation involves reversing replication and setting up the recovery plan to run in the opposite direction [20].

7 Methodology

This section will describe the emulation of a data center infrastructure by running a VMware platform, using HPE server besides the procedures to obtain high availability on sites level and avoid disasters by implementing VMware disaster recovery solution on HPE server.

The following flowchart outlines the implementation steps of creating two ESXi hosts to represent the datacenter in the protected site named DC-A and the recovery site named DC-B. A vCenter is then deployed to manage each ESXi host, and link each site's vCenter into a single control interface using the Enhanced Linked Mode feature. Then configure four virtual machines in DC-A with the names Win-1/2/3/4 running Windows to represent the enterprise's applications and create a virtual machine in DC-B running Windows Server to provide infrastructure services represented in DNS and NTP and then start deploying the disaster recovery solution with Install SRM and vSphere replication in each data center in addition to start the virtual machines replication process.

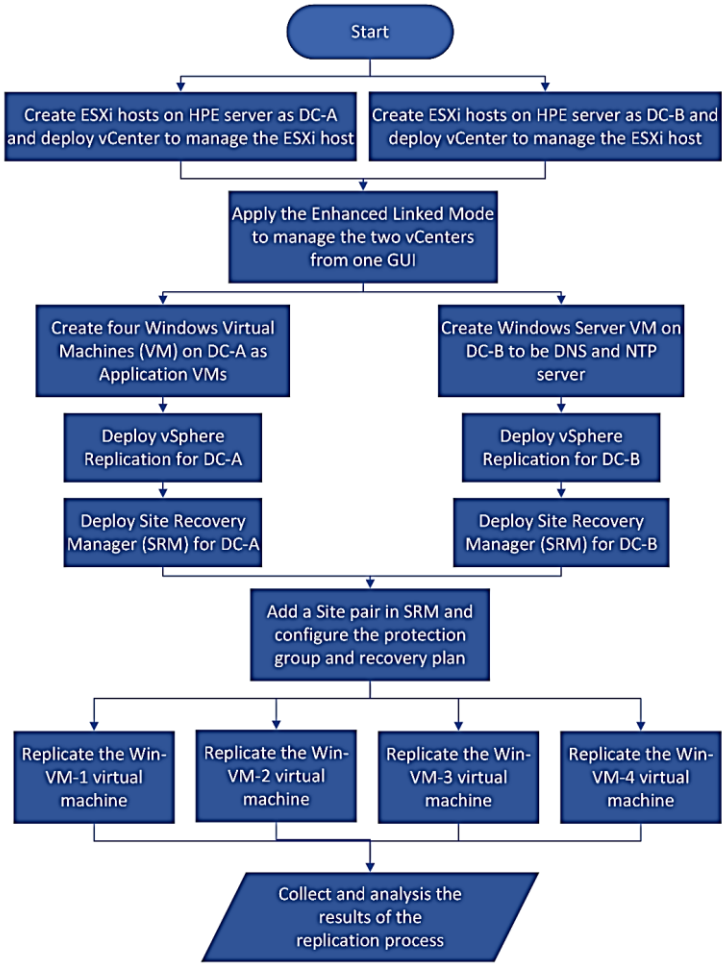

Fig. 16. Flowchart for the workflow

8.1 Hardware Component

- **HPE SimpliVity 380 server**

HPE SimpliVity 380 is a compact, scalable 2U rack-mounted building block that delivers server, storage, and storage networking services. Adaptable for diverse virtualized workloads, the secure 2U HPE ProLiant DL380 Gen10 delivers excellent performance with the right balance of expandability and scalability.

SimpliVity supports VMWare (ESXi), Microsoft (Hyper-V) hypervisors and is compatible with the Citrix Rady HCI Workspace Appliance Program. It does not have a native hypervisor. SimpliVity offers global disk deduplication, high availability, native backup, storage management, and replication. Figure (17) shown the HPE SimpliVity 380 server.

Fig. 17. HPE SimpliVity 380

- **Data Centers Architecture**

For the emulation, two data centers will be established as main data center and DR data center. The main data center is the protected site and the DR data center is recovery site. The figure (18) shows the structure of the data centers and the virtual machines that will be established in each of them. Table (1) contains the specifications of the two data centers that will be created and the resources that will be given to each virtual machine with the design IP addresses.

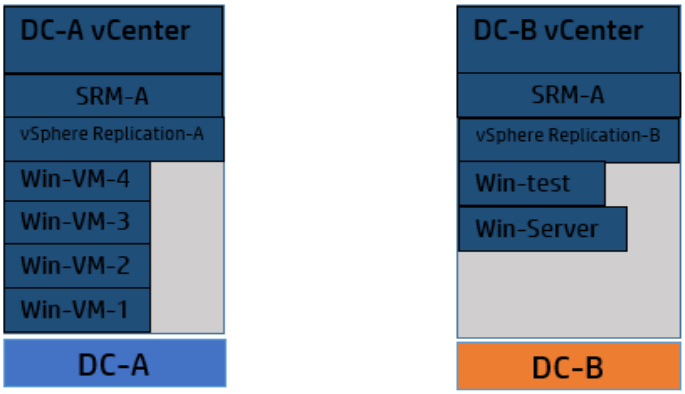

Fig. 18. Data Centers Architecture

Table 1. Virtual Machines Specifications

VM	RAM	CPU	Hard Disk	IP Address
DC-A	50GB	4	500GB	10.11.10.204(S)
DC-B	50GB	4	500GB	10.11.10.205(S)
DC-A vCenter	12GB	2	100GB	10.11.10.207(S)
DC-B vCenter	12GB	2	100GB	10.11.10.208(S)
SRM-A	12GB	4	33GB	10.11.11.226(S)
SRM-B	12GB	4	33GB	10.11.11.227(S)
vSphere Replication-A	8GB	2	33GB	10.11.11.228(S)
vSphere Replication-B	8GB	2	33GB	10.11.11.229(S)
Win-VM-1	4GB	2	48GB	10.11.11.220(S)
Win-VM-2	4GB	2	48GB	10.11.11.221(S)
Win-VM-3	4GB	2	48GB	10.11.11.222(S)
Win-VM-4	4GB	2	48GB	10.11.11.223(S)
Win-test	4GB	2	48GB	10.11.11.224(S)
Win-Server	4GB	2	90GB	10.11.11.225(S)

8.2 VMware Site Recovery Installation and Configuration

VMware Site Recovery uses the host-based replication feature of vSphere Replication and the orchestration of VMware Site Recovery Manager.

To start using VMware Site Recovery, several procedures need to be performed in a predefined order as shown in the workflow diagram in figure (19).

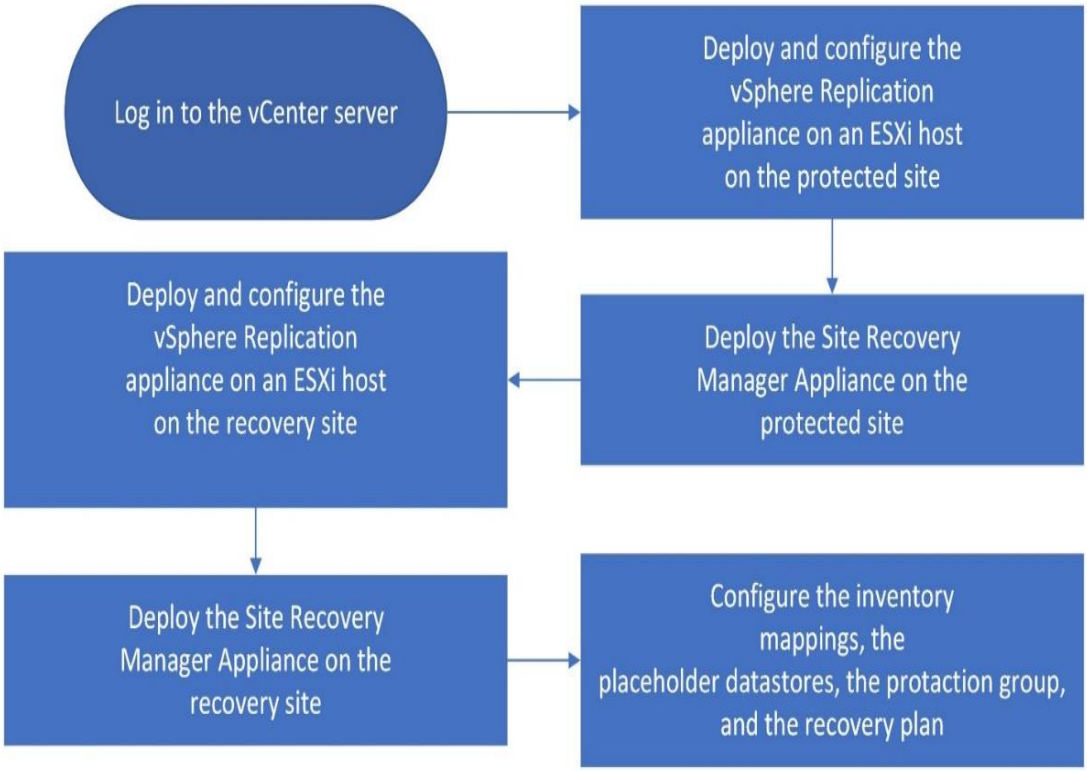

Fig.19. VMware Site Recovery Workflow

Before performing the replication process, the hard disk size of the Windows virtual machines in the main site will be increased by adding a new hard disk for each VM with different size as shown in Table (2) for comparisons of differences in the replication process in terms of sync duration and sync size.

Table 2. Virtual Machines Hard Disks Size

Virtual Machines	Hard Disk 1 Size	New Hard Disk Size	Total Size
Win-VM-1	48GB	-	48GB
Win-VM-2	48GB	50GB	98GB
Win-VM-3	48GB	100GB	148GB
Win-VM-4	48GB	160GB	208GB

8 Simulation Results of Disaster Recovery Plan

This section discusses performing systems simulation when the main data center failure occurs and activating disaster recovery by running the recovery plan, and the phases that virtual machines will go through during this process as shown in table (3). Moreover, what would materialize to the virtual machine's replication process if the hard disk capacity of VM were larger than the data store at the recovery site.

Table 3. Recovery Plan Steps

Recovery Plan Steps
1. Pre-synchronize storage
1.1. Protection Group Main-DR
2. Shut down VMs at protected site
3. Resume VMs suspended by previous recovery
4. Restore recovery site hosts from standby
5. Restore protected site hosts from standby
6. Prepare protected site VMs for migration
6.1. Protection Group Main-DR
6.1.1. Prepare VMs for migration
6.1.1.1. Win-VM-1
6.1.1.2. Win-VM-2
6.1.1.3. Win-VM-3
6.1.1.4. Win-VM-4
6.1.2. Change storage to read-only
7. Synchronize storage
7.1. Protection Group Main-DR
8. Suspend non-critical VMs at recovery site
9. Change recovery site storage to writable
10. Power on priority 4 VMs
10.1. Win-VM-1
10.2. Win-VM-2
10.3. Win-VM-3
10.4. Win-VM-4

9.1 Virtual Machines Replication Status

According to figure (20), the replication process was successful completed. Each VM's replication time was varied since each has a distinct disk size, and the time for each VM was computed as listed in the table (4**).**

Table 4. Replication Time for VMs

Virtual Machines	Replication Time	Disk Size (DC-A)	Disk Size (DC-B)
Win-VM-1	36:34 min	48GB	48GB
Win-VM-2	50:34 min	98GB	98GB
Win-VM-3	65:08 min	148GB	148GB

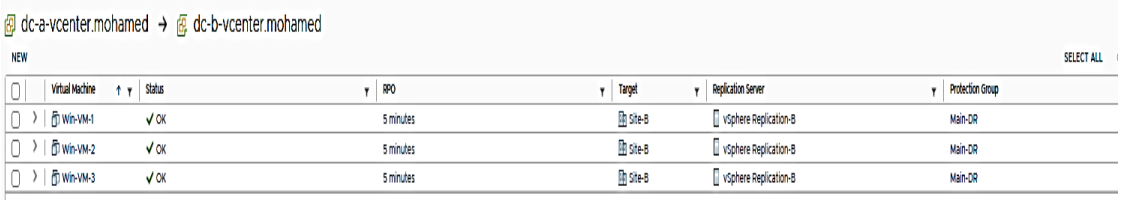

Fig. 20. Replication Status

As shown in Table (4), the disk size of the VMs in the recovery site (DC-B) is the same as the disk size in the main site (DC-A), which means that the VMs replication has been successfully completed. The following figures shows the size of the virtual machines' disks at both sites.

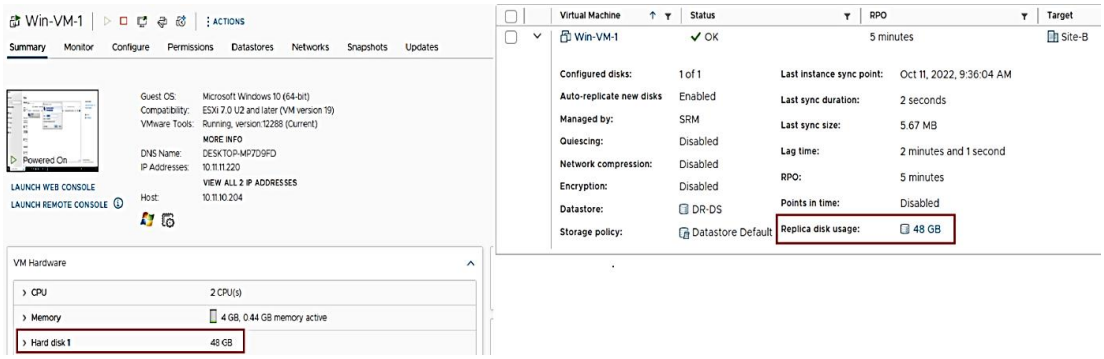

Fig. 21. Win-VM-1 Disk Size in both Sites

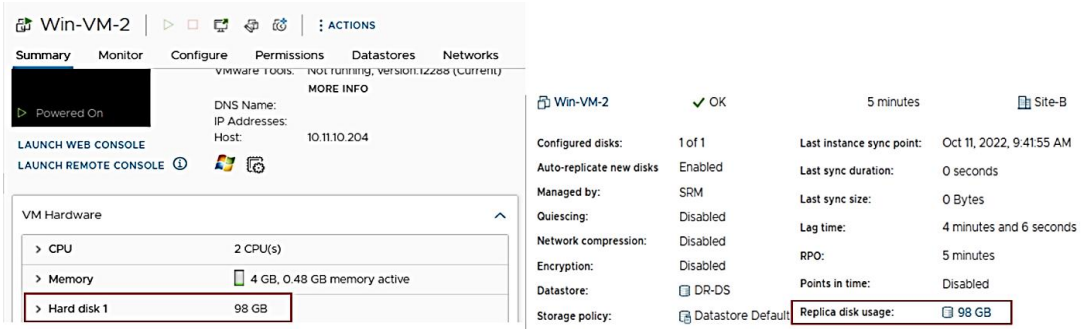

Fig. 22. Win-VM-2 Disk Size in both Sites

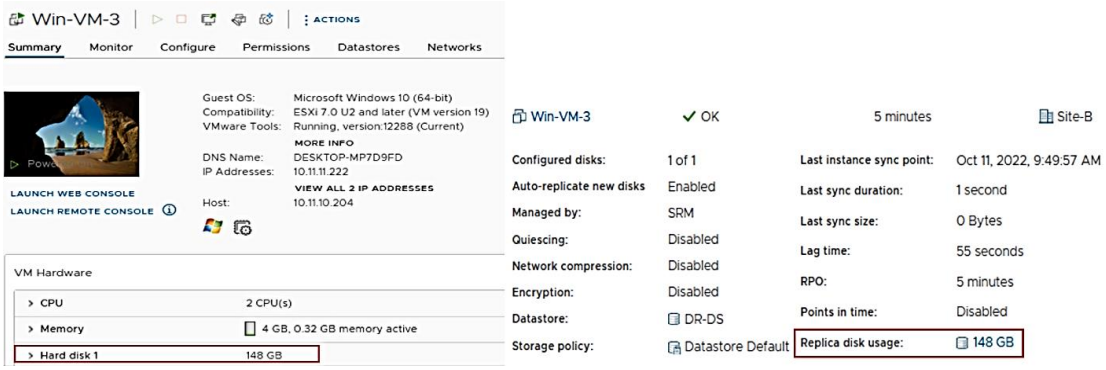

Fig. 23. Win-VM-3 Disk Size in both Sites

For Win-VM-4, the replication process will start and the process will be monitored to see if it succeeds safely as shown in figure (24).

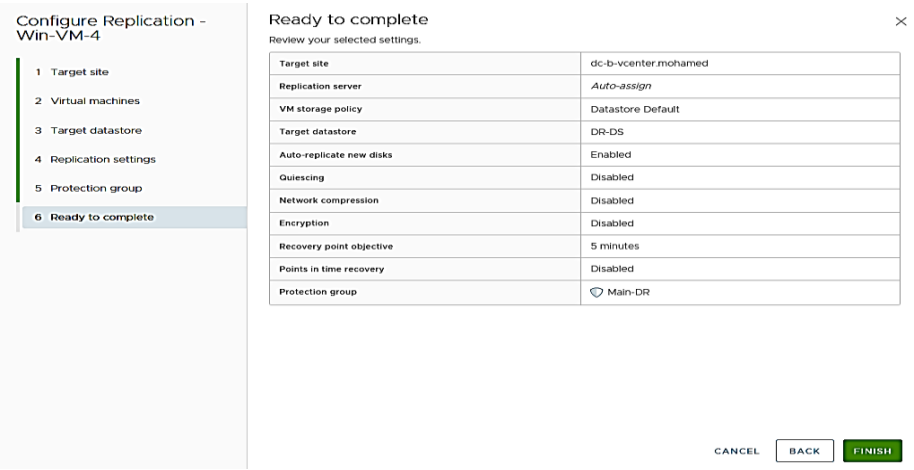

Fig. 24. Win-VM-4 Replication Process

As shown in Figure (25) below, Win-VM-4's replication fails because the data store that is used for replication does not have enough space to complete the operation. DR-DS is the data store that is used for replication process for all VMs. As shown in figure (26), the remaining space in the data center after the virtual machine is replicated is 204GB and Win-VM-4 disk size is 208GB for this reason the replication process can't be done.

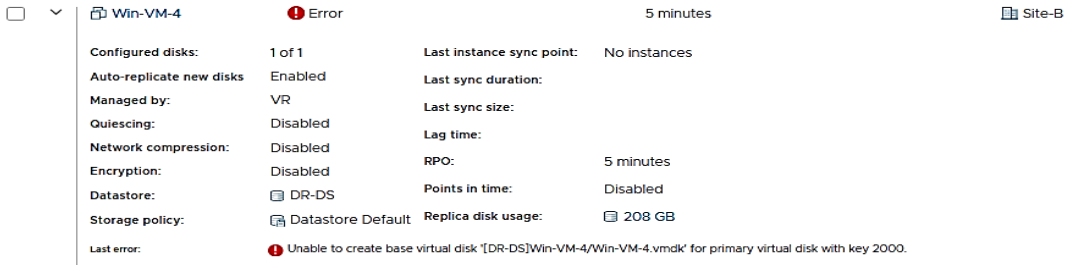

Fig. 25. Error in Replication Process

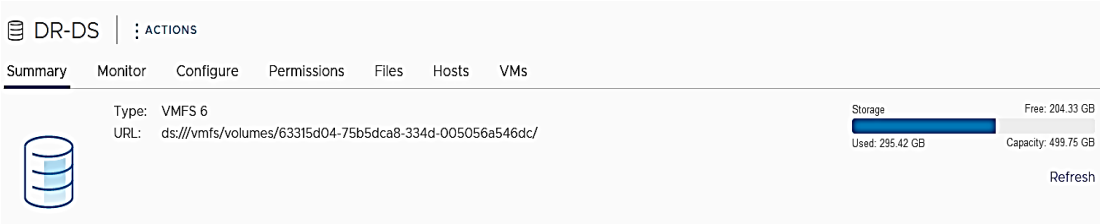

Fig. 26. DR-DS Data store Capacity

To solve this problem, the capacity of the DR-DS data store will be increased to match the needs of Win-VM-4 for replication process. 150 GB added to the data store, bringing the total capacity of DR-DS data store to 650GB as shown in figure (27).

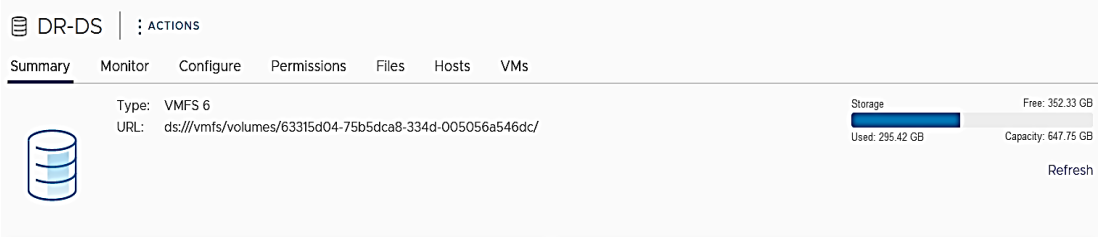

Fig. 27. DR-DS Data store Capacity

After the DR-DS data store capacity is successfully increased, the Win-VM-4 replication process has been restarted and the process continues successfully as shown in figure (28).

Fig. 28. Win-VM-4 Replication Process

The Win-VM-4 replication process was completed successfully in 78:21 minutes. As shown in the figure (29), all the VMs are now replicated in the DR site in a protected group with a recovery plan, and the disaster recovery solution is ready to use.

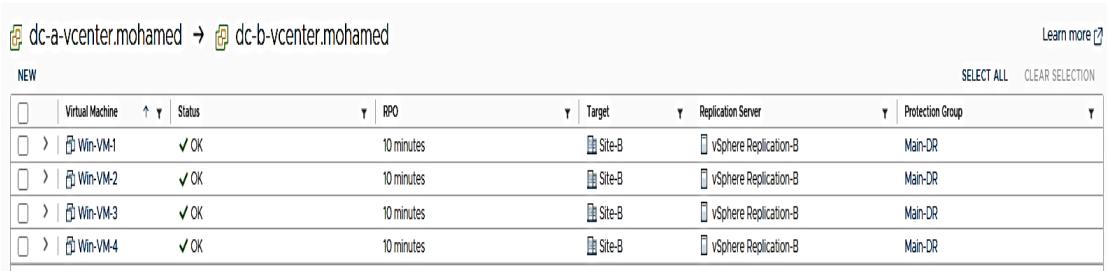

Fig. 29. Virtual Machine Replication Status

9.2 Recovery Plan Run

Recovery plan can be run under planned circumstances to migrate virtual machines from the protected site to the recovery site. If the protected site experiences an unforeseen event that might result in data loss, you can also run a recovery plan under unplanned circumstances.

When run a recovery plan to perform planned migration and disaster recovery, Site Recovery Manager makes changes at both sites that require significant time and effort to reverse. Because of this time and effort, you must assign the privilege to test a recovery plan and the privilege to run a recovery plan separately.

When a recovery plan creates, it must be tested before being used for planned migration or for disaster recovery. By testing a recovery plan, you ensure that the virtual machines that the plan protects recover correctly to the recovery site. If recovery plans not tested, an actual disaster recovery situation might not recover all virtual machines, resulting in data loss.

The Main-DR recovery plan successfully verified with elapsed time equal to 3 minutes and 32 seconds with zero error; the table below summarizes the assessment process, the status of each stage, and the time it took to complete the process.

Table 5. Main-DR Recovery Plan Run Summary

Test Summary	
Operation	Test
Storage Options	None
Started By	VSPHERE.LOCAL\Administrator
Start Time	2022-10-14 15:12:59 (UTC 0)
End Time	2022-10-14 15:16:30 (UTC 0)
Elapsed Time	00:03:32
Result	Success
Errors	0
Warnings	0
Successfully recovered VMs	4
VMs recovered with errors	0
Powered On VMs	4
Powered Off VMs	0
Successfully IP customized VMs	0
VMs with IP customization errors	0

10 Conclusion

VMware disaster recovery solution was introduced which consists of SRM and vSphere replication that provides high availability to data centers, ensure recoverability, and business continuity. Two data centers have been created running on the VMware platform that simulate a real environment by established four virtual machines running the windows operating system in the main datacenter to represent the applications and systems of any organization. In the DR data center, a windows server virtual machine created to serve as a DNS server so that each IP address has a domain name and as an NTP server so that all components are time synchronized. After implementing the disaster recovery solution and getting the results, it turns out that the duration of the replication process between virtual machines varies with the size of the hard disk memory and the solution has proven effective in terms of availability, recoverability, business continuity and reduce downtime with elapsed time equal to 3 minutes and 32 seconds without any error.

11 Future work

Some additional study could be motivated to investigate and analyze SRM through comparison performance analysis with:-

- Combining SRM with NetApp to provide recovery solutions through deployment of VMware vCenter Site Recovery Manager, with NetApp storage systems.
- SRM performance under different types of recovery, such as Planned migration or Bi-directional replication data flows

Conflict of Interest

This statement is to certify that all authors have seen and approved the manuscript being submitted and they declare no competing interest.

References

1. I. Ali and N. Meghanathan "Virtual Machines and Networks – Installation, Performance, Study, Advantages and Virtualization Options" International Journal of Network Security & Its Applications (IJNSA), Vol.3, No.1,January 2011 DOI: 10.5121/ijnsa.2011.3101 1
2. www.techtarget.com/searchdisasterrecovery/definition/disaster-recovery (accessed on 8th November 2022)
3. www.vmware.com/topics/glossary/content/disaster-recovery.html (accessed on 8th November 2022)
4. docs.vmware.com/en/Site-Recovery-Manager/index.html (accessed on 2nd November 2022)
5. M., Pokharel, S., Lee, J. S. Park, Disaster recovery for sys-tem architecture using cloud computing. In the 10th IEEE/IPSJ International Symposium on Applications and the Internet (SAINT), July 2010, 304-307.
6. S. Subashini, V. Kavitha, A survey on security issues in service delivery models of cloud computing. Journal of Network and Computer Applications, 34(1), (2011), 1-11.
7. S. Snedaker, Business continuity and disaster recovery planning for IT professionals. Newnes, 2013
8. S. Sengupta, K. M., Annervaz, Multi-site data distribution for disaster recovery—A planning framework. Future Generation Computer Systems, 41, (2014), 53
9. www.vmware.com/solutions/virtualization.html (accessed on 29th September 2022)
10. https://www.techtarget.com/searchitoperations/definition/virtualization (accessed on 29th September 2022)
11. www.redhat.com/en/topics/virtualization/what-is-virtualization (accessed on 30th September 2022)
12. www.vmware.com/products/esxi-and-esx.html (accessed on 29th September 2022)
13. www.techtarget.com/searchvmware/resources/VMware-ESXi-vSphere-and-vCenter (accessed on 29th September 2022)
14. VCenter Server and Host Management- vCenter Server 7.0 (modified on 2nd April 2021)
15. www.vmware.com/topics/glossary/content/disaster-recovery.html (accessed on 3rd November 2022)
16. Disaster Recovery as-a-Service Solution Brief (April 2020)
17. vSphere Replication for Disaster Recovery to Cloud - vSphere Replication 6.5 (VMware, Inc. 2017)
18. Site Recovery Manager Installation and Configuration (Modified on 14 FEB 2022)
19. docs.vmware.com/en/vSphere-Replication/8.4/vsphere-replication-84-admin.pdf (accessed on 3rd November 2022)
20. https://core.vmware.com/resource/site-recovery-manager-technical-overview (accessed on 3rd November 2022)