Protection of VMware Virtual Machines using Microsoft Azure Site Recovery

Vasudha N[1] Ramakrishna KT[2]

[1]Student [2]Faculty

[1]M.Sc. in Storage and Cloud Technology Jain University, Bangalore, India [2]INurture Education Solutions PVT LTD Jain University, Bangalore, India

Abstract— With the rapid growth of cloud computing paradigm, disaster recovery using cloud resources has become an attractive approach. Data Protection is on the priority in every IT industry in case of disaster, which was not proved efficient with previous methods of backup and replication. Azure Site Recovery (ASR) brings a very apt implementation helpful during the recovery of Disaster. In this paper, we explain the implementation of Azure Site Recovery in Enhanced mode. The approach is suited for protection of on-premises Virtual machines. It gives a solution on how LOB can be maintained, without any downtime with 99.9% surety.

Key words: Cloud Computing, Business Continuity, Disaster Recovery, Azure Site, VM Protection, VMware

I. Introduction

A key challenge in public and private enterprises is providing data protection using Disaster Recovery (DR) Service [1], allowing faster recovery but with price tag. Today, cloud-based DR solutions replace legacy. An important part of any organization's business continuity and disaster recovery (BCDR) strategy is to make sure on how to keep corporate workloads and applications up and running when planned and unplanned events occur.

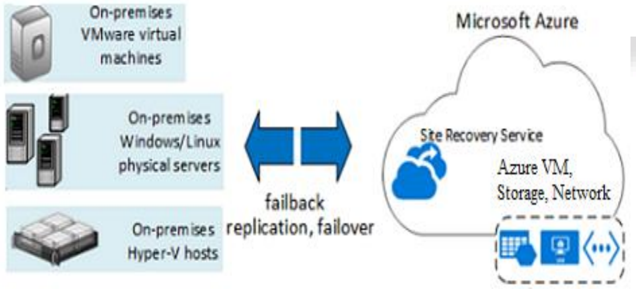
On-premises VMware virtual machines
On-premises Windows/Linux physical servers
On-premises Hyper-V hosts
failback
replication, failover
Microsoft Azure
Site Recovery Service
Azure VM, Storage, Network

Fig. 1: Replications to Azure

In such cases Azure service [2] help replication of on-premises physical servers and virtual machines to the cloud (Azure) or to a secondary site. The flexibility to the service includes applications that are running on Windows or Linux-based OS, running on physical servers, VMware or Hyper-V VMs [3]. Site Recovery helps do this by replication, failover, and recovery of workloads and applications so that they'll be available from a secondary location if your primary location goes down.

Using Azure as a destination for disaster recovery eliminates the cost and complexity maintaining a secondary site, and replicated data is stored in Azure Storage.

II. Vmware Virtual machines

A. Virtual machine (VM) is a software program which imitates a physical computer running Operating System or an application. A VM is known as a "guest" and is created within another computing environment referred as a "host." Multiple virtual machines can exist within a single host at one time.

The client operating systems, such as Microsoft Server 2008, Microsoft Small Business Server, Linux varieties, etc. are then set up as virtual machines, working directly with the VMware layer rather than with the actual hardware. This allows replacement of hardware to be very simple. If the hardware is replaced, VMware is reconfigured for the new hardware, and the virtual guest operating systems see no change whatsoever and are immediately able to boot and operate.

VMware allows the enterprise to replace many disparate, underused devices with a few virtual hosts. This greatly reduces system downtime, allows for simple movement of virtual clients from one hardware host to another and allows for scheduled hardware repair or replacement with downtime by moving those clients to another hardware host on the cluster.

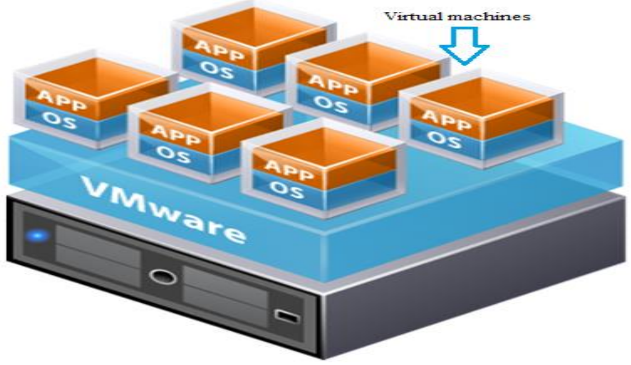

Fig. 2: Vmware Virtual Machines on a Physical Server

It also allows the IT administrator to very quickly add virtual servers as required without the need to purchase additional hardware.

Upgrading hardware becomes a simple process. Removing the requirement of the Operating System needing to work directly with the hardware makes disaster recovery or replacement of failed servers simple.

III. Disaster recovery

A disaster is a serious disruption of the functioning of a community, material, economic losses and impacts, which exceeds the ability of the affected surroundings to cope using its own resources.

Disaster recovery (DR) [4] is a set of policies and procedures to enable the recovery or continuation of vital technology infrastructure and systems following a natural or manmade disaster. Disaster recovery focuses on the IT or technology systems supporting critical business functions, which involves keeping all essential aspects of a business functioning despite significant disruptive events. Disaster recovery is hence called a subset of business continuity.

A. Traditional Disaster Recovery

Traditional disaster recovery in IT environment is using tape backup or a secondary data center. The locations are denoted as hot sites, they are locations that are already functioning with the entire infrastructure needed for a firm to continue working, and allow a firm's operations to be resumed immediately. This is the most extensive type of disaster recovery plan, as it depends on the implementation of duplicate infrastructure. Cold sites provide an alternative site from the day-to-day office if it is not usable due to an emergency. In both of these scenarios, a firm would need to install and provide either some or all of the necessary equipment and functions required to restore operations in order to continue working

B. Cloud Based Disaster Recovery

Cloud computing gives organizations the opportunity to rethink many traditional IT practices, but it may be a particularly good fit for disaster recovery and business continuity. Cloud gives them the ability to store data some place remote, store it online, and to typically recover faster than from tape. This service is now provided by many IT companies as Disaster Recovery as a service (DRaaS) [5].

The cloud gives companies backup of data, fail-over of servers, and the ability to have a secondary center far enough away to allow for regional disaster recovery. The cloud gives small and medium-sized business the same capabilities that larger companies have had for years. This approach saves costs, eliminates downtime, and uses proven failover and automated testing, an organization's DR plans will act as the insurance policy it should be.

Back up to and restore to the cloud. Is an approach where, data isn't restored back to on-premises infrastructure; instead, it's restored to virtual machines in the cloud. This requires both cloud storage and cloud compute resources. Replication to cloud virtual machines can be used to protect both cloud and on-premises production instances.

IV. MICROSOFT AZURE

Microsoft Azure, formerly known as Windows Azure, is Microsoft's public cloud computing platform. Users can pick and choose from various services to develop and scale new applications, or run existing applications, in the public cloud.

Microsoft Azure recovery services are a boon to the IT industries. The company through these services toward enterprises making the move to hybrid cloud deployments. Using Azure Backup, businesses can ensure they SharePoint, Exchange, SQL Server, Windows clients and Hyper-V virtual machines applications are always backed up. Microsoft was also the first major cloud provider to adopt the new international cloud privacy standard, ISO 27018.

A. Site

A site is a location which holds the data. In context to the recovery Services there are two things commonly talked about firstly, primary location/onpremisis or the site which holds the data on a physical hardware in a particular geographical location and secondary location or the Azure cloud which holds the replicated data in multiple locations.

B. Site Recovery Services

Site recovery orchestrates replication and failover but doesn't intercept the application data or have any information about it also contributes to application-level protection, its synchronous replication with RPOs is as low as 30 seconds to meet the needs of most critical business applications recovery plans that enables to recover an entire application stack with a single click, and include external scripts and manual actions in the plan.

Azure Site Recovery [6] provides two modes in which the recovery can be configured the prior one which was used initially that needed a continuous running of virtual machines in Azure that is the legacy mode in which the configuration server and the master target server which has to be deployed in Azure using the standard virtual machines whereas in Enhanced mode which is the latest mode and also the recommended mode that has to be implemented the difference is that the configuration and the master target server is deployed on premises.

V. PROTECTION OF VIRTUAL MACHINES

The general method of protection of virtual machines can be defined in the following ways:

- Agentless
- Light Agent

Agentless, this approach entails having a dedicated virtual machine with the antivirus engine installed on it. This machine does the malware scanning on the rest of the virtual infrastructure by connecting to the rest of the virtual machines using native VMware vShield technology. vShield also interacts with the antivirus's system management so it knows the settings and applied policies, when to turn protection on and off, how to optimize, and so on.

Light Agent, this approach, an agent which is light on its resources is installed on every protected virtual machine. This agent replaces vShield and connects the virtual machine to the antivirus engine that resides on the Security Virtual Appliance.

For maximum, high-performance protection of a Windows guest OS a *light agent* is needed. On other operating systems (Linux, OS X) a product for endpoints is better; however, application is limited by considerations of productivity and the antivirus's interaction with features of the virtual environment itself. In VMware it is possible to edit the protection properties of any virtual machine or template in a protection group. You can change the resource mappings, attached storage devices and their data stores, and other properties that control the configuration with which Site Recovery Manager recovers the virtual machine.

The Azure method of protecting virtual machines is through site recovery. The protection can be prioritized on recovery point objective (RPO) and the recovery time objective (RTO) are two very specific parameters that are closely associated with recovery.

The RTO is how long you can basically go without a specific application. This is often associated with your maximum allowable or maximum tolerable outage. The RPO is how much of data loss is allowed or afford to lose. RPO and RTO, really influence the kind of redundancy or backup infrastructure you will put together. The tighter the

RPO and RTO, the more money you will spend on your infrastructure.

VI. Architecture

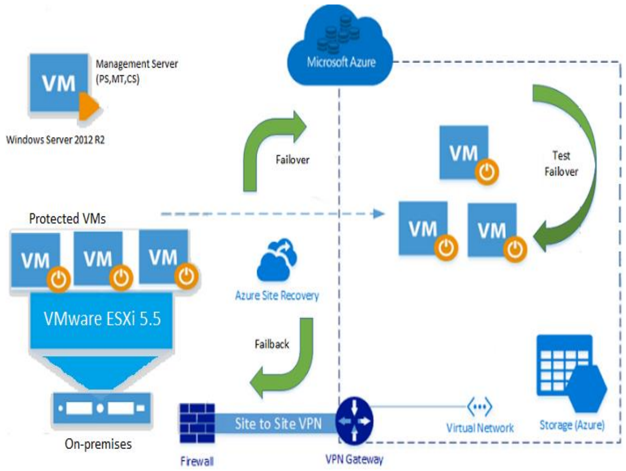

Fig. 3: Protecting VMware VMs using Azure Site Recovery

In the above diagram mentioned we can see the all the components showing the working of the Protection of VMware VMs using Azure Site Recovery.

In Enhanced mode the management server consisting of all the three process server, master target server, configuration server runs on Windows Server 2012 R2 will be deployed on-premises. The Azure site contains the VPN for connection and a storage for the replication data.

A. Components of Azure Site Recovery

- An on-premises management server: The management server runs Site Recovery components:
- Process server: Acts as a replication gateway. It receives data from protected source machines, optimizes it with caching, compression, and encryption, and sends replication data to Azure storage.
- Configuration server: Coordinates communication and manages data replication and recovery processes.
- Master target server: Handles replication data during failback from Azure.
- The Mobility service: This component is deployed on each machine (VMware VM or physical server) that you want to replicate to Azure.
- Azure account: A Microsoft Azure account.
- Azure storage: An Azure storage account to store replicated data. Replicated data is stored in Azure storage and Azure
- Azure network: An Azure virtual network with VPN connection that Azure VMs will connect to when failover occurs.

B. Functionality of the Components

The Process server talks with the mobility service which will be installed on all VMware VM's, the process server will take the data and make encryption, compression and caching then send the data to the Master target server. Also this server is responsible of the automatic discovery of VMware virtual machines.

The configuration Server is the brains, it co-ordinates communication between all components both on-premises and in Azure. Each Configuration Server can support up to 100 source virtual machines.

The on-premises ESXi (5.5) is the primary component on which the VMware VMs are running and which is to be protected. This is being managed by Vcenter server (5.5) which has to be added in Azure in order to discover the VMs.

The Master Target server receives incoming replication traffic from the Process Server, Each protected VM is added as a VHD using 'blob' storage.

The Azure account subscription enables to create a recovery vault which contains options to protect the virtual machines. An Azure storage has to be created which stores the VMs data, which is blob storage. In Azure account it is also mandatory to create a virtual network with VPN connection or express route connection, in this case it is a point to site VPN. This network ensures the security and helps in failover and failback of the data.

VII. Implementation

The implementation of Azure Site Recovery using Enhanced mode. This mode eliminates the virtual machines running on Azure. The following are the steps depicting the same:

Step 1: Install an ESXi server 5.5 on-premises

Step 2: Install a Vcenter server 5.5 to manage the ESXi

Step 3: Configuration of management server

Step 4: Creating a vault in Azure

Step 5: Create a Azure Virtual Network

Step 6: Create an Azure storage account with a friendly name

Step 7: Download the unified setup from the vault along with the registration key

Step8: Run the unified setup on the management server

Step 9: Once the CS, MT and PS server are installed register the server with the Azure recovery service vault

Step 10: Add the Vcenter server credentials in Azure vault

Step 11: Create a protection group in the vault

Step 12: Add the VMs that needs to be protected to the protection group and Enable the protection

Step 13: Create a recovery plan and add the protected VMs to the plan

Step 14: Run a test failover

Step 15: Go to the virtual machines tab in the portal which will contain the tested failover VM

Step 16: Any changes in the Azure VM can be re-protected and failed back to the on-premises server.

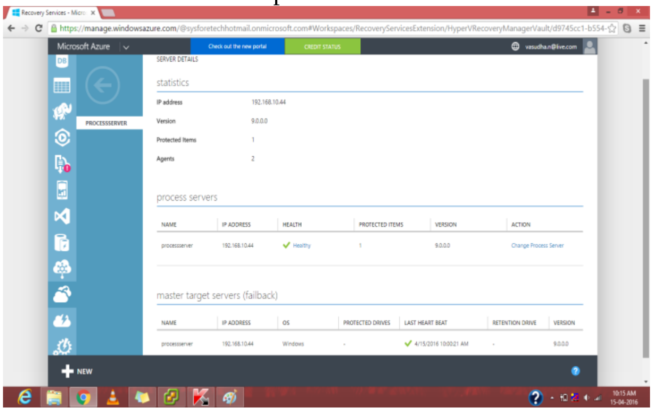
Recovery Services - Micro
https://manage.windowsazure.com/@sysforetechhotmail.onmicrosoft.com#Workspaces/RecoveryServicesExtension/HyperVRecoveryManagerVault/d9745cc1-b554-
Microsoft Azure
Check out the new portal
CREDIT STATUS
vasudha.n@live.com
DB
PROCESSSERVER
NEW
10:15 AM
15-04-2016

SERVER DETAILS

statistics

IP address	192.168.10.44
Version	9.0.0.0
Protected Items	1
Agents	2

process servers

NAME	IP ADDRESS	HEALTH	PROTECTED ITEMS	VERSION	ACTION
processserver	192.168.10.44	✓ Healthy	1	9.0.0.0	Change Process Server

master target servers (failback)

NAME	IP ADDRESS	OS	PROTECTED DRIVES	LAST HEART BEAT	RETENTION DRIVE	VERSION
processserver	192.168.10.44	Windows	-	✓ 4/15/2016 10:00:21 AM	-	9.0.0.0

Fig. 4: Server details on the recovery vault

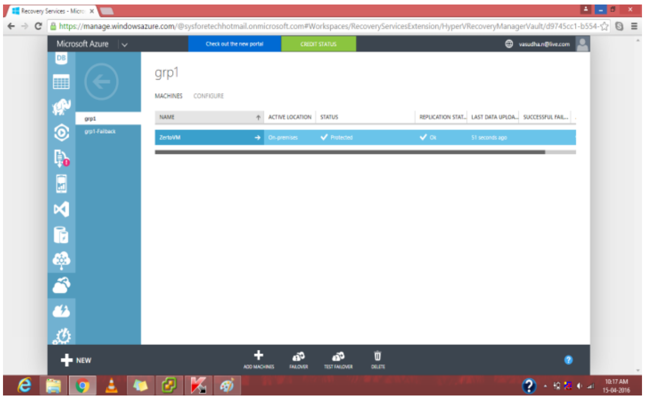
Recovery Services - Micro
https://manage.windowsazure.com/@sysforetechhotmail.onmicrosoft.com#Workspaces/RecoveryServicesExtension/HyperVRecoveryManagerVault/d9745cc1-b554-
Microsoft Azure
Check out the new portal
CREDIT STATUS
vasudha.n@live.com
grp1
grp1-Failback
NEW
ADD MACHINES
FAILOVER
TEST FAILOVER
DELETE
10:17 AM
15-04-2016

grp1

MACHINES CONFIGURE

NAME	ACTIVE LOCATION	STATUS	REPLICATION STAT...	LAST DATA UPLOA...	SUCCESSFUL FAIL...
ZertoVM	On-premises	✓ Protected	✓ Ok	51 seconds ago	

Fig. 5: Protecting VM in Azure

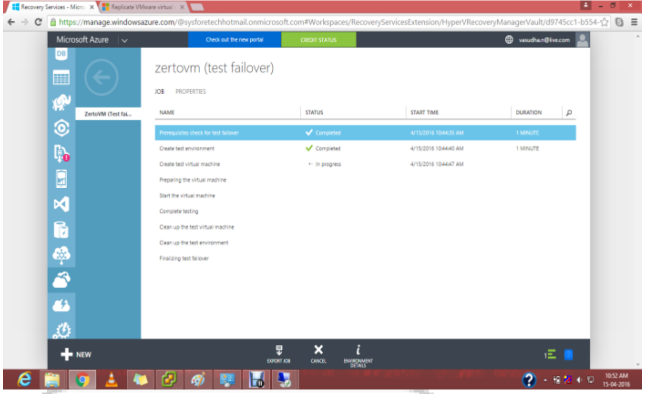
Recovery Services - Micro
Replicate VMware virtual
https://manage.windowsazure.com/@sysforetechhotmail.onmicrosoft.com#Workspaces/RecoveryServicesExtension/HyperVRecoveryManagerVault/d9745cc1-b554-
Microsoft Azure
Check out the new portal
CREDIT STATUS
vasudha.n@live.com
ZertoVM (Test fai...
NEW
EXPORT JOB
CANCEL
ENVIRONMENT DETAILS
10:52 AM
15-04-2016

zertovm (test failover)

JOB PROPERTIES

NAME	STATUS	START TIME	DURATION
Prerequisites check for test failover	✓ Completed	4/15/2016 10:44:35 AM	1 MINUTE
Create test environment	✓ Completed	4/15/2016 10:44:40 AM	1 MINUTE
Create test virtual machine	In progress	4/15/2016 10:44:47 AM	
Preparing the virtual machine			
Start the virtual machine			
Complete testing			
Clean up the test virtual machine			
Clean up the test environment			
Finalizing test failover			

Fig. 6: Testing failover

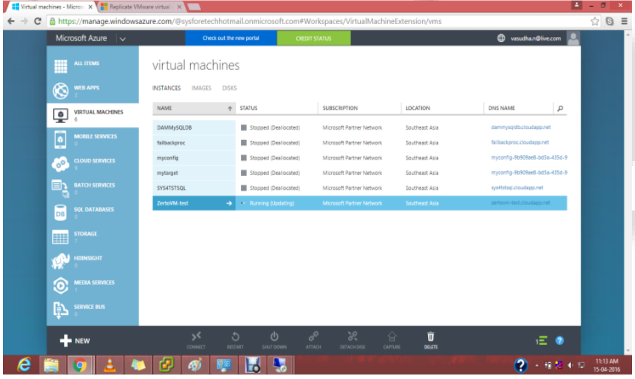

virtual machines

INSTANCES IMAGES DISKS

NAME	STATUS	SUBSCRIPTION	LOCATION	DNS NAME
DAMMySQLDB	Stopped (Deallocated)	Microsoft Partner Network	Southeast Asia	dammysqldb.cloudapp.net
failbackproc	Stopped (Deallocated)	Microsoft Partner Network	Southeast Asia	failbackproc.cloudapp.net
myconfig	Stopped (Deallocated)	Microsoft Partner Network	Southeast Asia	myconfig-9b909ae8-bd5a-435d-9
mytarget	Stopped (Deallocated)	Microsoft Partner Network	Southeast Asia	myconfig-9b909ae8-bd5a-435d-9
SYS4TSTSQL	Stopped (Deallocated)	Microsoft Partner Network	Southeast Asia	sys4tstsql.cloudapp.net
ZertoVM-test	Running (Updating)	Microsoft Partner Network	Southeast Asia	zertovm-test.cloudapp.net

Fig. 7: On-premises VM running on Azure

VIII. CONCLUSION

VMware VM on-premises were protected successfully using Microsoft Azure Site Recovery. Test fail over was successfully executed. The VMs were able to run and continue to function without any glitch on Azure. It is easy to configure replication to Azure in a few simple steps without incurring the cost and complexity that traditional replication solutions entail.

REFERENCES

[1] Yu Gu, Dongsheng Wang , and Chuanyi Liu, DR-Cloud: Multi-Cloud Based Disaster Recovery Service, ISSN 007 – 02141 02/10 llp p 1 3-2 3 Volume 19, Number 1, February 2014

[2] Dinesh Agarwal; Dept. of Comput. Sci., Georgia State Univ., Atlanta, GA, USA; Sushil K. Prasad, AzureBench: Benchmarking the Storage Services of the Azure Cloud Platform, published in: Parallel and Distributed Processing Symposium Workshops & PhD Forum (IPDPSW), 2012 IEEE 26th International.

[3] http://www.infoworld.com/article/2614229/virtualization/virtualization-showdown--microsoft-hyper-v-2012-vs--vmware-vsphere-5-1.html

[4] Bajpai, P. Rana, and S. Maitrey, Remote mirroring: A disaster recovery technique in cloud computing, International Journal of Advance Research in Science and Engineering, vol. 2, no. 8, 2013.

[5] Zia Saquib; Computer Networks and Internet Engineering, Centre for Development of Advanced Computing, India ; Veena Tyagi ; Shreya Bokare ; Shivraj Dongawe, Published in: Advanced Computing Technologies (ICACT), 2013 15th International Conference on 2013

[6] https://azure.microsoft.com/en-gb/services/site-recovery/

[7] J. Zhang and N. Zhang, cloud computing-based data storage and disaster recovery, in IEEE International Conference on Future Computer Science and Education (ICFCSE), 2011, pp. 629-632.

[8] Manish Pokharel, Seulki Lee, Jong Sou Park, "Disaster Recovery for Systems Architecture Using Cloud Computing,", IEEE/IPSJ Int. Symp. Applications and the Internet, 2010, pp. 304-307.

[9] Pete Manca, "Future of disaster recovery as aservice" .http: //www. drj.com/articles/ online-exclusive/thefuture-of-disaster-recovery-as-a-service.html, September 2012.

[10] The Vmware website. [Online] Available: http://www.vmware.com/products/high-availability/.

[11] Technical White paper on “VMware vCenter Site Recovery Manager ™ 5.0 Evaluation Guide”

[12] https://eugene.kaspersky.com/2014/05/13/three-ways-to-protect-virtual-machines-kaspersky-security-for-virtualization/

[13] Vic Winkler, "Cloud Computing: Virtual Cloud Security Concerns". Technet Magazine, Microsoft. http://technet.microsoft.com/enus/magazine/hh641415.aspx. Retrieved 12 February 2012.

[14] http://www.ijarcsse.com/docs/papers/Volume_3/11_November2013/V3I11-0211.pdf

[15] http://www.emc.com/collateral/white-paper/h14528-azure-site-recovery-san-replication-with-vmax3-wp.pdf

[16] T. Wood, E Cecchet, K. K. Ramakrishnan, P. Shenoy, J. van der Merwe, and A. Venkataramani, “Disaster recovery as a cloud service: economic benefits & deployment challenges”, Proc. 2nd USENIX Conference on Hot topics in cloud computing (HotCloud'10), Berkeley, CA, USA, 2010, pp. 8-8.