Azure Site Recovery: How to configure it when you have VMWare infrastructure

What is BCDR?

A business continuity and disaster recovery (BCDR) strategy help keep your business up and running. During planned downtime and unexpected outages, BCDR keeps data safe and available, and ensures that apps continue running. In addition to platform BCDR features such as regional pairing, and high availability storage, Azure provides Recovery Services as an integral part of your BCDR solution. Recovery services include:

- Azure Backup backs up your on-premises and Azure VM data. You can back up a file and folders, specific workloads, or an entire VM
- **Azure Site Recovery** provides resilience and disaster recovery for apps and workloads running on on-premises machines, or Azure IaaS VMs. Site Recovery orchestrates replication and handles failover to Azure when outages occur. It also handles recovery from Azure to your primary site.

How does Site Recovery do disaster recovery?

1. After preparing Azure and your on-premises site, you set up and enable replication for your on-premises machines.
2. Site Recovery orchestrates initial replication of the machine, in accordance with your policy settings.
3. After the initial replication, Site Recovery replicates delta changes to Azure.
4. When everything's replicating as expected, you run a disaster recovery drill.
    - The drill helps ensure that failover will work as expected when a real need arises.
    - The drill performs a test failover without impacting your production environment.
5. If an outage occurs, you run a full failover to Azure. You can fail over a single machine, or you can create a recovery plan that fails over multiple machines at the same time.
6. On failover, Azure VMs are created from the VM data in Managed disks or storage accounts. Users can continue accessing apps and workloads from the Azure VM
7. When your on-premises site is available again, you fail back from Azure.
8. After you fail back and are working from your primary site once more, you start replicating on-premises VMs to Azure again

What do I need to set up in Azure before I start?

In Azure you need to prepare the following:

1. Verify that your Azure account has permissions to create VMs in Azure.
2. Create an Azure network that Azure VMs will join when they're created from storage accounts or managed disks after failover.
3. Set up an Azure Recovery Services vault for Site Recovery. The vault resides in the Azure portal, and is used to deploy, configure, orchestrate, monitor, and troubleshoot your Site Recovery deployment.

What do I need to set-up on-premises before I start?

On-premises here's what you need to do:

1. You need to set up a couple of accounts:
   - If you're replicating VMware VMs, an account is needed for Site Recovery to access vCenter Server or vSphere ESXi hosts to automatically discover VMs.
   - An account is needed to install the Site Recovery Mobility service agent on each physical machine or VM you want to replicate.
2. You need to check the compatibility of your VMware infrastructure if you didn't previously do that.
3. Ensure that you can connect to Azure VMs after a failover. You set up RDP on on-premises Windows machines, or SSH on Linux machines.

Enable Replication

1. Go to **Replicate application** > **Source**. After you enable replication for the first time, select **+Replicate** in the vault to enable replication for additional virtual machines.
2. In the **Source** page > **Source**, select the configuration server.
3. For **Machine type**, select **Virtual Machines** or **Physical Machines**.
4. In **vCenter/vSphere Hypervisor**, select the vCenter server that manages the vSphere host, or select the host. This setting isn't relevant if you're replicating physical computers.
5. Select the process server. If there are no additional process servers created, the inbuilt process server of configuration server will be available in the dropdown menu. The health status of each process server is indicated as per recommended limits and other parameters. Choose a healthy process server. A critical process server can't be chosen.

6. For Target, select the subscription and resource group where you want to create the failed over virtual machines. Choose the deployment model that you want to use in Azure for the failed over VMs.
7. Select the Azure network and subnet that the Azure VMs will connect to after failover. The network must be in the same region as the Site Recovery service vault.
8. For **Virtual machines** > **Select virtual machines**, select each virtual machine that you want to replicate. You can only select virtual machines for which replication can be enabled. Then select **OK**. If you can't see or select any particular virtual machine, see Source machine isn't listed in the Azure portal to resolve the issue.
9. For Properties > Configure properties, select the account that the process server uses to automatically install the Site Recovery Mobility service on the VM. Also, choose the type of target managed disk to use for replication based on your data churn patterns.
10. By default, all the disks of a source VM are replicated. To exclude disks from replication, clear the **Include** check box for any disks that you don't want to replicate. Then select **OK**. You can set additional properties later. Learn more about excluding disks.
11. From Replication settings > Configure replication settings, verify that the correct replication policy is selected. You can modify replication policy settings at Settings > Replication policies > *policy name* > Edit Settings. Changes applied to a policy also apply to replicating and new virtual machines.
12. If you want to gather virtual machines into a replication group, enable **Multi-VM consistency**. Specify a name for the group, and then select **OK**.
13. Select Enable Replication. You can track the progress of the Enable Protection job at Settings > Jobs > Site Recovery Jobs. After the Finalize Protection job runs, the virtual machine is ready for failover.