CHAPTER 1

CCA-MCP DRaaS Application Note

This document is an application deployment guide for Azure Site Recovery (ASR) on the top of Cisco Cloud Architecture for Microsoft Cloud Platform (CCA-MCP). While this guide provides an overview on ASR deployment, it is highly recommended to utilize Microsoft professional services and Cisco Advanced Services for DRaaS technology deployment due to the technology complexity and variations on enterprise customer's needs and requirements.

The guide includes prerequisites for the covered scenarios and explains how to set up a Site Recovery vault, get the Azure Site Recovery Provider installed on source and target VMM servers, register the servers in the vault, configure protection settings for VMM clouds that will be applied to all protected virtual machines, and then enable protection for those virtual machines.

CCA MCP Overview

The CCA-MCP infrastructure is the foundation on which cloud services are offered. The base infrastructure consists of a set of data center devices that are setup, connected, and configured prior to adding tenant services.

Service Providers build data centers using physical components to implement compute, storage and data center networking to create a pool of resources used to offer services to tenants. Tenant services are offered using these physical resources, and provisioned and managed using automation software to enable consumption of these services. When tenants are on boarded, cloud containers are created from the pool of resources, to provide a slice of resources that include compute, storage and networking. This container is securely isolated from other tenants that are consuming similar services, thereby providing isolation for multi-tenant services.

Figure 1-1 Cisco Cloud Architecture for the Microsoft Cloud Platform

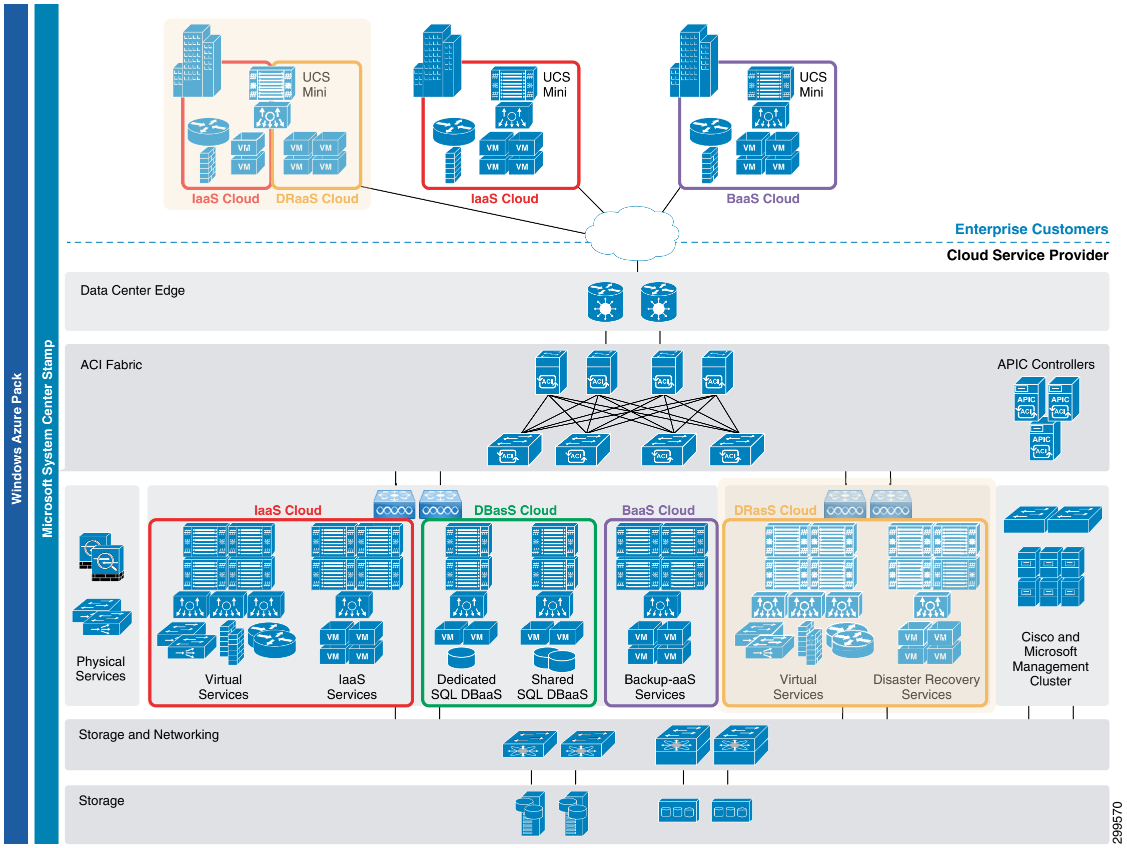

Enabled cloud services in the CCA-MCP Solution are the Infrastructure as a Service and Platform/Software as a service. Each of these services are described in a service configuration guide, and require the data center physical infrastructure to be built and the resource pools created and ready to onboard these services.

The architecture of this solution uses a layered approach enabling a modular design. This enables scalable solution deployment with expansion capability in modular units.

1. Data Center Network
2. Compute for Tenant workloads
3. Storage and SAN
4. Service Tiers and differentiated services
5. Cloud Management

Business Value

Most organizations consider Disaster Recovery (DR) a complicated process that is troublesome to manage and test. DR planning puts CIOs in a difficult position: most realize that ensuring application continuity in the case of a disaster is valuable to the business, but hesitate at the cost of the additional hardware and software needed, not to mention the time required to develop and test DR plans. Hence, more and more organizations started to subscribe to DRaaS service to reduce costs, complexity, and increase critical business applications availability.

Given the increased complexity of hybrid cloud solutions, service providers often find themselves playing the role of a systems integrator, wrapping together solutions from multiple vendors and providers to stand up a single, unified solution. That complexity also makes it more difficult to provision and monetize value-added services such as Disaster Recovery.

For the first time, Microsoft and Cisco have come together to provide cloud service providers with a joint-engineered, cloud-enabled platform solution that simplifies service delivery. Cisco Cloud Architecture for Microsoft Cloud Platform brings together Cisco's world-class hardware and Microsoft's enterprise-ready cloud software into a single, fully-integrated solution, making it easier for cloud providers to deliver cloud services such as DRaaS, while lowering deployment risk, simplifying operations, and improving TCO.

Cisco and Microsoft have invested the time, money and resources to help cloud providers offer their customers:

- **Full Featured Disaster Recovery—**With the integration of Windows Azure Pack into the Cisco Cloud Architecture for the Microsoft Cloud Platform, service providers can access Azure Site Recovery to support Disaster Recovery solutions for workloads like Exchange, SharePoint, SQL Server
- **Rapid Deployment—**Cisco Cloud Architecture for the Microsoft Cloud Platform facilitates deployment of DRaaS across your network infrastructures
- **Automated VM Protection and Replication—**The integration of Windows Azure Pack and Cisco Cloud Architecture enables service providers to offer policy-based replication and protection of VMs using a few simple steps
- **One-Click Orchestrated Recovery—**Azure Site Recovery's Disaster Recovery Plans provide a customizable framework to implement automated and custom recovery sequencing across public, private and hybrid cloud environments.

The below section provides benefits to the Service Providers and their customers

For Cloud Service Providers

- A business-critical service offering that promotes high revenues and margins
- Opportunities for additional services sales
- Expansion and monetization of existing cloud investments

For Service Provider Customers

- Reliable disaster recovery for both physical and virtual environments
- Reduced capital and operating systems
- Real continuous data protection
- Fast recovery times with reliable recovery points

Azure Site Recovery Overview

Cisco partnered with Microsoft to build a service provider cloud-based DR solution focused on mission critical workloads, Our primary priority: Make DR available to everyone, available everywhere, and easy to use. Arguably, that's three primary priorities – and we are pleased to include Microsoft Azure Site Recovery (ASR) with our Cisco Cloud Architecture for Microsoft Cloud Platform. ASR is a data protection and Disaster Recovery solution which can help enterprises protect important services by coordinating the automated replication and recovery of protected instances at a secondary location.

Azure Site Recovery on-premise to on-premise scenario provides a well-documented DRaaS solution for cloud providers. Disaster recovery is simplified and incorporated into the overall design, keeping both tenant and management systems highly available. ASR sends only management metadata to Azure to orchestrate DRaaS setup and recovery. Tenant and management data is transferred directly from the main to the DR site, and never goes to Azure. ASR plans orchestrate the recovery of resources at a designated site (Figure 1-2). ASR further simplifies the disaster recovery process by enabling testing of failovers and restorations of systems.

Figure 1-2 Microsoft Azure Site Recovery

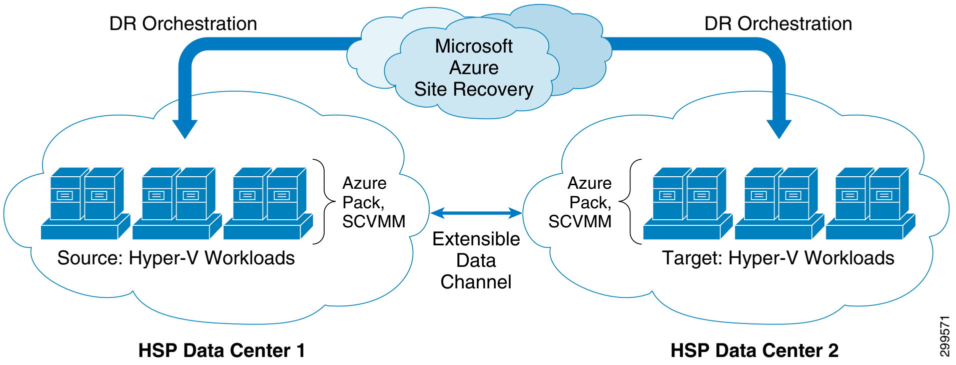

A recovery plan can be used for a planned failover, unplanned failover as well as for DR drills using test failover. A recovery plan continues to address the following needs for the user:

- Defining a group of virtual machines that failover together.
- Defining the dependencies between the virtual machines so that the application comes up accurately.
- Automating the recovery along with custom manual actions so that tasks other than the failover of the virtual machines can also be achieved

Microsoft's Azure Site Recovery (ASR) on-premise to on-premise scenario can be applied to both Incloud (IaaS onboarded workloads) and Remote (private cloud workloads) DRaaS use-cases.

The on-premise to on-premise deployment scenario will be covered in this solution guide - Virtual machines running in an environment using Hyper-V can be replicated between two data centers. Azure Site Recovery monitors the health of applications running in the primary site, stores the recovery plan, and executes it when needed. In the event of a site outage at the primary data center, VMs are recovered in an orchestrated fashion to help restore service quickly. This process can also be used for testing recovery, or temporarily transferring services.

Deployment Scenarios

The following are the two main use cases for DRaaS service delivery:

- In-Cloud DRaaS for IaaS Workloads with ASR and Azure Pack.
- Replicate IaaS on-boarded virtual machines from Primary on-premise SP site to Secondary on-premise SP site with Hyper-V replication. Users configure and enable protection settings in Azure Site Recovery vaults and VMM. Virtual machine data is replicated from a source Hyper-V host server to a target host server orchestrated through ASR.

Figure 1-3 In-Cloud DRaaS for IaaS Workloads with ASR and Azure Pack

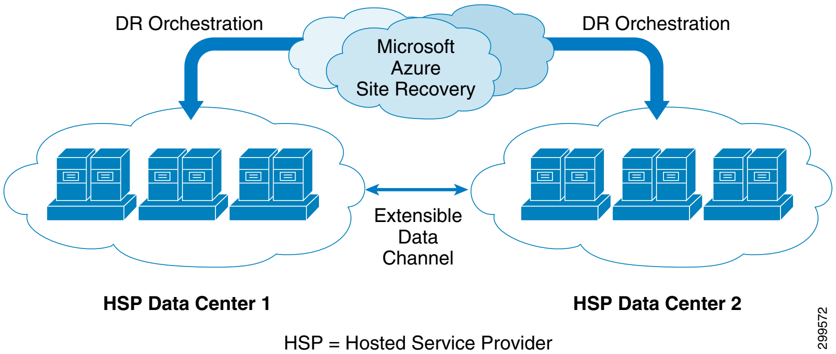

Fail Safe

After you've set up protection for virtual machines and physical servers they begin replicating to the secondary location. After replication is in place you can run failovers as the need arises. Site Recovery supports a number of types of failover.

Table 1-1 Site Recovery Failover Options

Failover	When to run	Details	Process
Test Failover	Run to validate your replication strategy or perform a disaster recovery drill	No data loss or downtime No impact on replication No impact on your production environment	Start the failover Specify how test machines will be connected to networks after failover Track progress on the Jobs tab. Test machines are created and start in the secondary location Azure - connect to the machine in the Azure portal Secondary site - access the machine on the same host and cloud Complete testing and automatically clean up test failover settings.

Table 1-1 Site Recovery Failover Options (continued)

Failover	When to run	Details	Process
Planned Failover	Run to meet compliance requirements Run for planned maintenance Run to fail over data to keep workloads running for known outages - such as an expected power failure or severe weather reports Run to failback after failover from primary to secondary	No data loss Downtime is incurred during the time it takes to shut down the virtual machine on the primary and bring it up on the secondary location. Virtual machines are impact as target machines becomes source machines after failover.	Start the failover Track progress on the Jobs tab. Source machines are shut down Replica machines start in the secondary location Azure - connect to the replica machine in the Azure portal Secondary site - access the machine on the same host and in the same cloud Commit the failover
Unplanned Failover	Run this type of failover reactive manner when a primary site becomes inaccessible because of an unexpected incident, such as a power outage or virus attack You can run an unplanned failover can be done even if primary site isn't available.	Data loss dependent on replication frequency settings Data will be up-to-date in accordance with the last time it was synchronized	Start the failover Track progress on the Jobs tab. Optionally try to shut down virtual machines and synchronize latest data Replica machines start in the secondary location Azure - connect to the replica machine in the Azure portal Secondary site access the machine on the same host and in the same cloud Commit the failover

Remote DRaaS for Enterprise Private Cloud Workloads with ASR and Azure Pack

Replicate enterprise private cloud virtual machines to SP cloud site with Hyper-V replication. Users configure and enable protection settings in Azure Site Recovery vaults and VMM. Virtual machine data is replicated from a source Hyper-V host server to a target host server orchestrated through ASR.

This Remote DRaaS ASR scenario has not gone through Cisco validation yet, we can offer the following design considerations:

- Both enterprise and SP side SCVMMs should have Internet access directly or via proxy to communicate with Microsoft ASR.
- Since SCVMMs are in different domains, trust between the hosts need to be certificate based. As the Active Directory (AD) of Enterprise and Service Provider AD will not have two way trust, hence Kerberos based authentication will not work.
- We need to call out the Management control. The ASR Vault is owned and managed by the Service Provider and all DR options are also controlled by the Service Provider. Enterprise can only select the VM to be protected.
- Connectivity needs to be setup between hosts on Enterprise and hosts on Service Provider DC for replication traffic to flow.
- There is no in built WAP integration for this remote DRaaS scenario, primarily for following reasons
  - Enterprises might not have WAP deployment.

- Even if the enterprise has the user identity on Enterprise WAP will be their Active Directory users which is not available on Service Provider side. To enable access of VMs on the Service Provider WAP, the Service Provider needs to run a script to assign user role on the protected VMs in his DCs.

Figure 1-4 Remote DRaaS for Customer Workloads

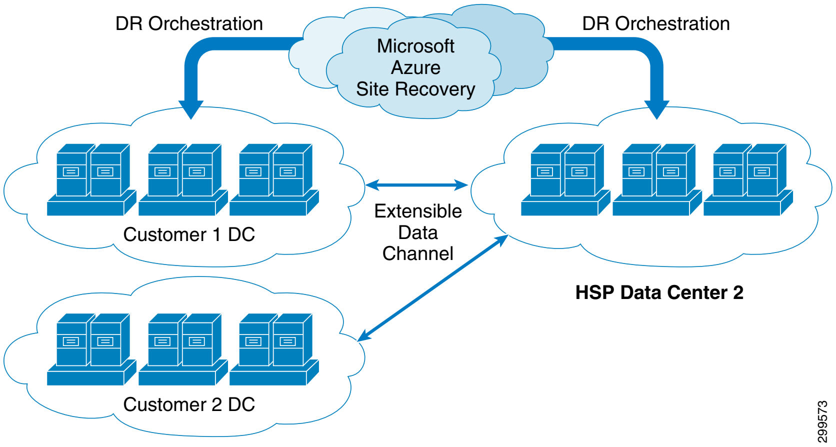

Both Remote DRaaS and InCloud DRaaS scenarios utilize on-premise to on-premise ASR deployment based on Hyper-V replica implementation.

Data Center Infrastructure

Note Network Container type depends on the Data Center architecture we choose for this validation and availability. There is no dependency on the Network container type for validating the ASR functionality.

Service Provider Data Center1

- Minimum two or more Hyper-V 2012 R2 Servers with 2TB storage capacity
- Minimum two or more Tenants: One Network Container per Tenant based on the DC Architecture used (VMDC x.x., CCA).
- At least two guest VM's to be protected, with multiple disks and preferably one VM running SQL.
- Management Software:
  - Systems Center
  - Azure Pack

Service Provider Data Center2

- Minimum two or more Hyper-V 2012 R2 servers with 2TB storage capacity
- Two-Four network containers needed for supporting the scenarios tested, containers need to match with the source containers with same number of VLAN's and similar network services

Network Connectivity

- Connectivity between the SCVMM servers from Source Data Center and Destination Data Center to Microsoft Azure over Internet.

- Communication between the SCVMM and ASR is over HTTPS.
- Dedicated replication network required between the source and destination Hyper-V servers.
- Source and the Destination Data Centers can be connected over S2S or L3VPN. Will L3VPN based connectivity for the PoC.

Prerequisites and Support

Microsoft Azure Requirements

An Azure account with Azure Site Recovery enabled.

VMM Prerequisites

- You'll need at least one VMM server.
- The VMM server should be running at least System Center 2012 SP1 with the latest cumulative updates.
- Any VMM server containing virtual machines you want to protect must be running the Azure Site Recovery Provider. This is installed during the Azure Site Recovery deployment.
- If you want to set up protection with a single VMM server you'll need at least two clouds configured on the server.
- If you want to deploy protection with two VMM servers each server must have at least one cloud configured on the primary VMM server you want to protect and one cloud configured on the secondary VMM server you want to use for protection and recovery
- All VMM clouds must have the Hyper-V Capacity profile set.
- The source cloud that you want to protect must contain the following:
- One or more VMM host groups.
- One or more Hyper-V host servers in each host group.
- One or more virtual machines on the host server.

Hyper-V Prerequisites

- The host and target Hyper-V servers must be running at least Windows Server 2012 with Hyper-V role and have the latest updates installed.
- If you're running Hyper-V in a cluster note that cluster broker isn't created automatically if you have a static IP address-based cluster. You'll need to configure the cluster broker manually.
- Any Hyper-V host server or cluster for which you want to manage protection must be included in a VMM cloud.
- **Capacity Planning**—For guidance use the Capacity Planner for Hyper-V Replica.

Figure 1-5shows the different communication channels and ports used by Azure Site Recovery for orchestration and replication.

Figure 1-5 On Premises to On Premises

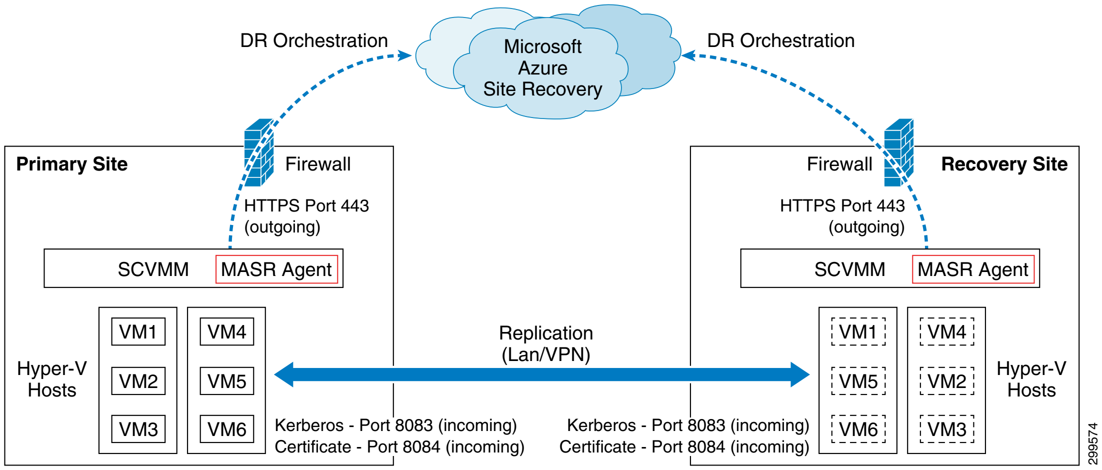
DR Orchestration
Microsoft Azure Site Recovery
DR Orchestration
Primary Site
Firewall
HTTPS Port 443 (outgoing)
SCVMM
MASR Agent
Hyper-V Hosts
VM1
VM2
VM3
VM4
VM5
VM6
Replication (Lan/VPN)
Kerberos - Port 8083 (incoming)
Certificate - Port 8084 (incoming)
Kerberos - Port 8083 (incoming)
Certificate - Port 8084 (incoming)
Firewall
Recovery Site
HTTPS Port 443 (outgoing)
SCVMM
MASR Agent
VM1
VM5
VM6
VM4
VM2
VM3
Hyper-V Hosts
299574

Prerequisites for On-Premises to On-Premises Protection

- Two data centers in the same Active Directory forest if you’re using Active Directory user subscription accounts.
- Each data center should have a VMM server running System Center 2012 R2.
- Each data center should be running Azure Pack with Update Rollup 4 or later.
- Each data center should have SMA installed and configured in the Azure pack deployment
- A Hyper-V host server located in each site, running Windows Server 2012, 2012 with SP1, or 2012 R2. The Hyper-V hosts must belong to VMM private clouds.
- An Azure account with Azure Site Recovery enabled. If you don’t have an account or the feature isn’t enabled, refer to Azure free trial and Azure Site Recovery Pricing details.
- VMM servers must have Internet access. In the browser on the server open http://manage.windowazure.com to confirm access.
- Check that Hyper-V host servers in both sites can connect. Ping using FQDNs to confirm. A primary Hyper-V host server should be able to connect to a secondary Hyper-V host server using port 8084.
- Each VMM cloud needs:
  - At least one VMM host group with at least one Hyper-V host server in each group.
  - One or more virtual machines on the host server.
- If you want to configure two-way protection you need at least two clouds on each VMM server.
- If you want to set up network mapping you’ll need to set up logical networks on the VMM servers. In the self-service Azure Pack portal tenants set up virtual networks based on those logical VMM networks. When a tenant configures a virtual network a matching VM network is created in the VMM console. You then select and map these VM networks in the Azure Site Recovery portal using the network mapping feature.

Network Mapping Requirements

- In order to deploy network mapping so that virtual machines are connected to a VM network after failover, the following is required:
- The virtual machines you want to protect on the source VMM server should be connected to a VM network. That network should be linked to a logical network that is associated with the cloud.

- The target cloud on the secondary VMM server that you use for recovery should have a corresponding VM network configured, and it in turn should be linked to a corresponding logical network that is associated with the target cloud.

Refer to detailed instructions on planning the environment.

Following are the high level steps required to enable protection of VMs, please following the inbuilt hyperlinks in each step for detailed procedure from Microsoft site:

1. Create a Vault—create a vault in the Azure Site Recovery portal.
2. Install and configure the Azure Site Recovery Provider—Install and configure the Azure Site Recovery Provider on the VMM server in each site. The Provider connects the server to the Azure Site Recovery portal.
3. Configure Cloud Settings—Configure protection settings for VMM clouds. The cloud that contains the virtual machines you want to protect is known as the primary cloud. The cloud that contains the Hyper-V host server to which the virtual machines will replicate is known as the secondary cloud. Each cloud can act as a primary cloud protecting a secondary cloud, or as a secondary cloud that's protected. A cloud can't be both primary and secondary.
4. Set Up the Runbooks—You configure and schedule a single master runbook to set up Azure Site Recovery protection. This master runbook in turn invokes a number of other runbooks.
5. Configure Plans—On the primary site you enable Azure Site Recovery protection on a public plan or add-on, and create a private plan with the same settings on the secondary site.
6. Tenant Steps—In order to set up virtual machine protection tenants will use the self-service Azure Pack portal to:
   - Subscribe to the plan or add-on—Tenants subscribe to a plan or add-on the primary datacenter that has virtual machine protection enabled.
   - Create a virtual machine—Tenants create a virtual machine or virtual machine role on the primary site, under the plan subscription.
   - Create VM networks—Tenants can create virtual networks on the primary site to specify how replica virtual machines will be connected to networks after failover. When a tenant creates a virtual network a VM network with the same settings is configured on the primary VMM server.
7. Set up Network Mapping—If the tenant has created virtual networks you can set up network mapping between VM networks on the primary and secondary VMM servers. Network mapping:
   - Ensures that virtual machines are connected to appropriate VM networks after failover. Replica virtual machines will connect to a secondary network that's mapped to the primary network.
   - Optimally places replica virtual machines on Hyper-V host servers. Replica virtual machines will be placed on hosts that can access mapped VM networks.

   If you don't configure network mapping replicated virtual machines won't be connected to any VM networks after failover. Read about Network mapping.
8. Verify user Accounts—Before you can replicate virtual machines you'll need to verify that user credentials associated with the plan or add-on subscription are valid on the primary and secondary sites.
9. Detecting and Replicating Virtual Machines—The runbooks automatically detect plans or add-on subscriptions that have protection enabled. The runbook automatically enables protection for virtual machines in the subscriptions, and initiates the initial replication.
10. Run a Failover—After the initial replication finishes you can run a test, planned, or unplanned failover whenever you need to.

Validation Details

Protection

- Protection will be set up for the guest VM using Azure site recovery or SCVMM. This involves a source side discovery, target side discovery, target cloud selection, replica VM and network configuration.
- ASR selects a suitable host on the recovery site, which has necessary Storage, Memory and network connectivity for the VM to come up.

Recovery

- Create recovery plan, once the VM's have been replicated, and are fully protected, the user needs to create a Recovery Plan, which will help to coordinate the failover between sites.
- Recovery plans consist of a series of steps, executed in a particular order. Please refer to Microsoft documentation for all the options available to choose.
- Execute recovery plan by selecting the type of recovery: Test, Unplanned or Planned failover.
  - **Test failover**—Supports testing without impacting production
  - **Planned Failover**—Supports planned maintenance as well as DR drill with actual failover
  - **Unplanned Failover**—Supports unplanned failover needs like power failures,

The following steps are needed to support In-Cloud DRaaS use case for IaaS workloads. The same procedure described above has to be followed with some additional steps listed below:

DR Infrastructure Setup

- Service Provider login on Azure portal to create ASR vault
- Download the on-premises agents and install on the SCVMM machines
- Register the agents to ASR vault
- Create the Clouds on both SCVMM machines
- Pair the clouds of in the ASR Portal
- Downloads ASR runbooks from script center and deploy on the Azure Pack

Plan and DR Add-On Creation in WAP

- Service Provider create a DR offer with a plan with DR enabled (or publish as Add-on)
- Service Provider create a private plan on secondary WAP
- Create a schedule for ASR runbook to get triggered periodically

For detailed information and procedure on how to enable DR for IaaS Workloads with ASR and Azure Pack refer the Azure blog.

Service Provider FAQ

Is ASR secure? What data do you send to Azure?

Yes, ASR is secure. It neither intercepts your application data nor has any information about what is running inside your virtual machines. No Internet connectivity is needed, either from the Hyper-V hosts or the virtual machines.

Since replication is between your own enterprise and service provider sites, no application data is sent to Azure. Only metadata, such as VM or cloud names, that is needed to orchestrate failover, is sent to Azure. ASR does not have the ability to intercept your application data, and that data always remains on-premises.

ASR is ISO 27001:2005-certified, and is in the process of completing its HIPAA, DPA, and FedRAMP JAB assessments.

Compliance requirements require that even metadata from our on-premises environments remains within the same geographic region. Can ASR ensure that we meet that requirement?

Yes. When you create a Site Recovery vault in a region of your choice, we ensure that all metadata that we need to enable and orchestrate your disaster recovery setup remains within that region's geographic boundary.

What versions of Windows Server hosts and clusters are supported?

Windows Server 2012 and Windows Server 2012 R2 can be used when you choose Hyper-V Replica to enable replication and protection between Hyper-V Sites.

What versions of Hyper-V guest operating systems are supported?

The most current list of supported guest operating systems is available in the topic titled About Virtual Machines and Guest Operating Systems.

Is the identity of my tenant shared with Azure?

No. In fact, your tenants do not need access to the ASR portal. Only the service provider administrator performs actions on the ASR portal in Azure.

What if Azure is down? Can I trigger a failover from VMM?

While Azure is designed for service resilience, we like to plan for the worst. ASR is already engineered for failover to a secondary Azure datacenter in accordance with the Azure SLA. We also ensure that even when this happens, your metadata and ASR Vault remain within the same geographic region in which you chose to your vault.

What if my ISP for the primary datacenter also experiences an outage during a disaster? What if my ISP for the secondary datacenter also experiences an outage?

You can use the ASR portal to perform an unplanned failover with a single click. No connectivity from your primary datacenter is needed to perform the failover.

It's likely that applications that need to be running in your secondary datacenter after a failover will need some form of Internet connectivity. ASR assumes that even though the primary site and ISP may be impacted during a disaster, the secondary site's SCVMM server is still connected to ASR.

Does this solution works for dedicated or shared infrastructure model?

ASR supports both dedicated as well as shared infrastructure models.

Do Tenant identities get shared with Azure?

No, in fact Tenants never have to go to Azure management portal. Only Service Provider admin uses ASR UI in Azure portal.

Will my Tenants get bill from Azure?

No Service Providers will get ASR bill from Microsoft and they will generate Tenant specific bills.

Will Tenant's application data go to the public cloud?

For Service Provider as target site - application data never goes to Azure. It is always sent encrypted over network link between two data centers.

References

Performance and scaling testing:

- On-premises to on-premises
- Network infrastructure considerations for Site Recovery
- Site Recovery Components
- Best Practices for Site Recovery deployment
- Set up protection between on-premises VMM sites—Setup by Step.
- Cisco Cloud Architecture for the Microsoft Cloud Platform: Infrastructure Foundation Guide, Release 1.0.