TECHNICAL PAPER
November 2024

Protecting and Recovering Mission-Critical Applications in a VMware Hybrid Cloud with VMware Live Site Recovery

vmware® by Broadcom

Contents

About this document

This guide provides a comprehensive documentation of the considerations and configuration steps required for using VMware® Live Site Recovery™ (VLSR) to protect and recover a reference multi-tiered set of business-critical applications from a source VMware datacenter (on-premises or cloud-based) to a supported target VMware datacenter (on-premises or cloud-based), with the least cost (time, financial, and administrative intervention) possible.

Audience

Technical architects, administrators, or operators can use this guide as a foundation to build similar solutions for their own enterprise infrastructure.

What we don't include in this document

This guide demonstrates how to use VMware Live Site Recovery to protect virtualized business critical applications on VMware vSphere® installed on-premises or in a hybrid cloud. Because we assume you're familiar with the general concepts of business continuity and recovery, we don't define or explain such concepts in detail. We also don't discuss or explain the setup, configuration, operation, or administration of VMware Live Site Recovery, virtualization, a VMware hybrid cloud or the applications and services hosted on or provided by the protected workloads.

We assume you've configured the infrastructure to perform these tasks:

- Installation, setup, configuration and/or administration of VMware vSphere infrastructure
- Installation, setup, configuration and/or administration of specific VMware vSphere-based Cloud infrastructure
- Installation, setup, configuration and/or administration of VMware Live Site Recovery
- Virtualizing Active Directory Domain Services on VMware vSphere
- Architecting Microsoft SQL Server on VMware vSphere
- Installation, setup, configuration and/or administration of Microsoft Active Directory Domain Services or Domain Controllers
- Installation, setup, configuration and/or administration of Microsoft SQL Server, Windows Failover Cluster or Always On
- VMware vSphere Client

This document doesn't include detailed descriptions of these topics.

Before you begin

Requirements

You should complete the following tasks before continuing with the instructions in this guide:

- Set up network connectivity for:
  - The **protected site:** This is the source infrastructure.
  - The **recovery site:** A new or existing VMware Cloud Foundation® environment or any of the publicly available brands of the VMware cloud infrastructure options, such as Azure VMware Solution (AVS), or Google Cloud VMware Engine (GCVE).

> **Note:** The VMware cloud brand and version dictates the type of network connectivity type required for VMware Live Site Recovery. Consult the cloud provider's guides for more information.

- Install VMware Live Site Recovery on both the protected and recovery sites.
- On each of the sites, install VMware vSphere Replication appliances in the same VMware vCenter® where the VMware Live Site Recovery instance is registered.
- Configure the VM IP addresses, DNS server IP addresses, network segment, and datastore required to complete the protection and recovery plans.
- Make sure all the VMs that will be protected and recovered have an up-to-date version of VMware Tools installed

> **Note:** This is a standard recommendation, but it's especially relevant if the VMs will be reconfigured or customized as part of the recovery process.

Terms to know

Cold site recovery

Distributing servers and services over multiple datacenters is a common business continuity and disaster recovery (BCDR) strategy, but a cloud-based solution can reduce the associated costs of maintaining a dedicated disaster recovery (DR) site like staffing, cooling, heat, and duplicate hardware. You can further reduce these associated costs by minimizing the actual utilization of the cloud-located resources until it is necessary to do so: when an actual disaster event has happened, or during a simulation, testing, or validation exercise. This type of "use only when needed" utilization is commonly described as having a cold site for BCDR. In this configuration, live (hot) devices, servers, and services aren't hosted in the target DR site. This saves you money and resources in your BCDR systems. We demonstrate how VMware Live Site Recovery achieves this cost-saving objective while providing a simplified, flexible, automated, and repeatable BCDR solution for your enterprise.

Application high availability

The ability of an application to function and deliver services even when one or more of its components fail is the focus of **application high availability (HA)**. The resilience of this application—whether native or through the use of

third-party solutions or add-ons—determines its capacity to withstand and recover from failures. The combination of Microsoft SQL Server Always On and Windows Server Failover Clustering (WSFC) provides application-level resilience in the scenario described in this guide. Because of these features, Microsoft SQL Server services can remain available after a brief interruption even when the original server providing the service has become unavailable for any reason. WSFC restores its resources on a functioning node in the event of the original server failure, typically without the need for administrative intervention.

Disaster recovery event

A **disaster recovery (DR) event** is a failure that affects more than one server or part of the system. A DR event is a collection of multiple HA events that can't be easily fixed by the resilience of an application, component, or service. Because it's not usually transient in nature, the effects of a DR event are more impactful, disruptive, and destructive. Recovering from a DR event is more difficult, more expensive, and slower than recovering from a HA event. This is because multiple layers of the infrastructure are affected. In turn, this means that planning and getting ready for a DR event costs more.

Architecture

We configured a VCF-based protected site on-premises with several VMs. These VMs are part of a multi-tiered mission-critical workloads. Some VMs are Domain Controllers providing Active Directory Domain Services (ADDS) for the infrastructure, and the rest are VMs running Microsoft SQL Server instances (MSSQL). We are replicating these VMs to a remote VCF-based infrastructure, configured as our cold DR site. The protected site could be on-premises, or in a VMware-supported cloud environment. The same is true for the recovery site.

The following figure shows our setup. For the exercises in this guide, your setup will be similar.

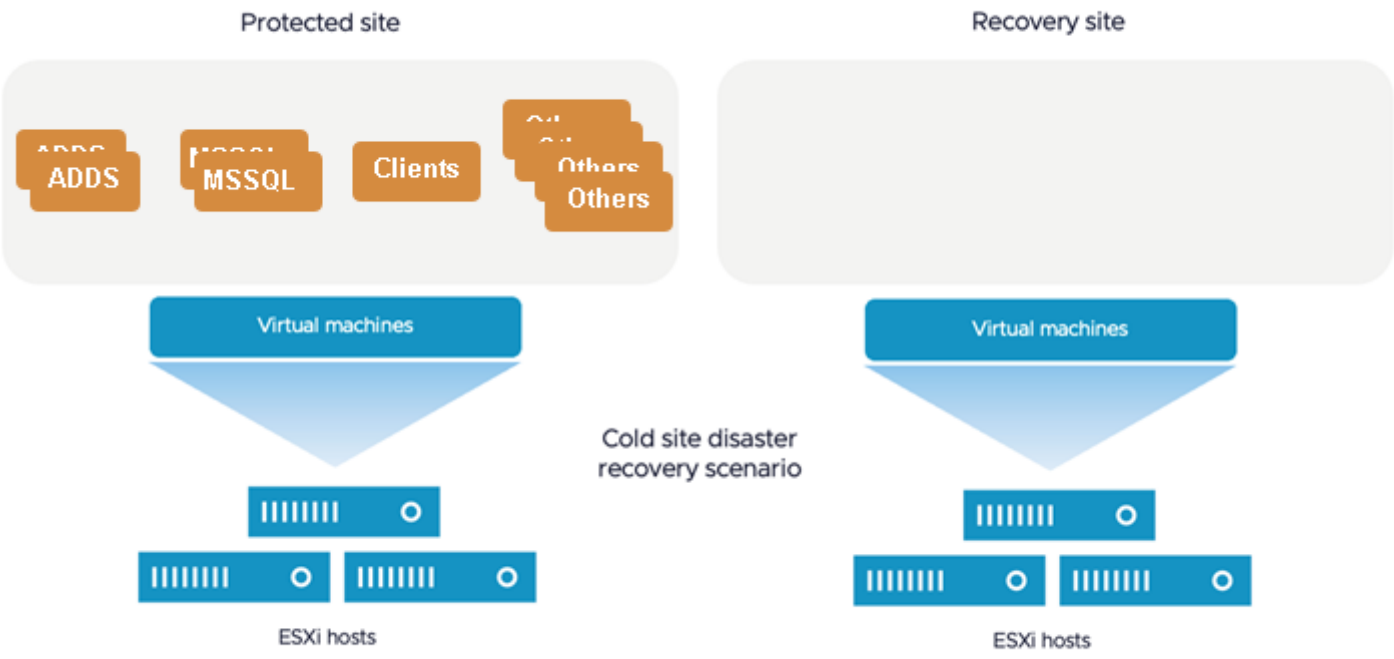

Applications

We chose three applications that show the capabilities of automation, orchestration, and the recovery tasks possible with VMware Live Site Recovery:

- Windows Active Directory Domain Controllers
- Microsoft SQL Server
- Windows client

Windows Active Directory Domain Controllers

Most BCDR plans include considerations and provisions for Domain Controllers because most applications depend on the services they provide, so they are common in most enterprise network infrastructure. Recovering modern versions of Windows Domain Controllers (anything newer than Windows Server 2008 R2) in the event of a disaster is somewhat difficult and can be complicated in our DR scenario. This is partly due to the security features Microsoft introduced into virtualized Domain Controllers beginning in Windows Server 2012. This guide addresses this issue and shows how VMware Live Site Recovery helps minimize these challenges.

Microsoft SQL Server

Because of its integration with so many front-end applications, services, and solutions, Microsoft SQL Server is one of the most prevalent business-critical applications found in any Microsoft-based corporate IT infrastructure. SQL Server has native, built-in resilience to maximize its availability and minimize the possibility of service disruption in the event of an outage. Combining Windows Server Failover Cluster (WSFC) with the SQL Server Always On feature is a high availability option that ensures faster service availability, particularly for databases, in the event of a node failure. Even then, this resilience is more useful and intended for high availability (which protects against component or service failures) rather than for disaster recovery events.

Windows client

We chose an ordinary Windows client from which we tested connectivity and access to the servers and services we recovered in our failure scenarios.

How to set up the business continuity/disaster recovery environment and workflow

Let's jump into the VMware Live Site Recovery configuration. For most exercises in this section, you'll need access to the vSphere Client. You'll also confirm the functionality of the recovered workloads inside these applications: Windows OS, Active Directory domain services (ADDS), and SQL Server.

Make sure you read the **Before you begin > Requirements** section above. This includes installing and configuring VMware Live Site Recovery.

Pair the recovery and protected sites

Here, we'll connect the vCenter and VMware Live Site Recovery instances on each site to one another. This is called pairing the sites.

1. In the vSphere Client on the recovery site's vCenter, select **Site Recovery**. The **Summary** page appears, as shown below.

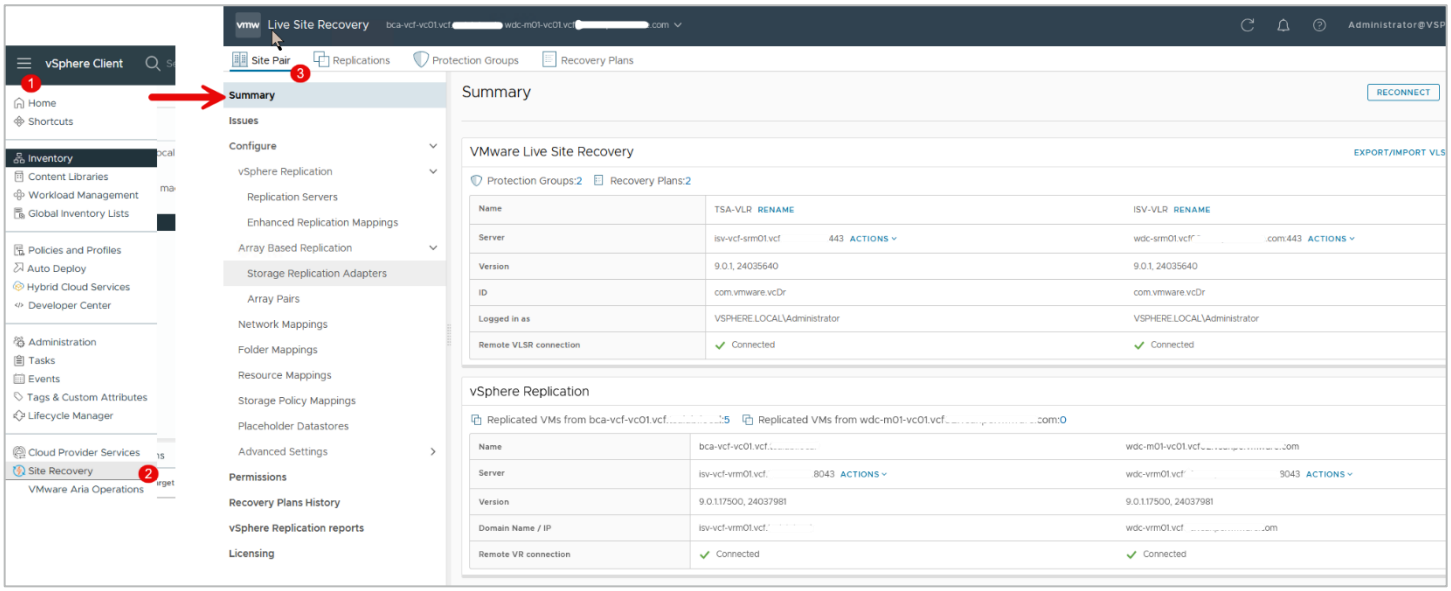

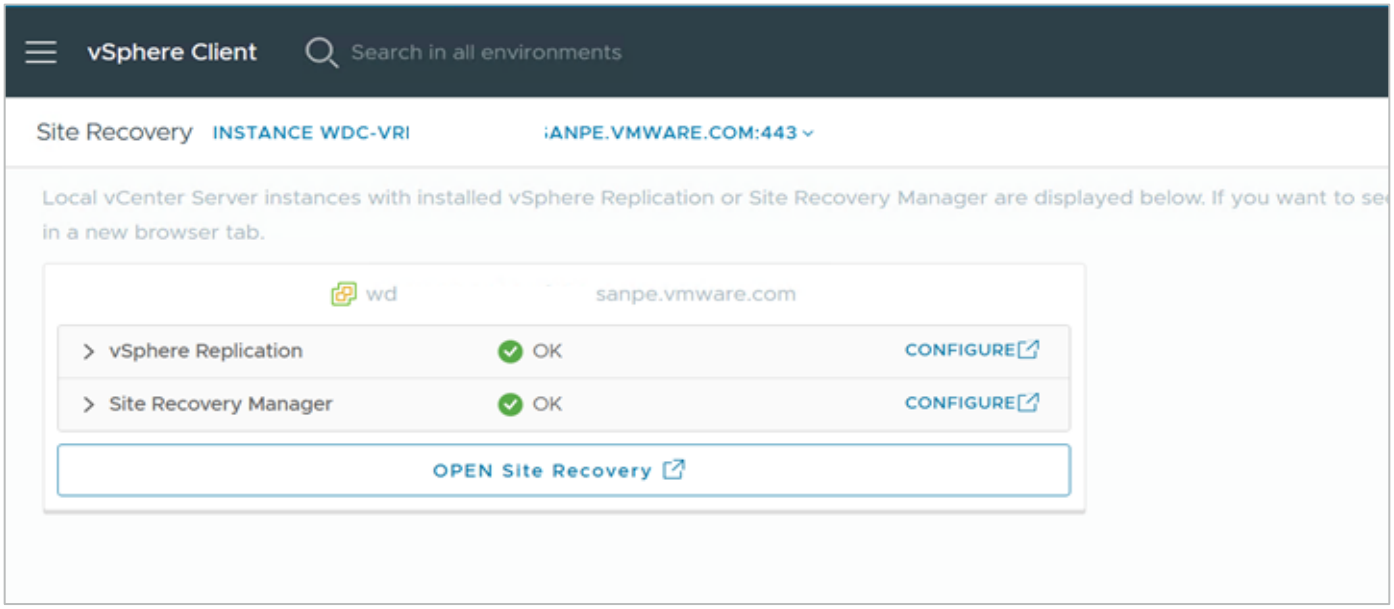

2. Click **OPEN Site Recovery** to access the protected site.

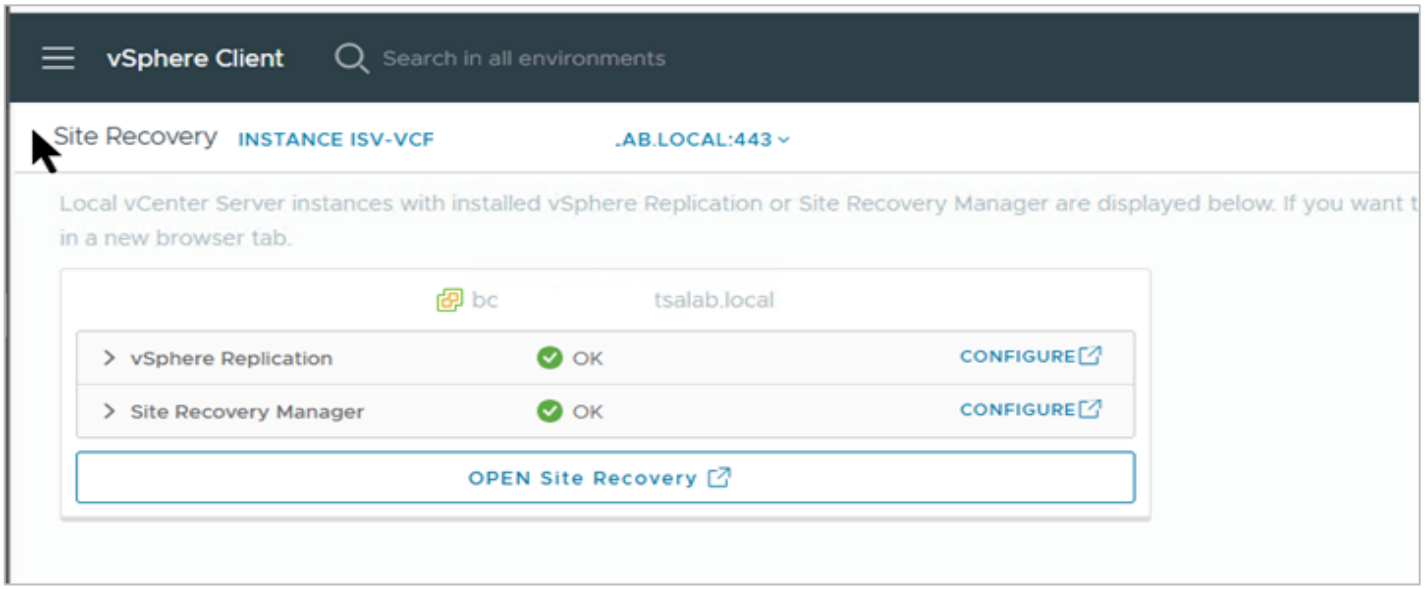

3. Click **NEW SITE PAIR**.

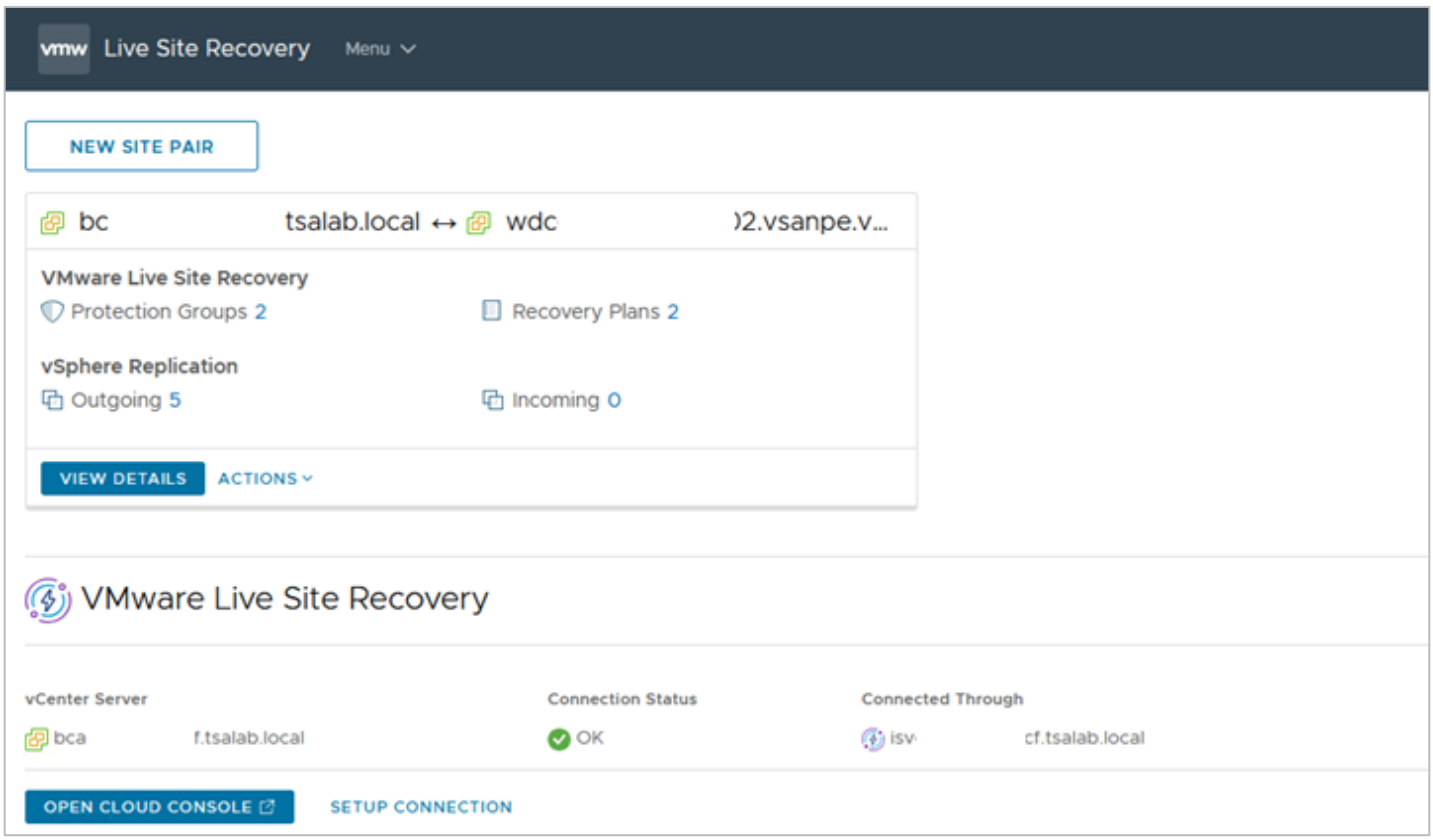

4. On the Pair type screen, select the applicable SSO domain and click **NEXT**.

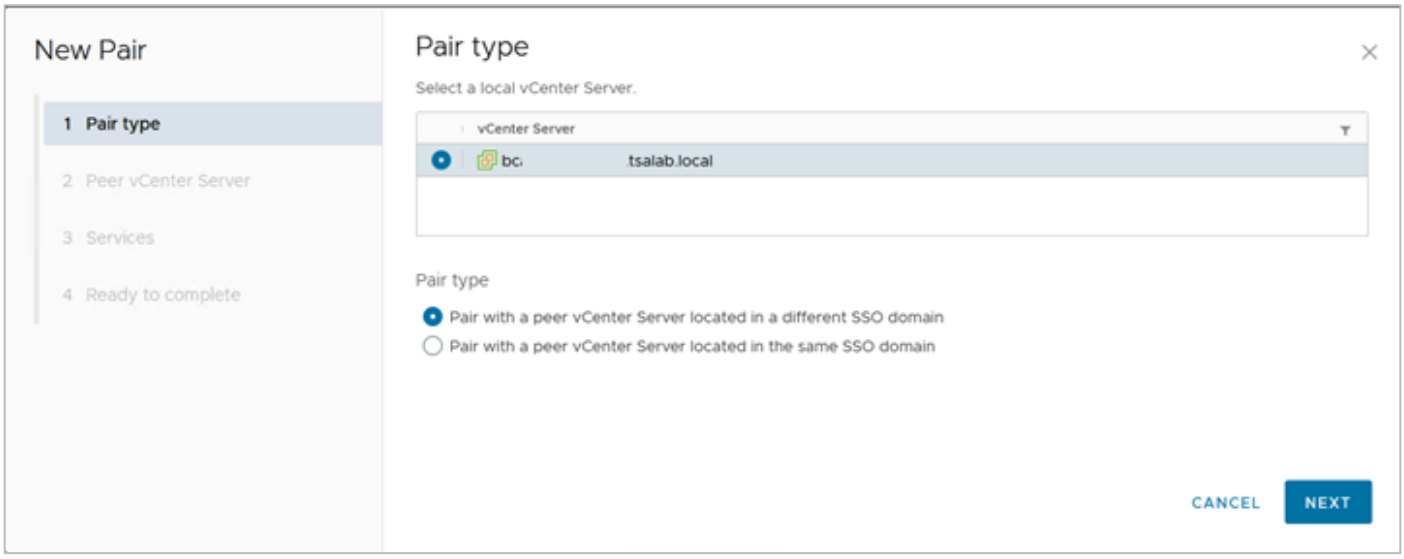

You'll be prompted for the vCenter credentials.

Because we're doing this from the recovery site's vCenter instance, the credentials we provide here will be for the **protected** site's vCenter instance.

5. Provide the remote vCenter's information and credentials, click **FIND VCENTER SERVER INSTANCES,** and click **NEXT**.

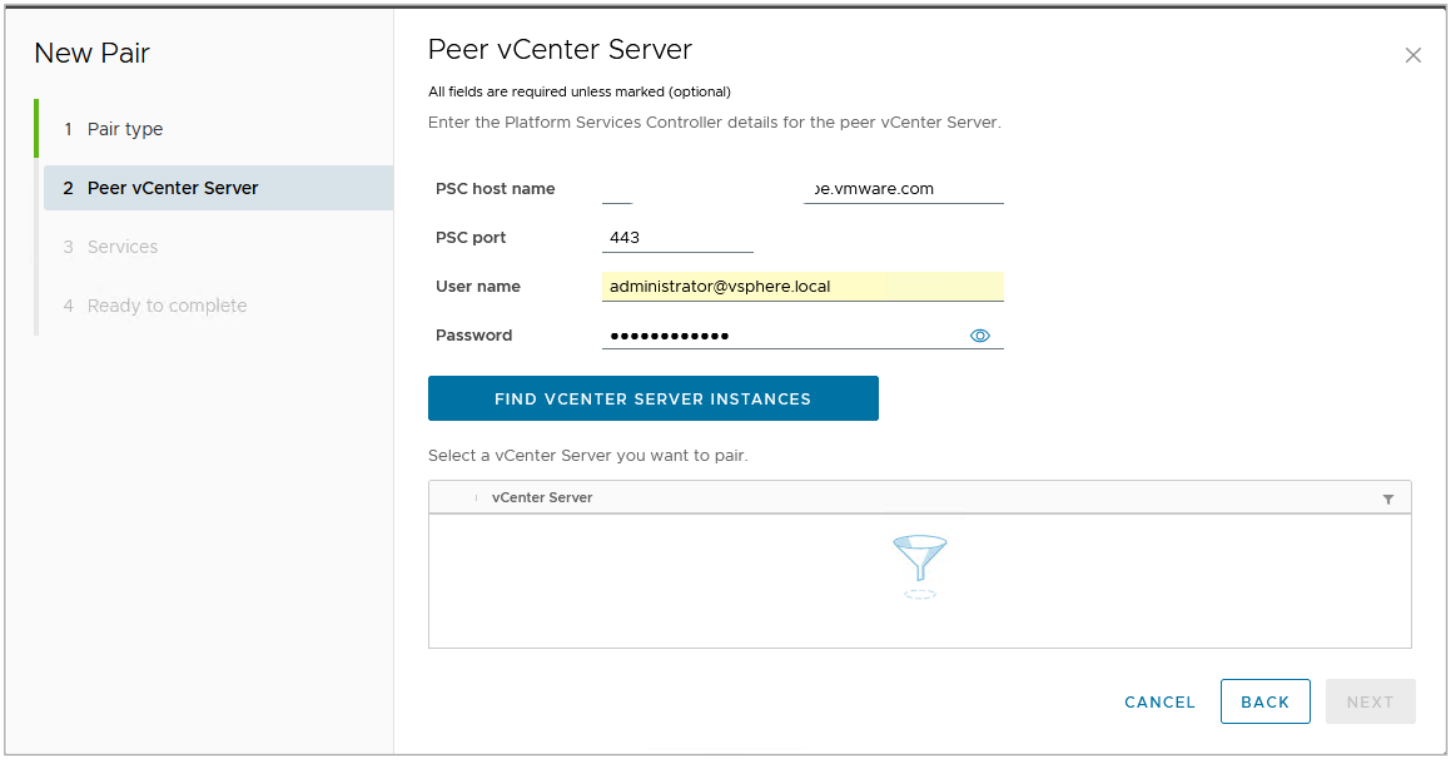

Note: If you're using default self-signed certificates in your environment, you'll need to click **Connect** to ignore the vCenter's self-signed certificate security warning to proceed.

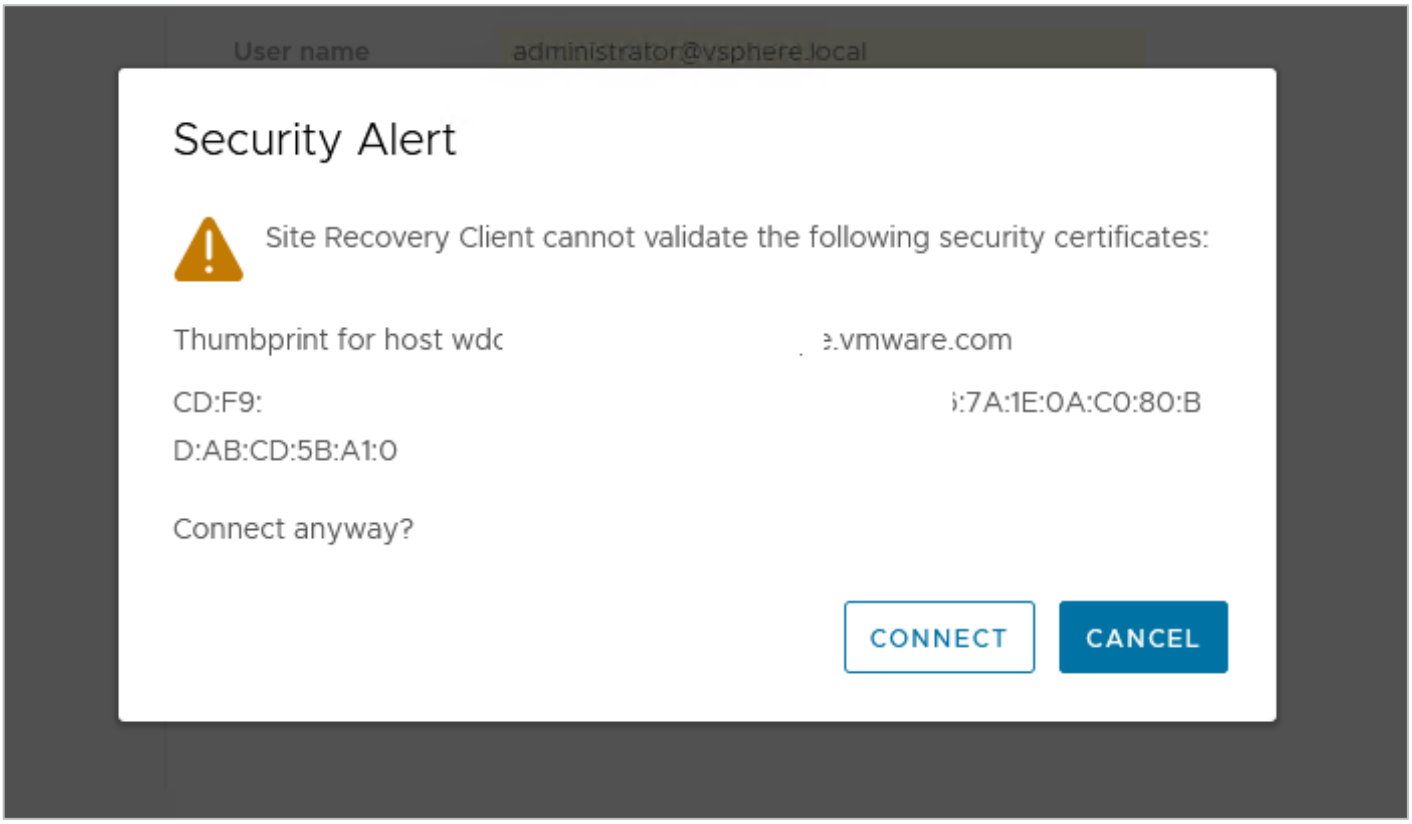

We used VMware vSAN for the storage subsystem in our environments. vSAN is the default storage option for all VMware vSphere cloud infrastructure. In this configuration, we see that the VMware Live Site Recovery and the VMware vSphere Replication appliance are both registered on our vCenter.

6. Click **Next** to continue.

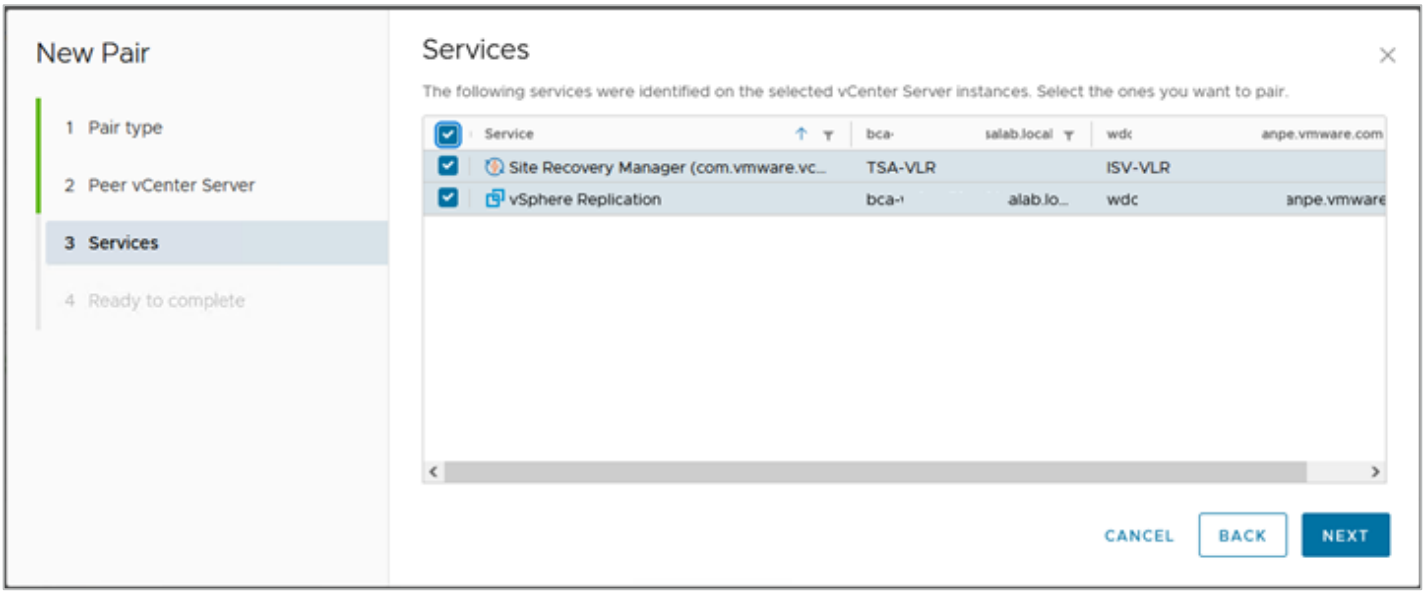

Note: You'll need to click **Connect** to ignore the vCenter's self-signed certificate security warnings if you're using them before you can proceed.

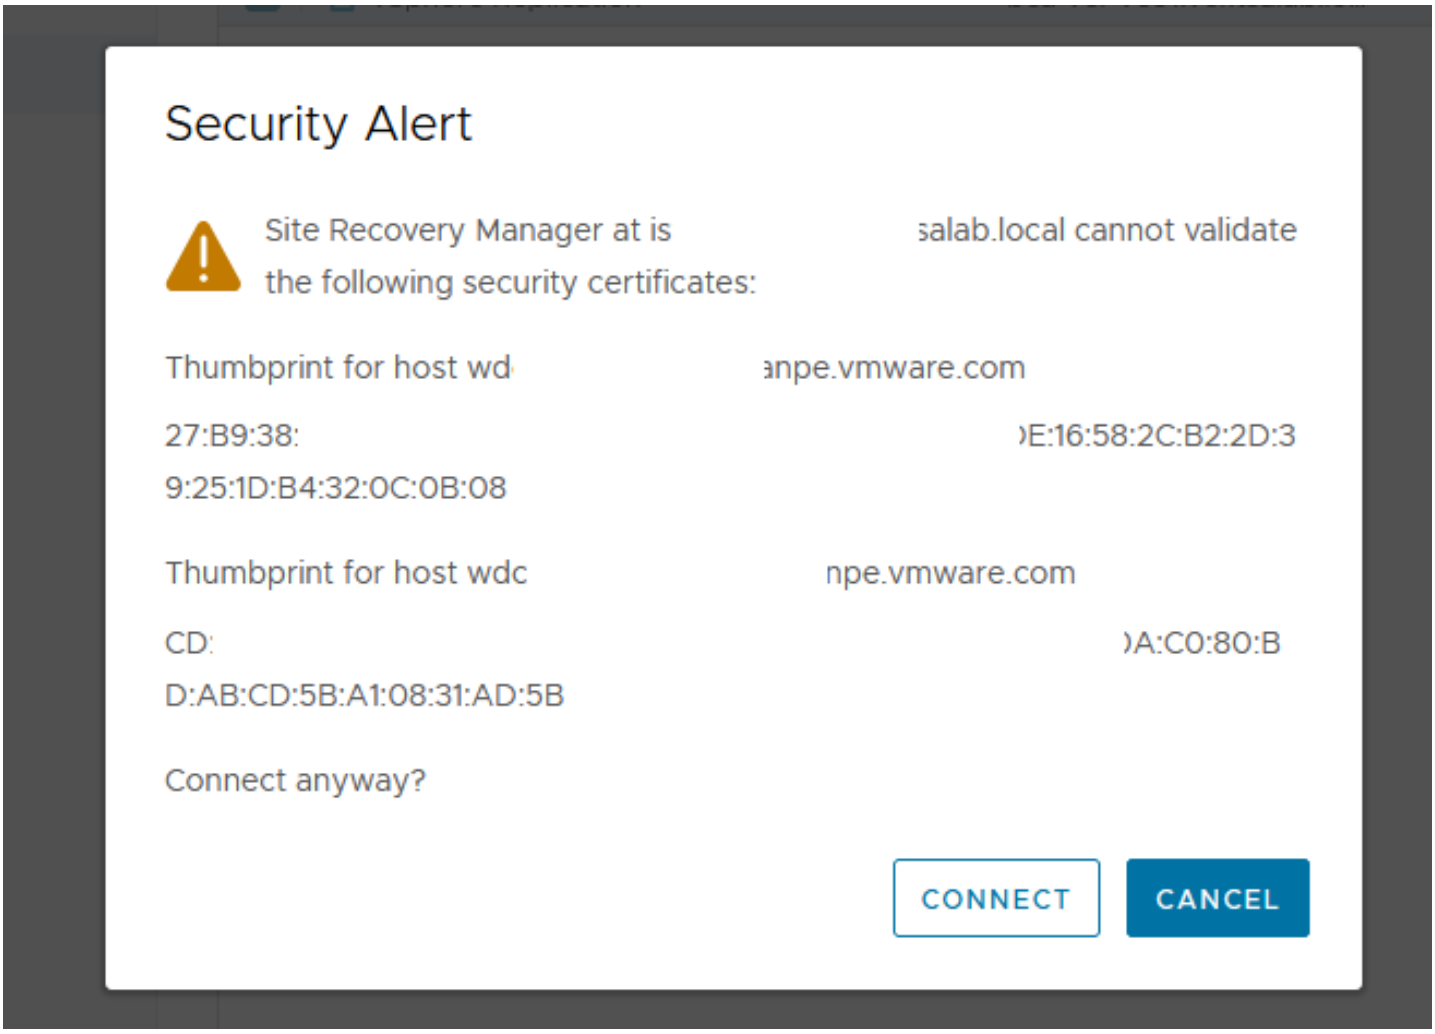

7. Click **Finish** to complete the site pairing process.

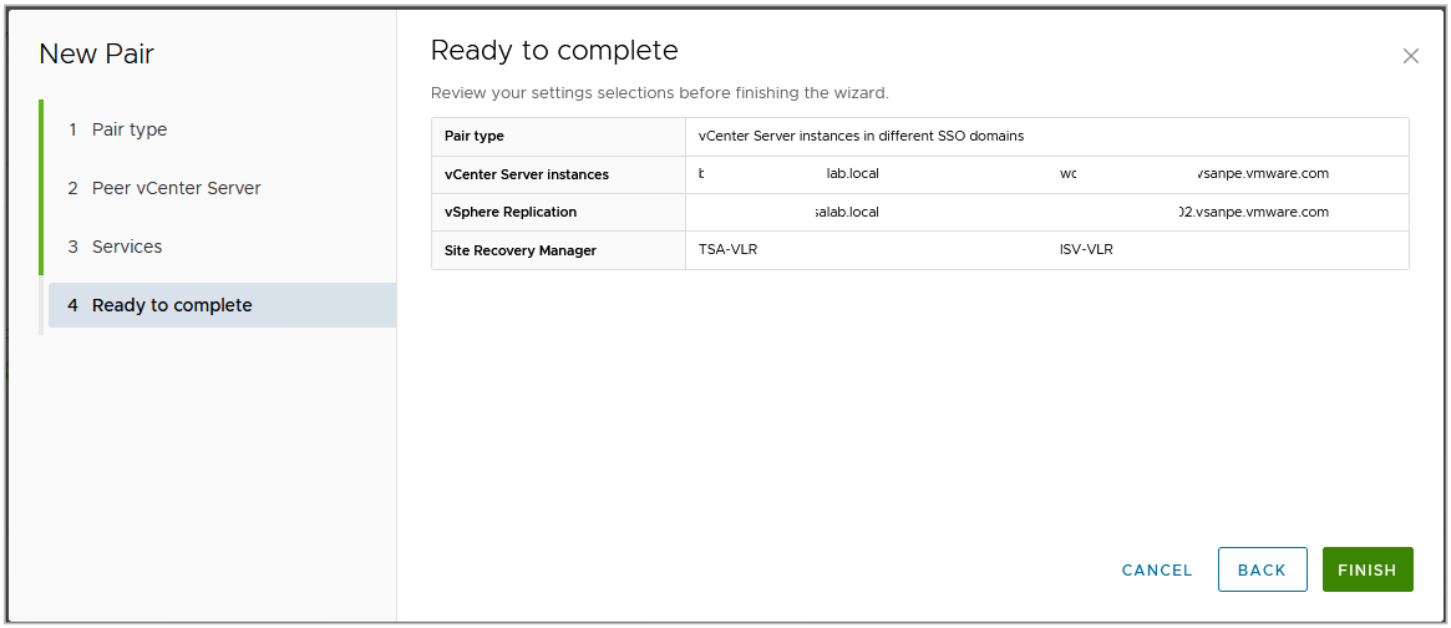

Now we're done with the site pairing exercise. We're ready to starting protecting our mission-critical workloads. The Site Recovery page will look like this screenshot.

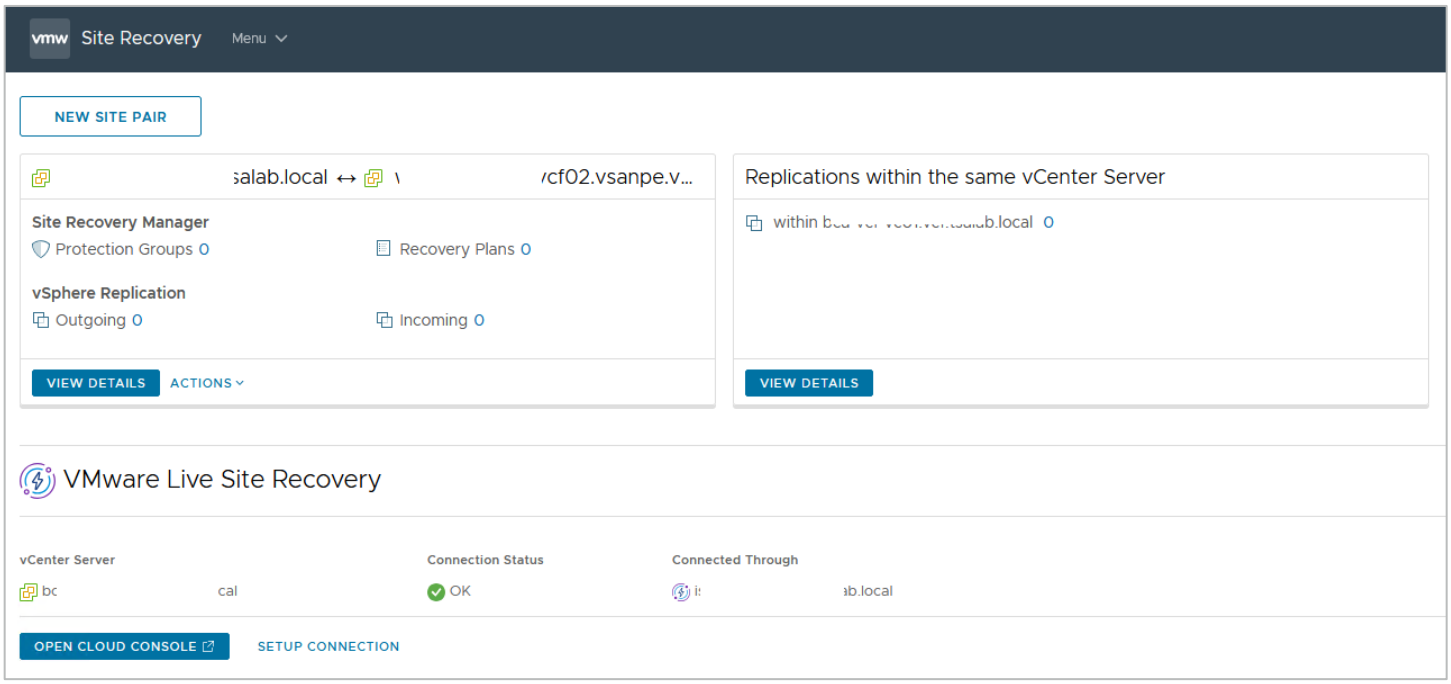

Set up site recovery protection

1. From the site pairing page, click **View Details.**

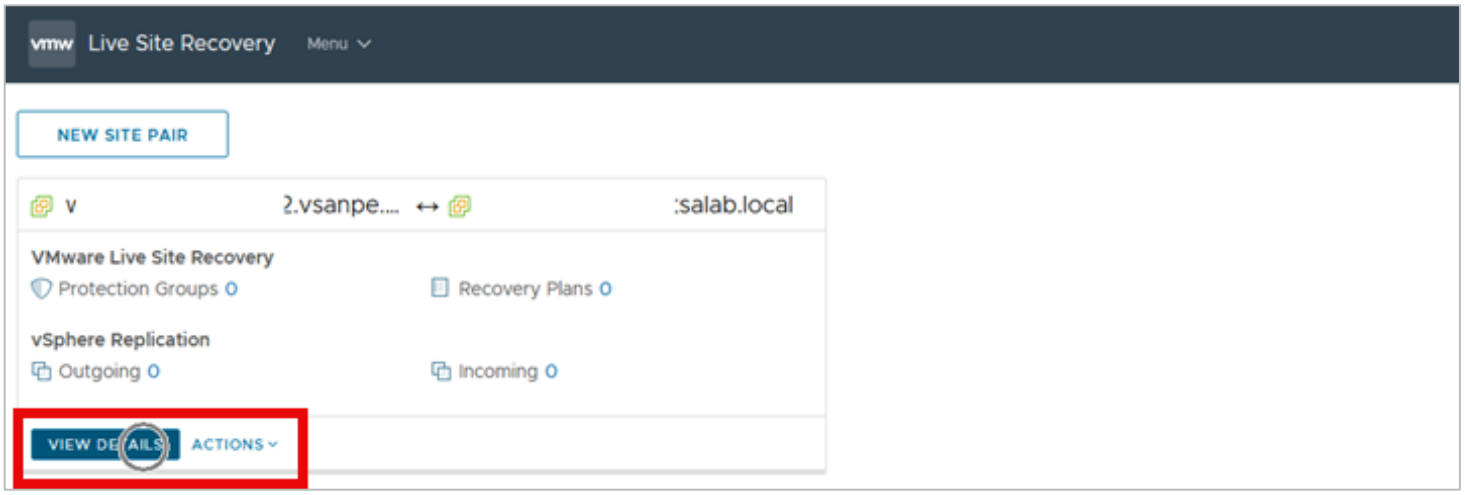

2. Provide the admin credentials for the **protected site's vCenter**, and then click **Login** to complete the initial pairing.

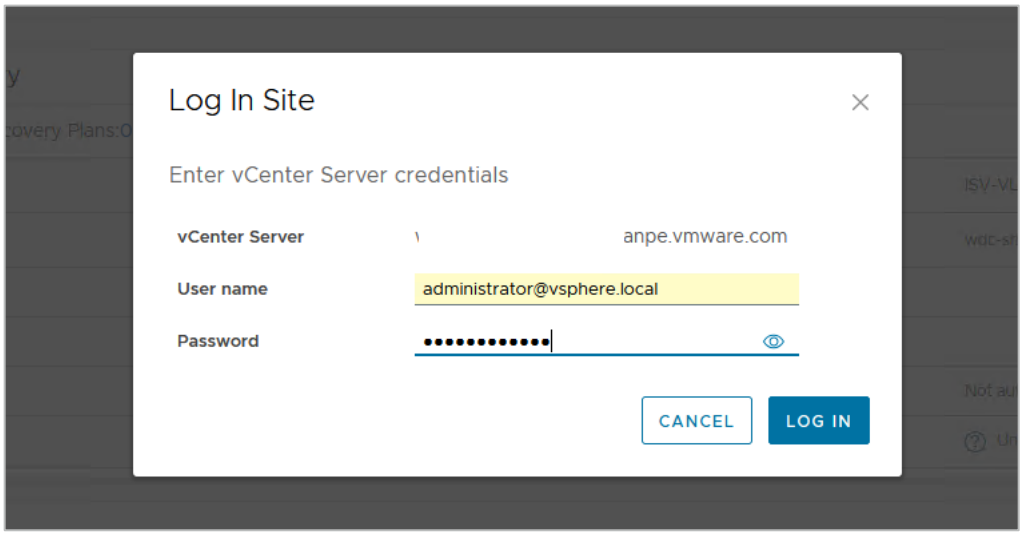

Here is our Live Recovery configuration and administration landing page.

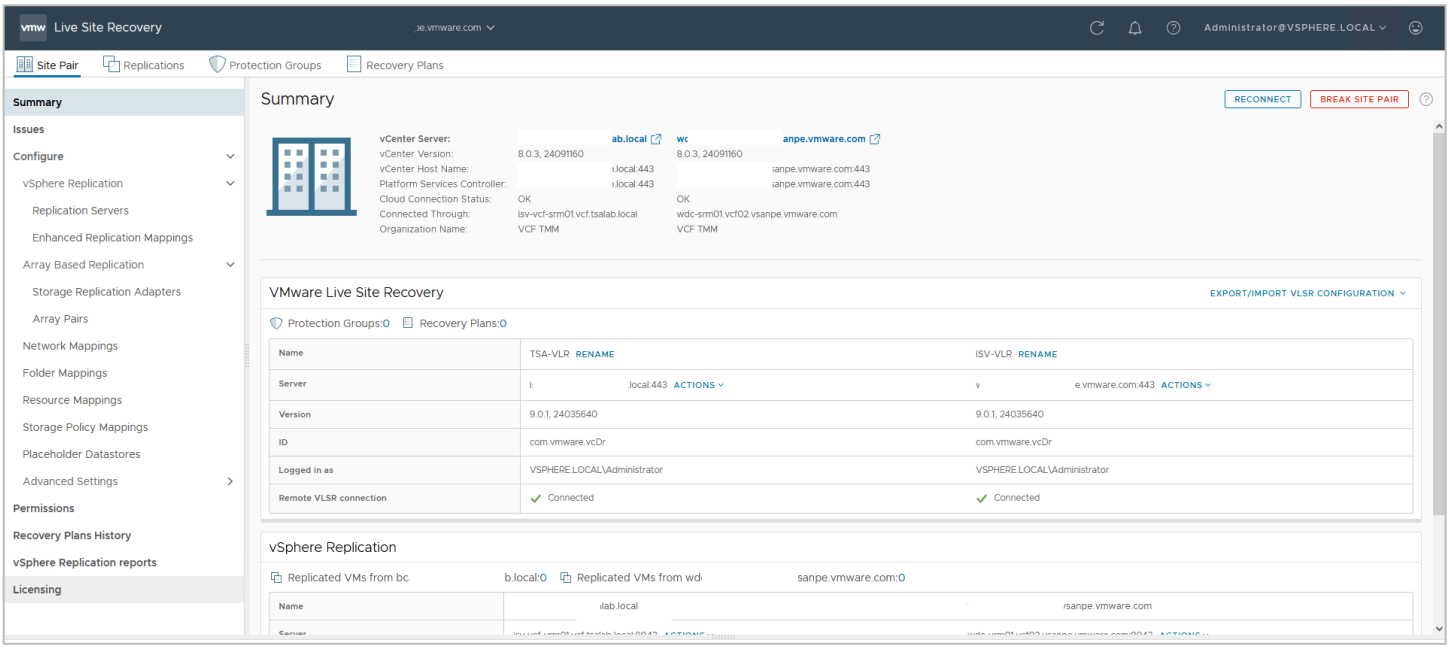

Factors influencing our design and configuration choices

VMware Live Site Recovery allows you to configure an orchestrated workflow of all the actions and steps required to recover a VM, including the guest operating system, applications, processes, and other components. VMware Live Site Recovery does this by using the features and capabilities of the VCF infrastructure and the storage subsystem to create a point-in-time copy of the VM from the source (protected site) to the target (recovery site). VMware Live Site Recovery can use either array-based replication or vSphere Replication to replicate VM data from the source site to the target site. For this paper, we used vSphere Replication.

vSphere Replication can replicate VMs between different storage types

Because vSphere Replication is host-based, it doesn't depend on the underlying storage, so it works with a variety of storage types, including vSAN, traditional SAN, NAS, and direct-attached storage (DAS). Unlike many array replication solutions, vSphere Replication can replicate VMs between the same or even different storage types, like vSAN to DAS, SAN to NAS, and SAN to vSAN, to name a few.

VMware Live Site Recovery can pre-configure recovery plans

When a real or simulated failure occurs at the protected site, admins can initiate the pre-configured recovery steps and actions in their recovery plans. These steps include, among others:

- The order in which VMware Live Site Recovery recovers the protected VMs.
- The network to which the recovered VMs are connected.

- Whether to customize or change the IP addresses for the recovered VM, or let them obtain such addresses from an available/accessible DHCP server.
- In-guest configuration scripts to run on the recovered VMs.

When an admin initiates this recovery plan, VMware Live Site Recovery prepares the VMs for recovery using the vSphere Replication-created copy of the VM data. The VMs are added to inventory, connected to the necessary resources (networks, folders, resource pools, and storage policies), powered on in the specified order, and customized as needed. If the workflow includes running scripts inside the VMs, the guest operating system is instructed to call and execute the scripts (of course, the scripts must exist on the VMs and be accessible during the recovery process).

Here's a description of the protection and recovery workflow you'll configure for the exercise:

- The SQL Server instances run on Windows VMs and are joined to the Active Directory domain services (ADDS) infrastructure. For this reason, you should have the Domain Controllers available and functional before powering on the SQL Server VMs.
- The SQL Server instances are clustered in a 3-node, Always-on Availability Group configuration. Clustering SQL Servers requires the use of a Windows Server Failover Cluster (WSFC). You'll use a file share witness (a folder located on one of the Domain Controllers) as the quorum option for this configuration.
- We specifically use availability groups in this guide and demonstration because (at the time of this writing):
  - The default storage option for VMware clouds is vSAN.
  - The default replication option for vSAN is vSphere Replication.
  - vSphere Replication doesn't currently have the capabilities to replicate disks used for shared-disk Windows clustering.
  - Although the scripts and all other required steps are similar, the factors mentioned above preclude the use of the steps documented in this guide for protecting and recovering Microsoft SQL Servers configured in shared-disk mode—Always On Failover Clustering Instance (FCI).
- In steady-state operation, applications, scripts, and processes access the SQL Server instance and the database through a common name: the **listener**. The listener is a host name that resolves to a specific IP address (or set of IP addresses). The Domain Controllers provide the DNS service, which manages this resolution. The listener must be available and accessible to the services provided by the SQL servers.
- Usually, the IP address segments in the protected site are different from the ones used in the recovery site.
- It is possible to extend the network segments from the protected site to the recovery site. Because the mechanism for achieving this configuration differs among the various VMware cloud brands, including it in this guide is impractical. For simplicity, the exercise includes a workflow for changing the IP addresses of recovered VMs to match those available at the recovery site.
- This IP address change means that you'll need to change the IP addresses of the VM (a trivial task in VMware Live Site Recovery) and the listener.
- VMware Live Site Recovery can't automatically change VM application configurations because it doesn't have knowledge of the applications that run inside the VM. For this purpose, you'll use the VMware Live Site

Recovery script-triggering feature to instruct the guest operating system to run a script that will change the IP address of the listener and update the DNS record after the recovery.

Stability and security of virtualized Domain Controllers

Around 2012, with compute resources growing and virtualization becoming mainstream in IT environments, dedicating a physical server to running a Domain Controller became impractical and inefficient from a cost and ROI perspective. But there were some issues with the stability and security of virtualized Domain Controllers. To address these problems, Microsoft implemented measures to make virtualized Domain Controllers safer and more stable, including guardrails to prevent a malicious actor from cloning or copying them.

VM-Generation ID makes Domain Controller virtualization safer

At a high level, a Domain Controller has a complete copy of the domain's users, passwords, and other secrets, making it difficult to minimize or mitigate an attack. VM-Generation ID (among other capabilities) helps protect virtualized Domain Controllers in several ways:

- It stores and tracks a unique counter for every copy of a virtualized Domain Controller. The hypervisor assigns a counter to the VM. In vSphere, this is the **VM Gen-IDx** value you see in a Windows VM's .vmx file.
- When the Domain Controller boots up, it reads this counter from its configuration file and then stores it internally.
- This counter persists over the lifetime of the VM unless a specific type of operation is performed on it. These actions alter the state and identity of the VM, so whenever any are performed, the hypervisor changes the counter.
- The next time the Domain Controller is powered on, Windows reads its generation ID, compares it to what was previously stored, and discovers there is a mismatch.
- When this happens, Windows immediately performs several steps in response to the disparity and triggers the VM-Generation ID safety measures. Refer to the following document for a more detailed discussion of virtualization-based safeguards: Safely virtualizing Active Directory Domain Services (AD DS).

Restoring a Domain Controller triggers VM-Generation ID change

- VLSR recovery workflow includes bringing up A REPLICATED COPY of a Protected VM at the Recovery Site when (in a real Disaster event) the real Domain Controller is unavailable) or in a simulated DR exercise (when the VM is recovered to a fenced-off "Test" network). Recovering a Domain Controller requires us to instantiate a replicated copy of a real Domain Controller. Such a "Copy" operation automatically changes the VM-generation ID of the Domain Controller, which then automatically triggers the Domain Controller safety responses from Windows.
- One of the responses is an instruction to the Domain Controller to (among other things) reset its InvocationID and discard its RID Pool. For all practical purposes, the Domain Controller is no longer a Domain Controller at this point, due to the change in its VM-generation ID. Windows then updates the VM-generation ID it had stored previously to match the new one provided by the hypervisor. The VM then obtains a new set of RID

Pool from the RID Master, and life is good. Well, we have abbreviated the complete narrative for our purposes, but what is of relevance to us for this Guide is that, in spite of the fact that recovering a Domain Controller with VLSR triggers Windows to invoke the Virtualized Domain Controller Safety feature, doing so is a supported, repeatable, more efficient, reliable, and faster option than anything else available as of the time of this writing.

Logical topology of the VMware Live Site Recovery infrastructure

Now that we know our desired outcome and the considerations governing our ability to achieve it, we are ready to proceed.

Here's an approximate representation of the logical topology of our VMware Live Site Recovery infrastructure:

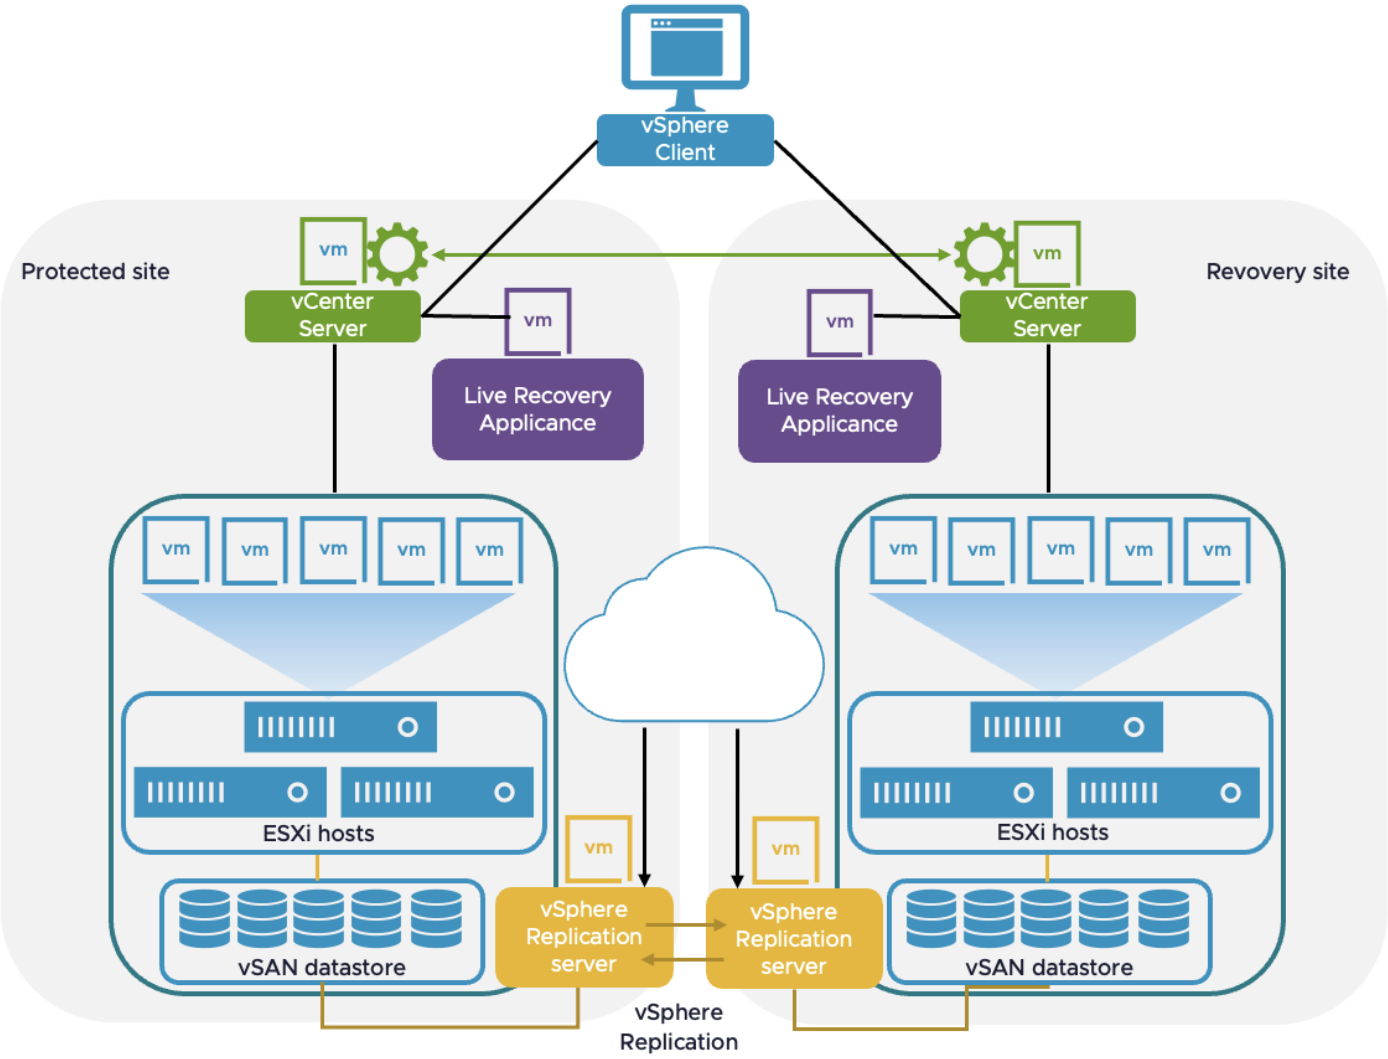

In our example, our storage platform is vSAN (the default option for vSphere-based Cloud platforms), so the VMware vSphere Replication Server will be responsible for replicating our protected VMs from the Protected Site to the Target Site (and vice versa). No special configuration is required for this part at this point.

Create mappings with network pairing

We'll ignore Array Based Replication because it doesn't apply to vSAN, which is the default storage option for VCF and all vSphere-based cloud offerings.

NOTE: VLSR supports non-vSAN platforms which provide their own VLSR-compatible storage replication adapters.

1. Select **Replication Servers**.

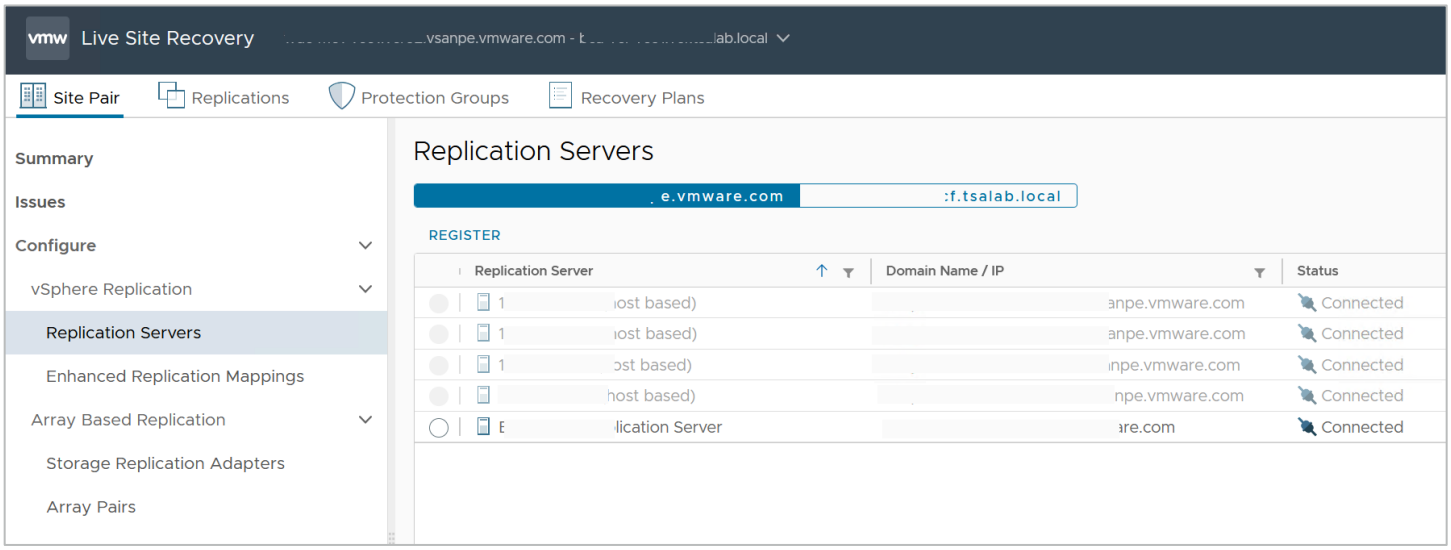

Network pairing lets you map the network segments on one side to a corresponding segment on the other.

2. Click **New** to begin creating a mapping.

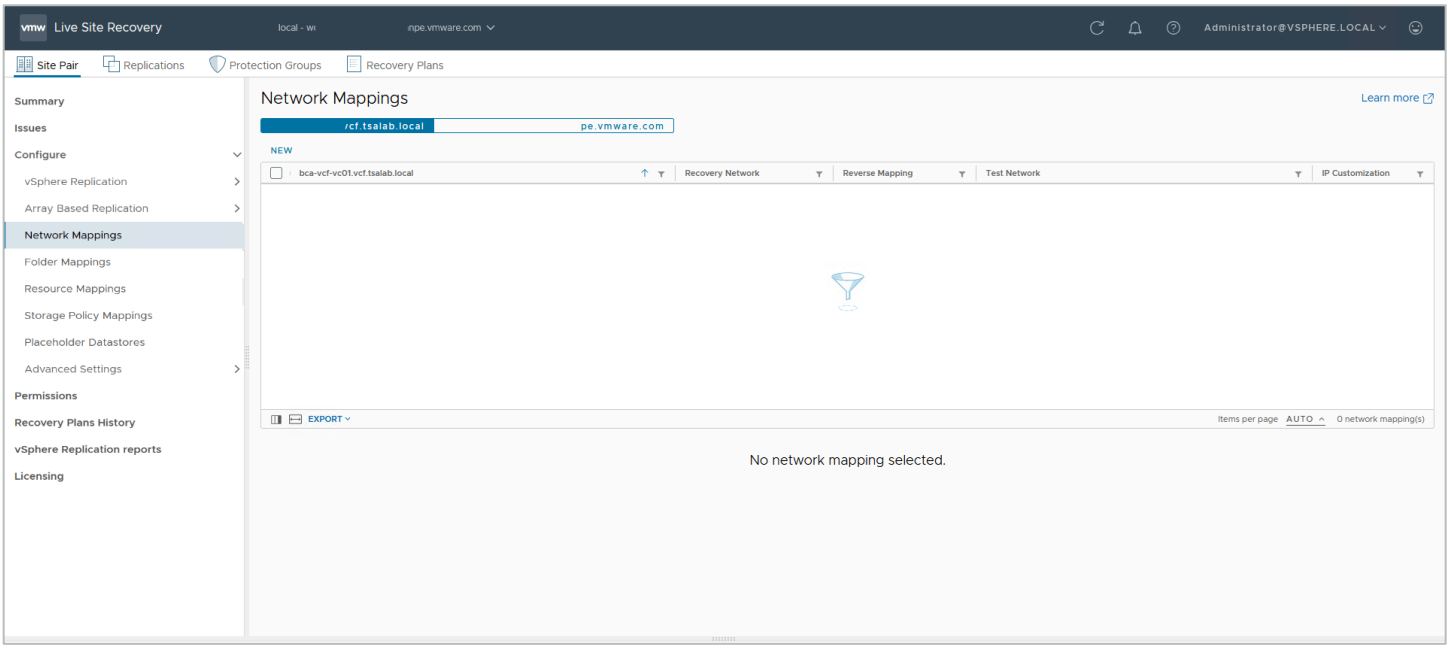

3. Select **Prepare mappings manually** and click **Next**.

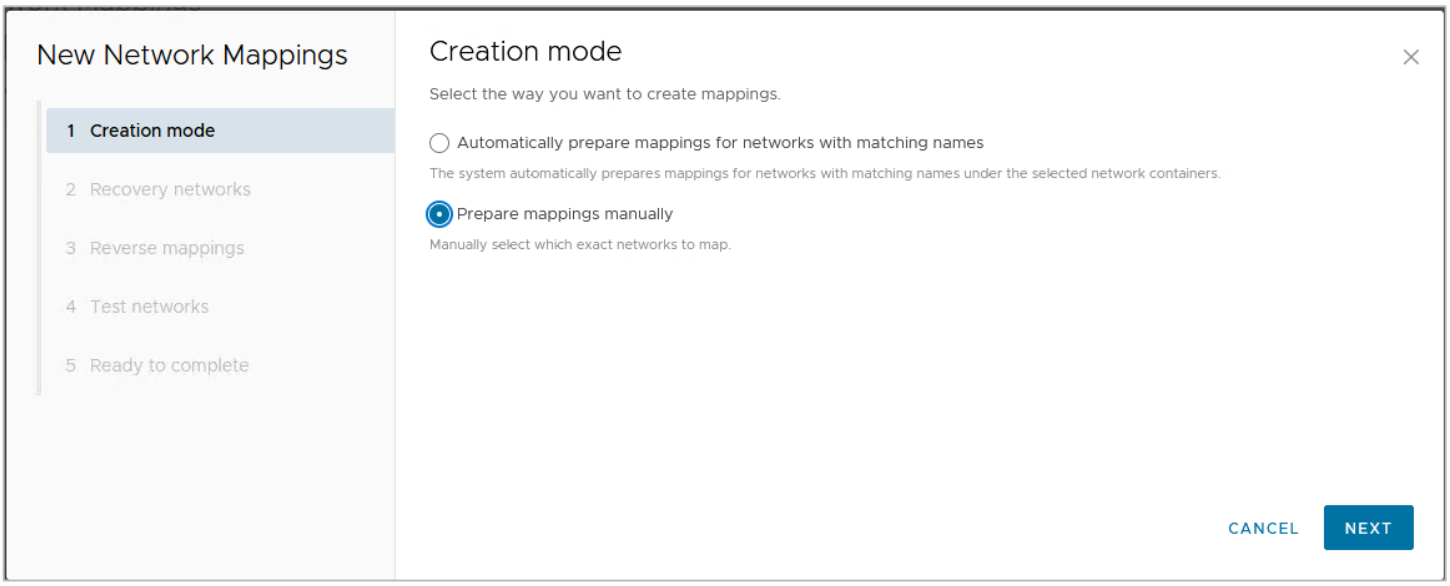

You have the option to create a mapping of the networks at either the virtual distributed switch (vDS) level as a unit, or you can create a mapping at the individual port group level. We'll demonstrate the fine-grained flexibility in VMware Live Site Recovery by mapping select port groups from the protected site to corresponding port groups on the recovery site.

4. Select the check box near each port group on the protected site and the corresponding port group on the recovery site you want to map each to. Then click **Add Mappings** and **Next**.

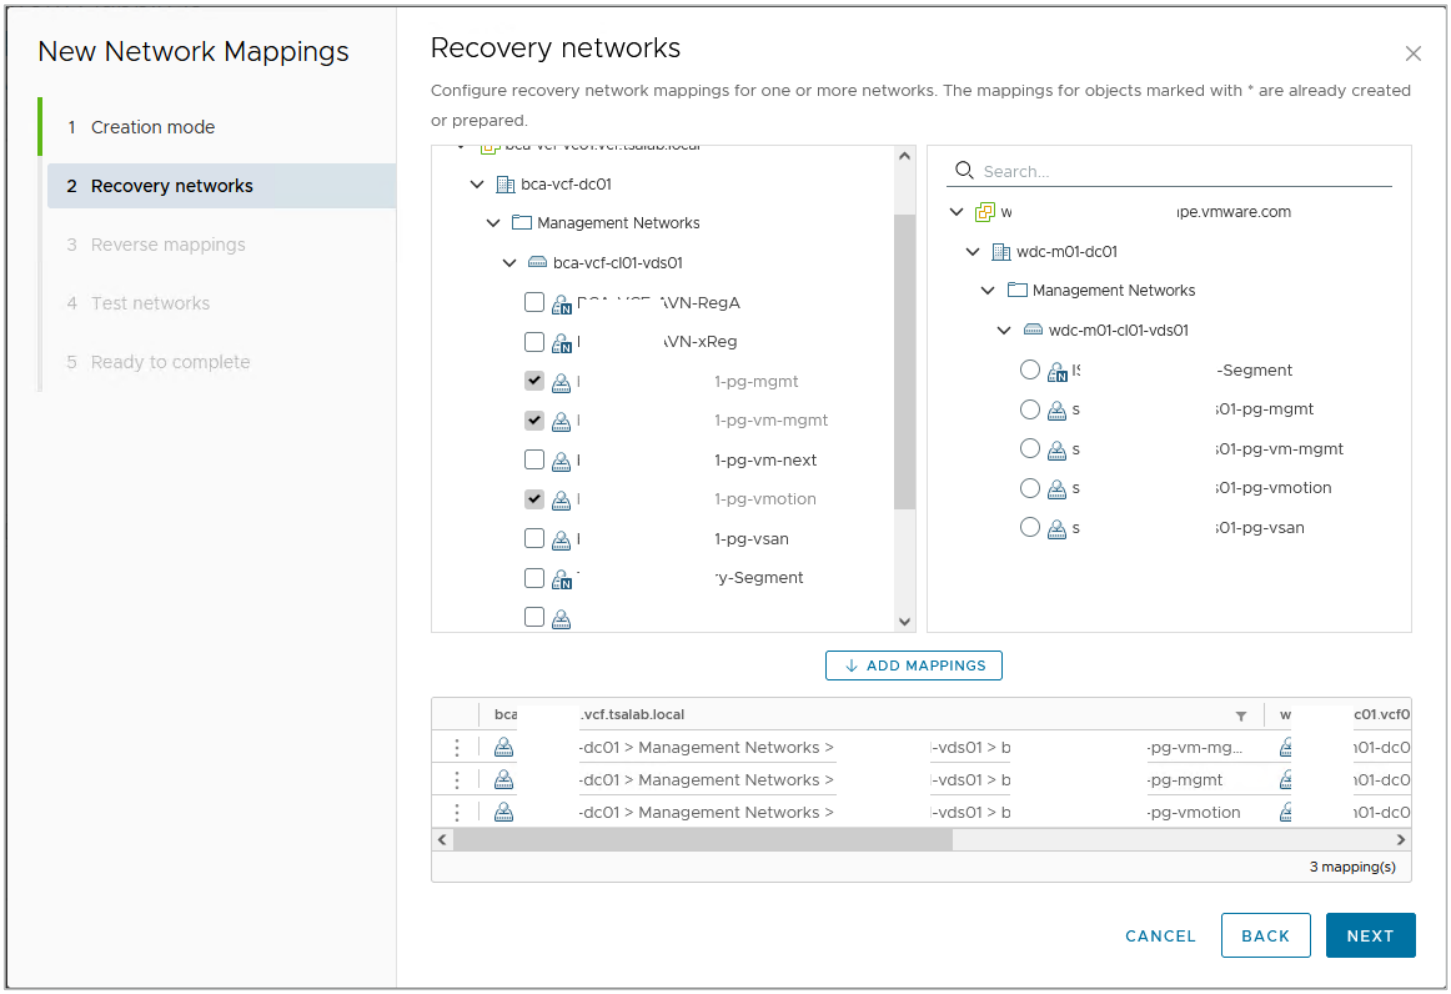

5. Check the option to automatically create a reverse mapping (so you don't have to do it manually) and click **Next**.

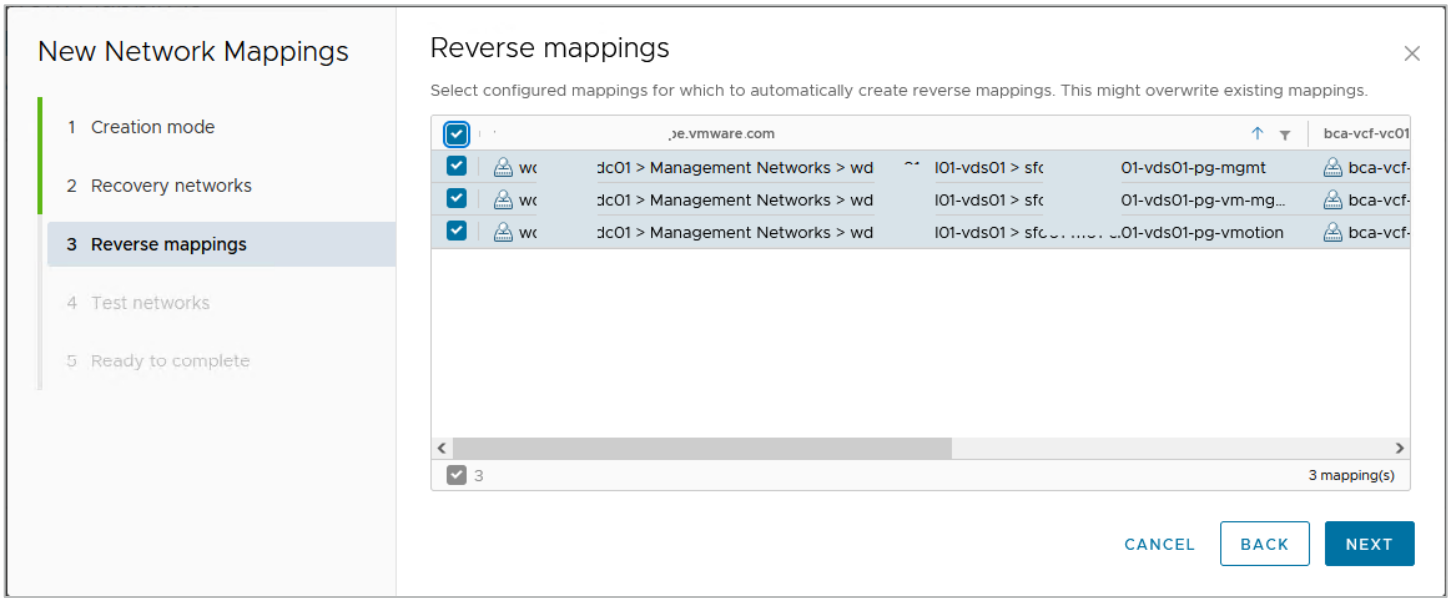

6. Click **Finish** to complete the configuration.

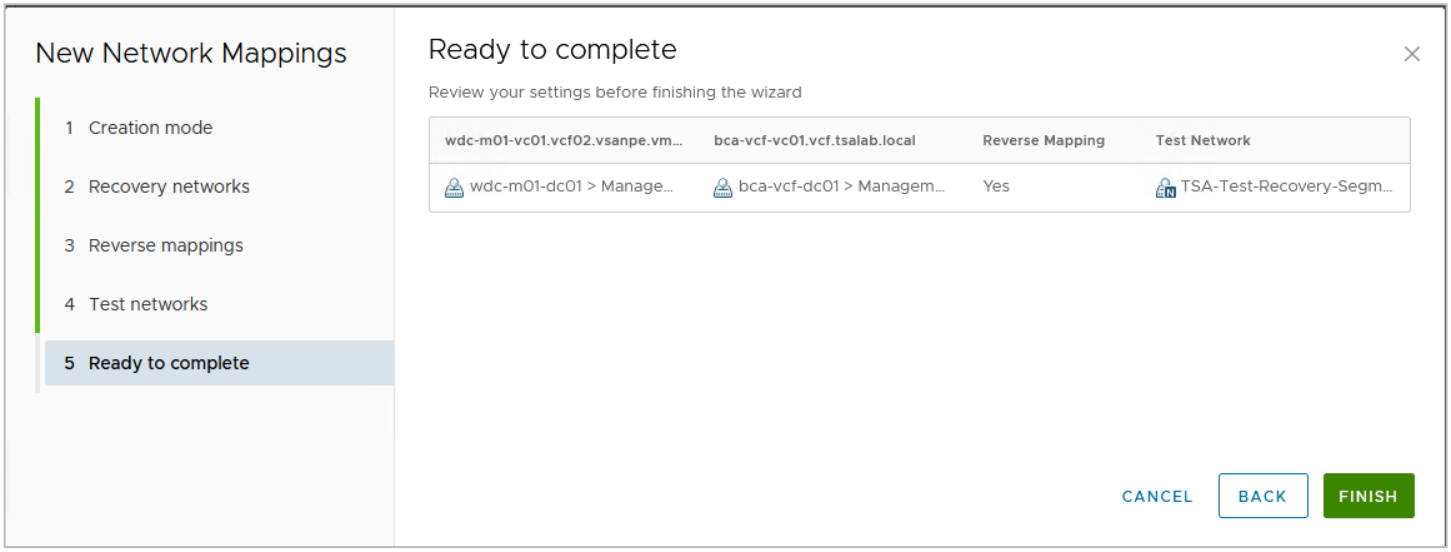

7. Click **Continue** to acknowledge and dismiss the warning about possible impact to existing protected VMs.

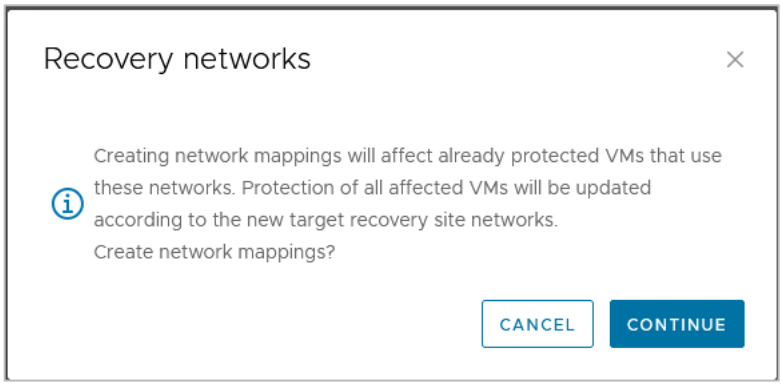

Create a test recovery network

One of the most compelling features in VMware Live Site Recovery and why it is much preferred over competing BDCR orchestration solutions (or a manual option) is being able to conduct simulated/test disaster recovery exercises without impacting the production environment. Admins can demonstrate and prove their infrastructure disaster recovery readiness by conducting a recovery of the protected workloads into the recovery site while the protected workloads continue to provide uninterrupted services at the protected site. VMware Live Site Recovery does this by bringing up a copy of the protected workload in an isolated network segment at the recovery site. VMware Live Site Recovery creates this isolated network by default, but admins can specify their own recovery test (aka "bubble") network. The default isolated network is inaccessible to anything outside of the bubble. But what if admins want to demonstrate the functionality and accessibility of recovered workloads to their auditors? They can do this by recovering the workloads into a specific network of their choice (assuming they have such a controlled network in place).

Create an isolated network port group

Creating an isolated network port group in a VCF infrastructure is a simple operation. Although NSX administrative tasks are outside the scope of this paper, let's briefly describe how to create such isolated segments for ease of reference and completeness.

1. In NSX Manager, from **Segments**, click **Add Segment**.

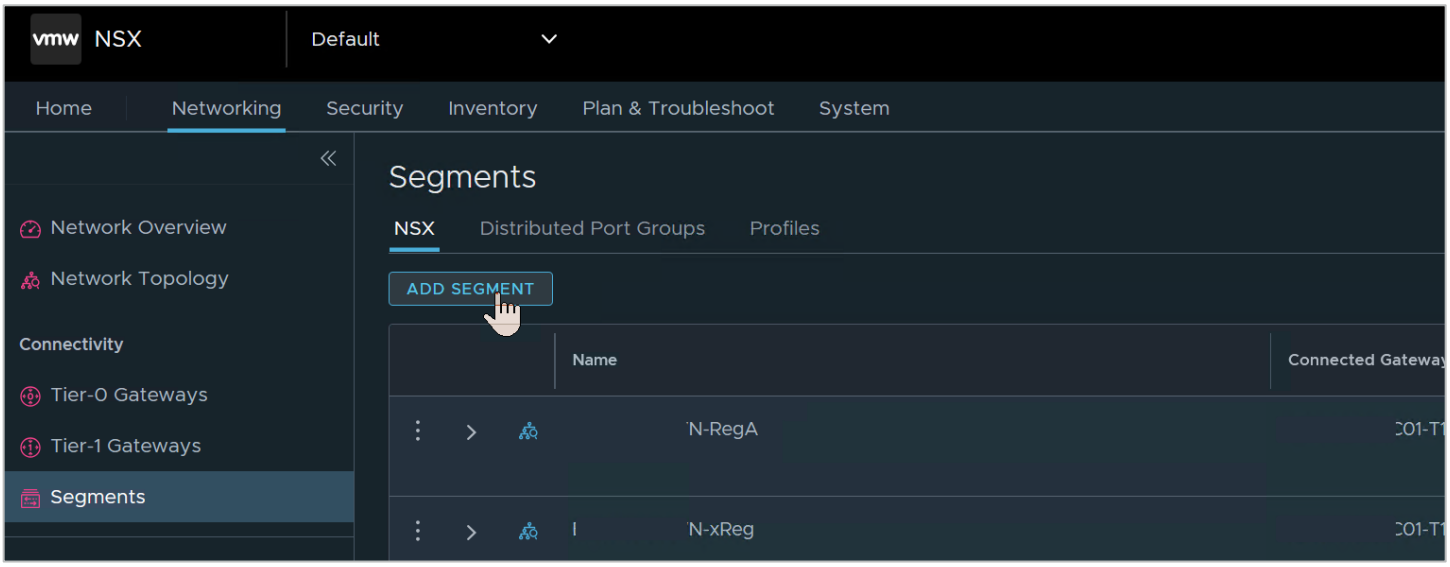

2. Give the new segment a descriptive name.
3. Don't specify a gateway in the **Connected Gateway** menu.
4. Select an appropriate **Overlay Transport Zone** for the segment.
5. Click **Save**.

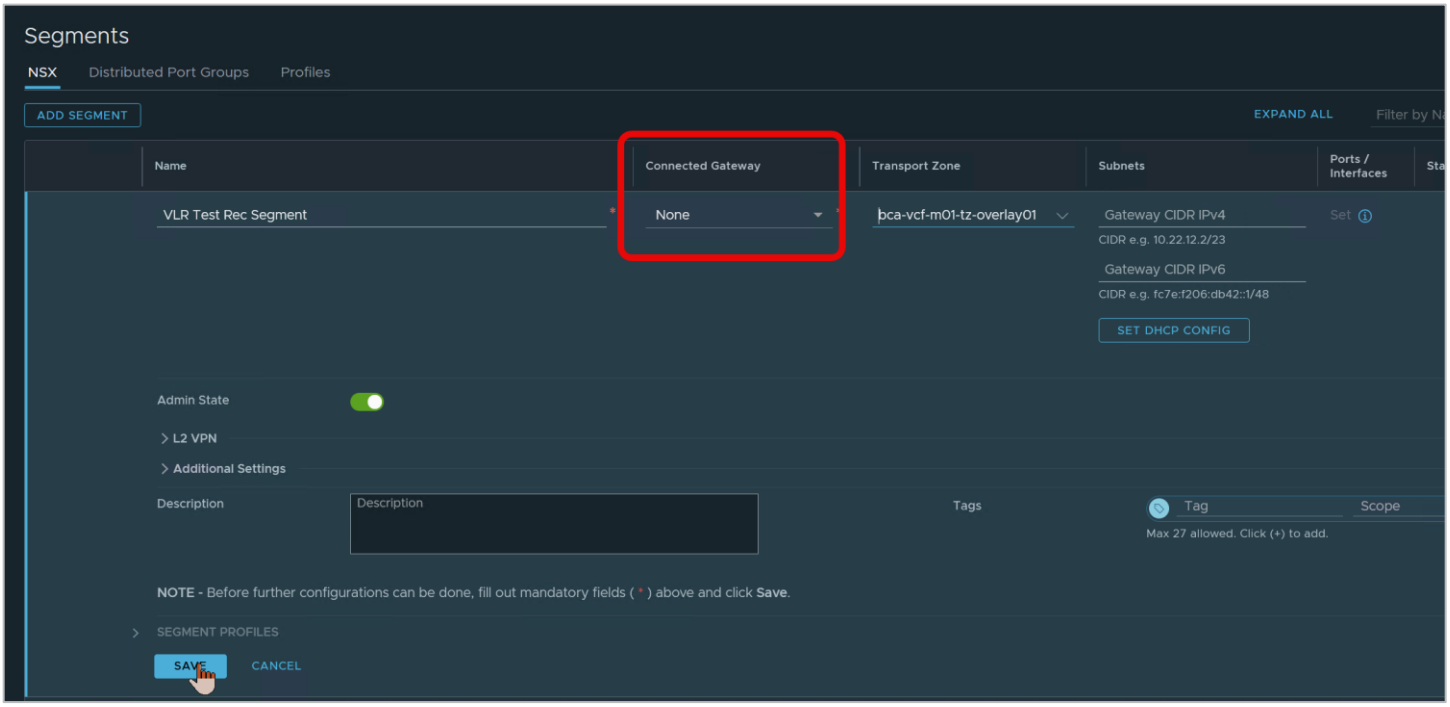

Specify the test recovery network

The segment you created in the previous section will appear as a port group on all the connected ESXi hosts in the cluster. Because it is unrouted, network traffic over the port group will be restricted to only the VMs directly connected to it.

You'll use this unrouted segment/port group to manually specify your desired isolated test recovery network in VMware Live Site Recovery.

By default, VMware Live Site Recovery automatically creates an internal, isolated network for test failovers. You'll change these.

1. Select the network mapping for which you want to specify a desired routed network.
2. Select the elipses **(…)** menu.
3. Select **Edit Test Network Mapping**.

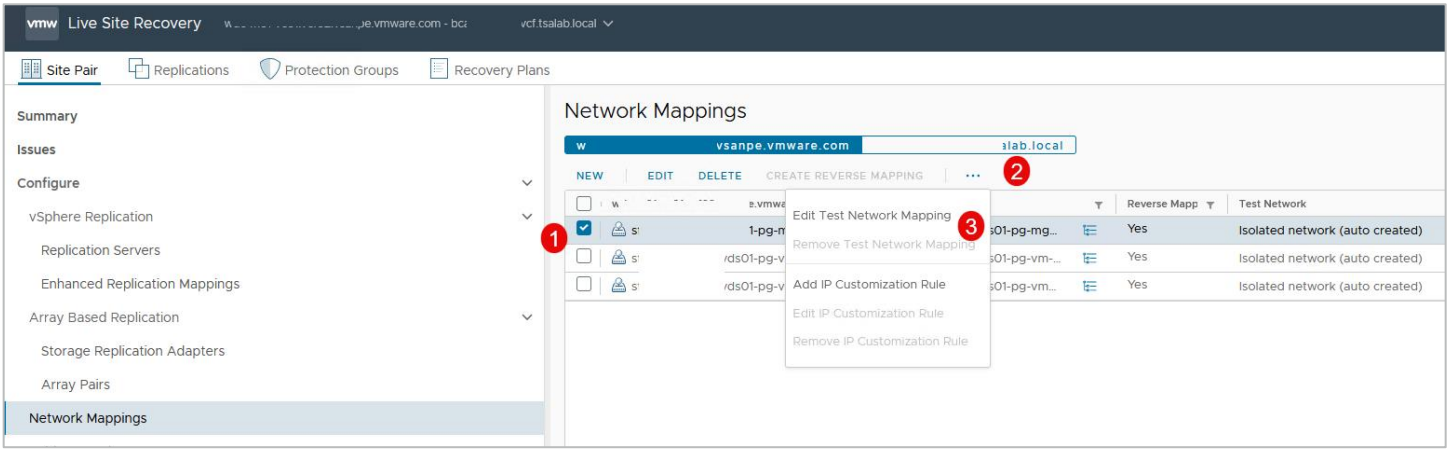

4. Select the **Select a specific network** option.
5. Select the pre-configured isolated network/segment and click **Save**.

6. When done, click **Next**.

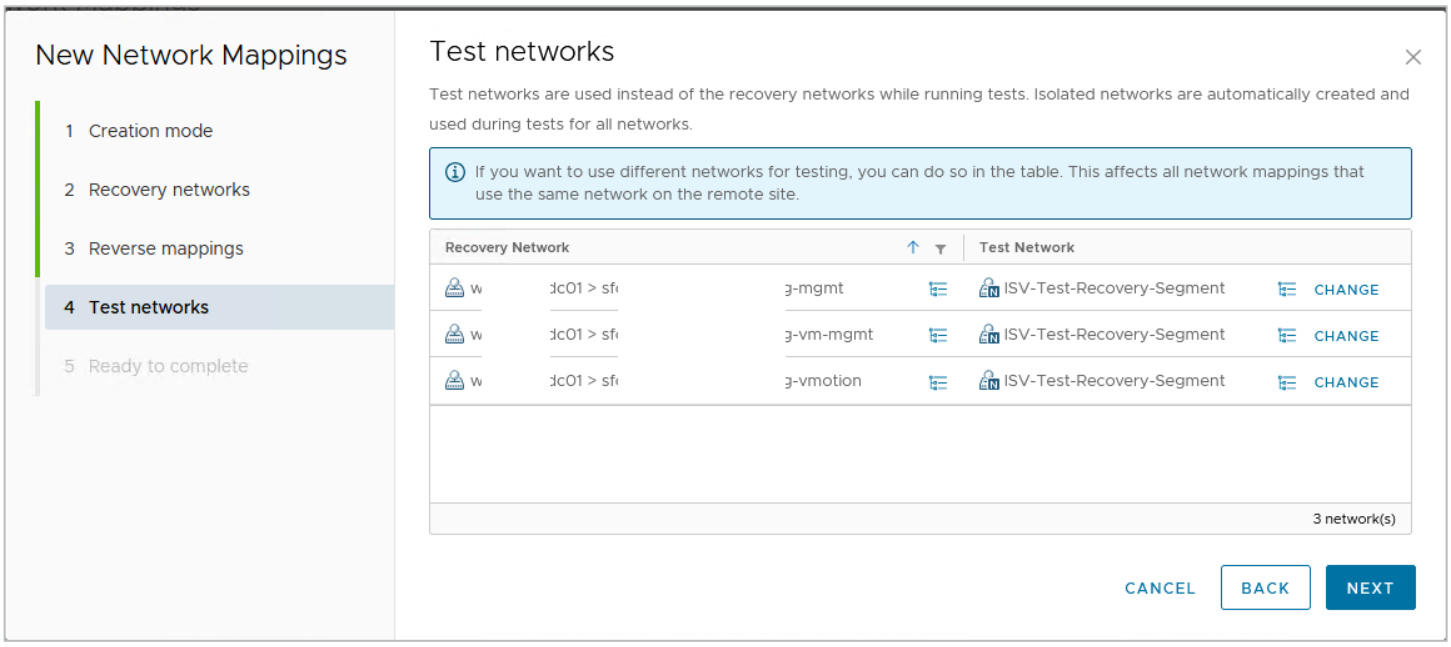

7. Click **Finish** to proceed.

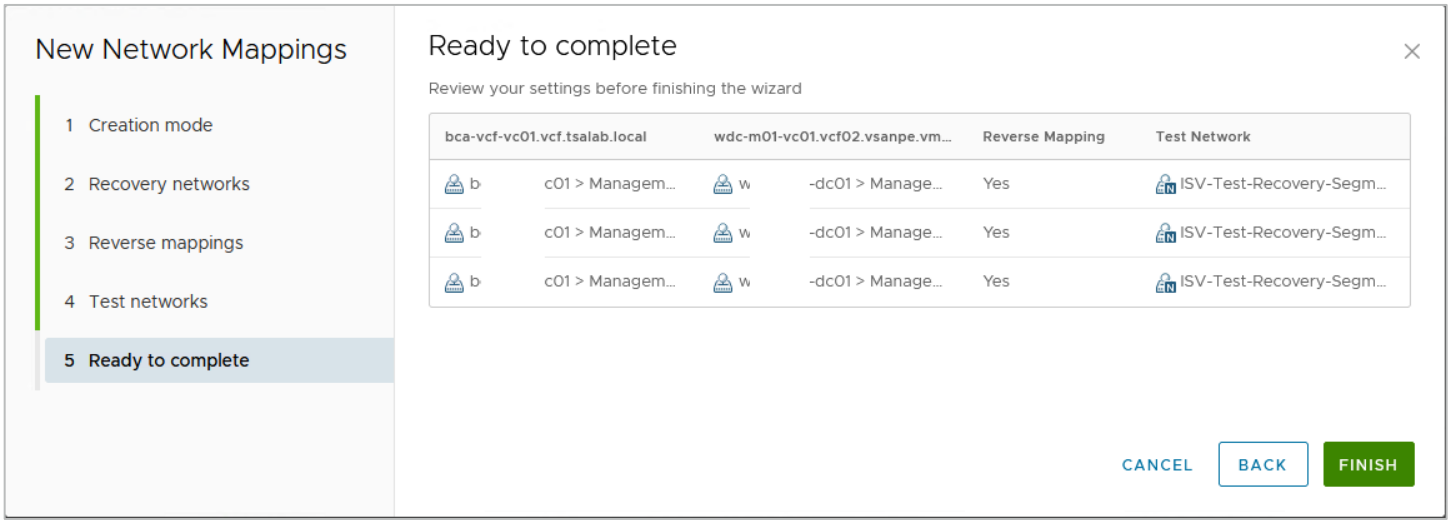

Create folder mappings

Folder mappings help to organize protected and recovered VMs in a logical and intuitive fashion, so let's create one:

1. Click **New** to begin.

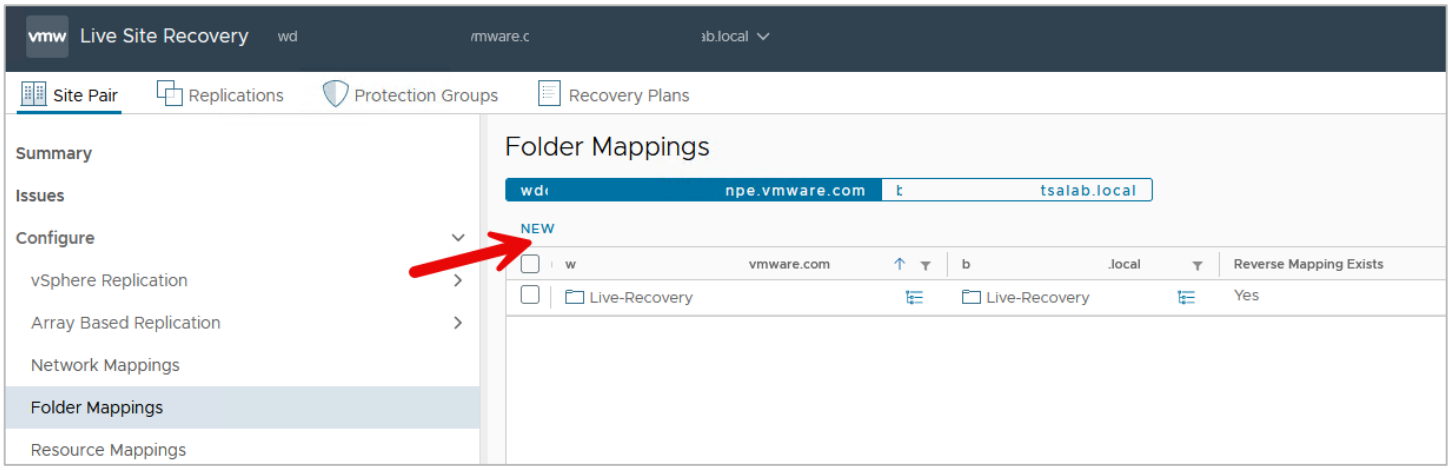

2. Click **Next**.

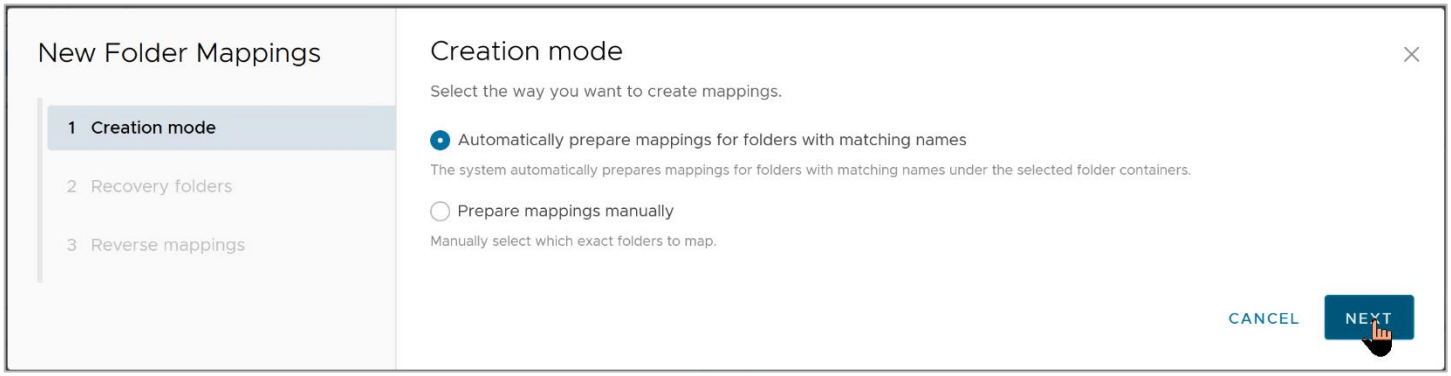

3. Select the VM folders to match up and click **Next**.

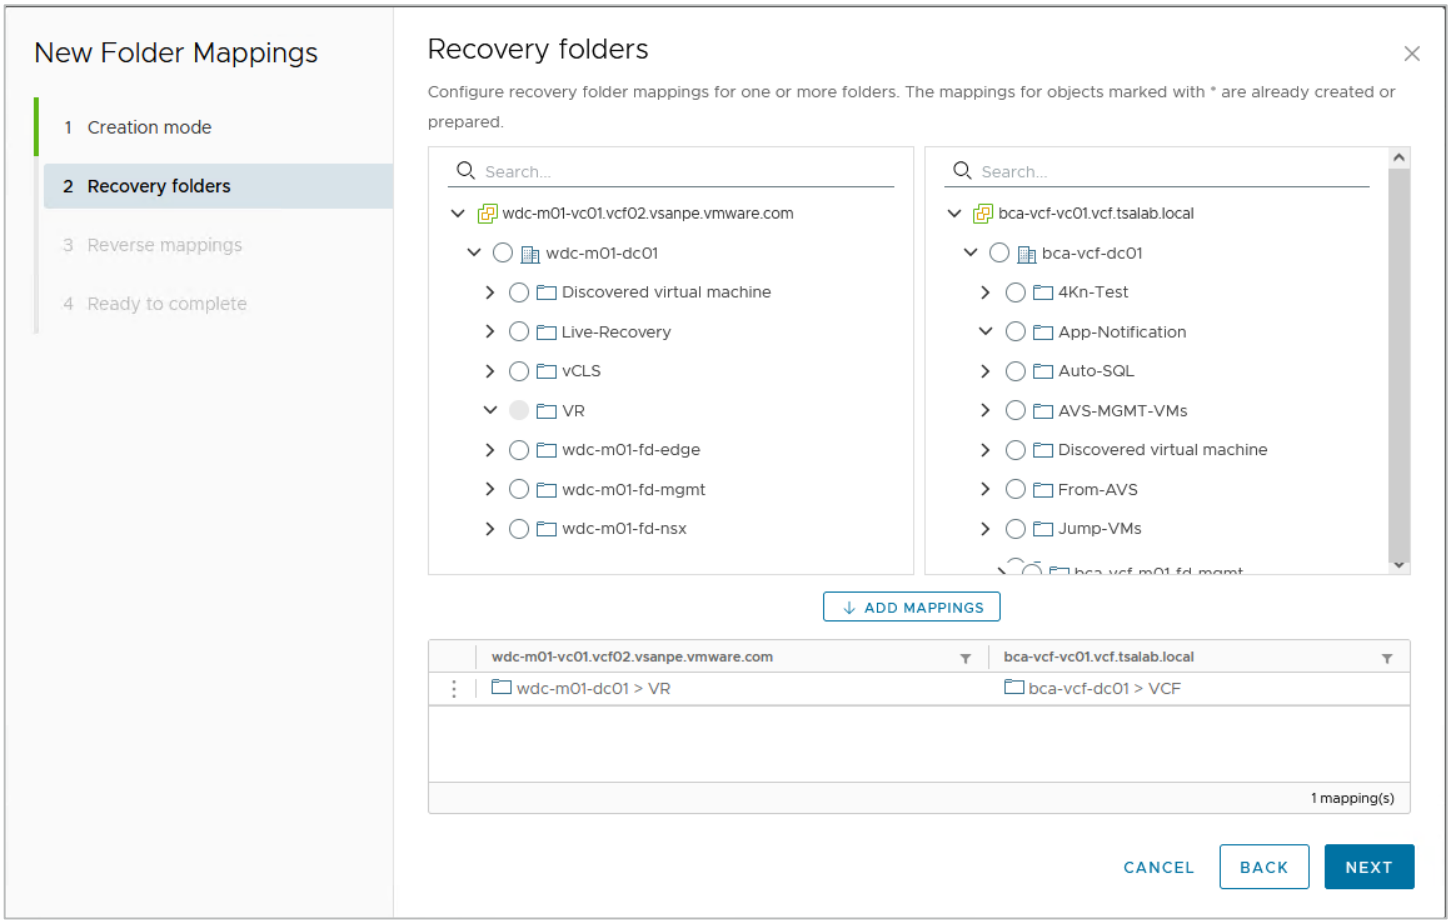

4. Select the check box or boxes to accept the option to create a matching folder map in the opposite direction automatically.
5. Click **Next**.

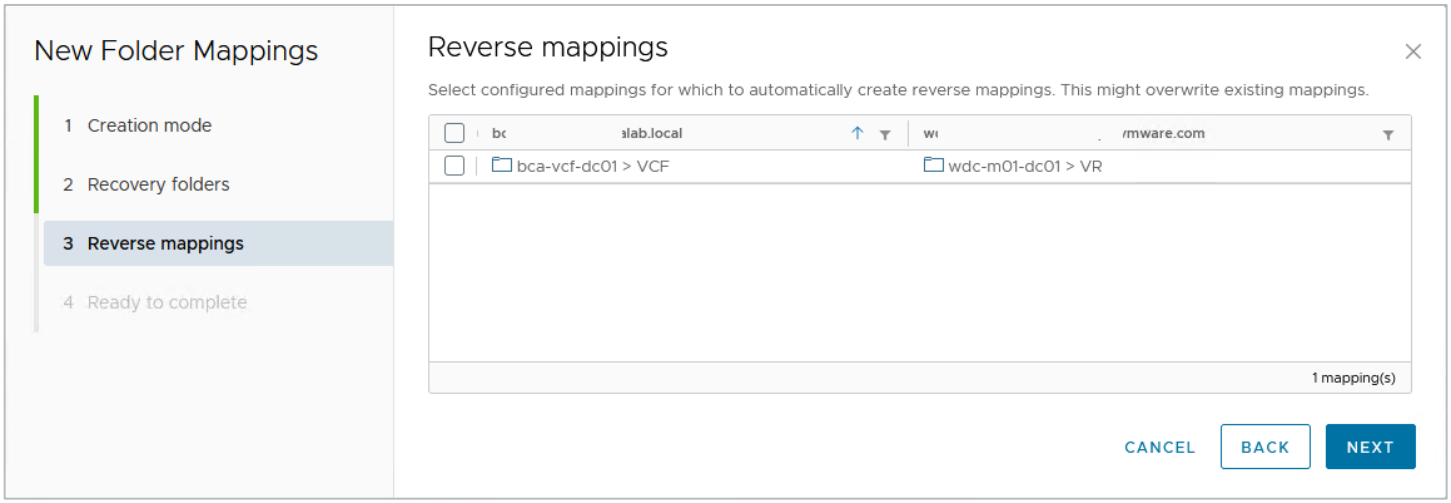

6. Click **Finish**.

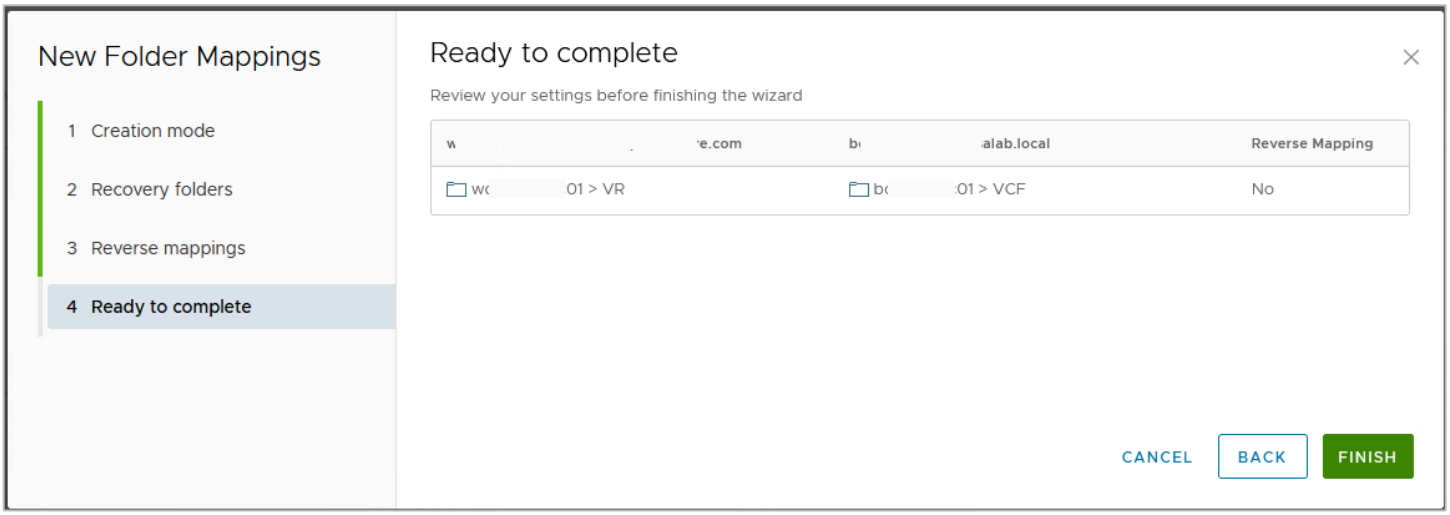

Create resource mappings

You'll map resources at the highest level possible (cluster level, in this case).

1. Click **New** to begin.

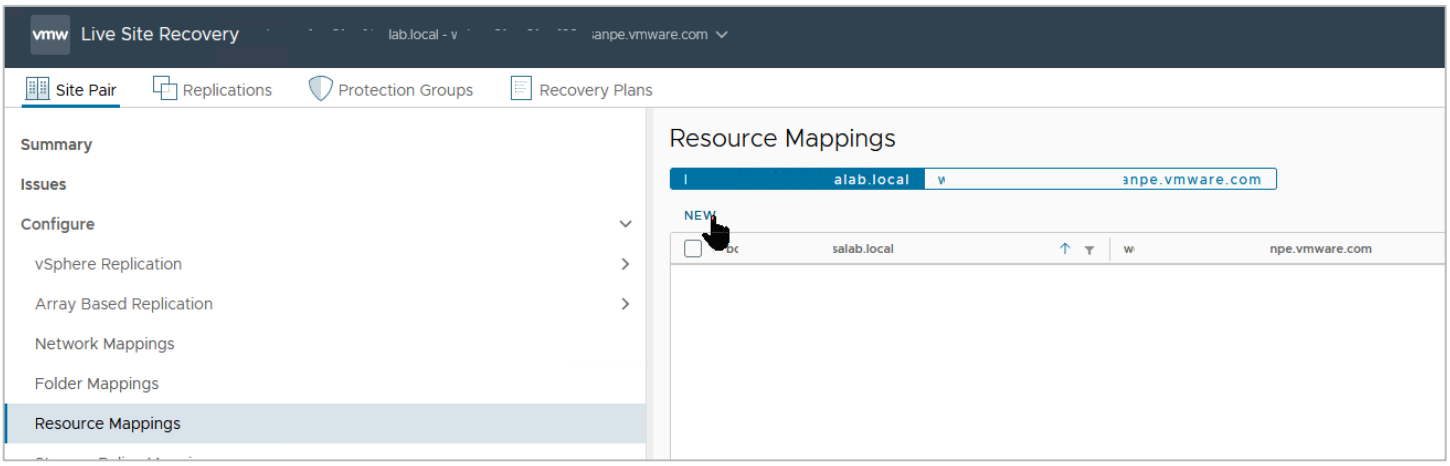

2. Select the cluster containing your protected workloads and map it to the cluster you would like them placed in at the recovery site.
3. Click **Add Mappings** and click **Next**.

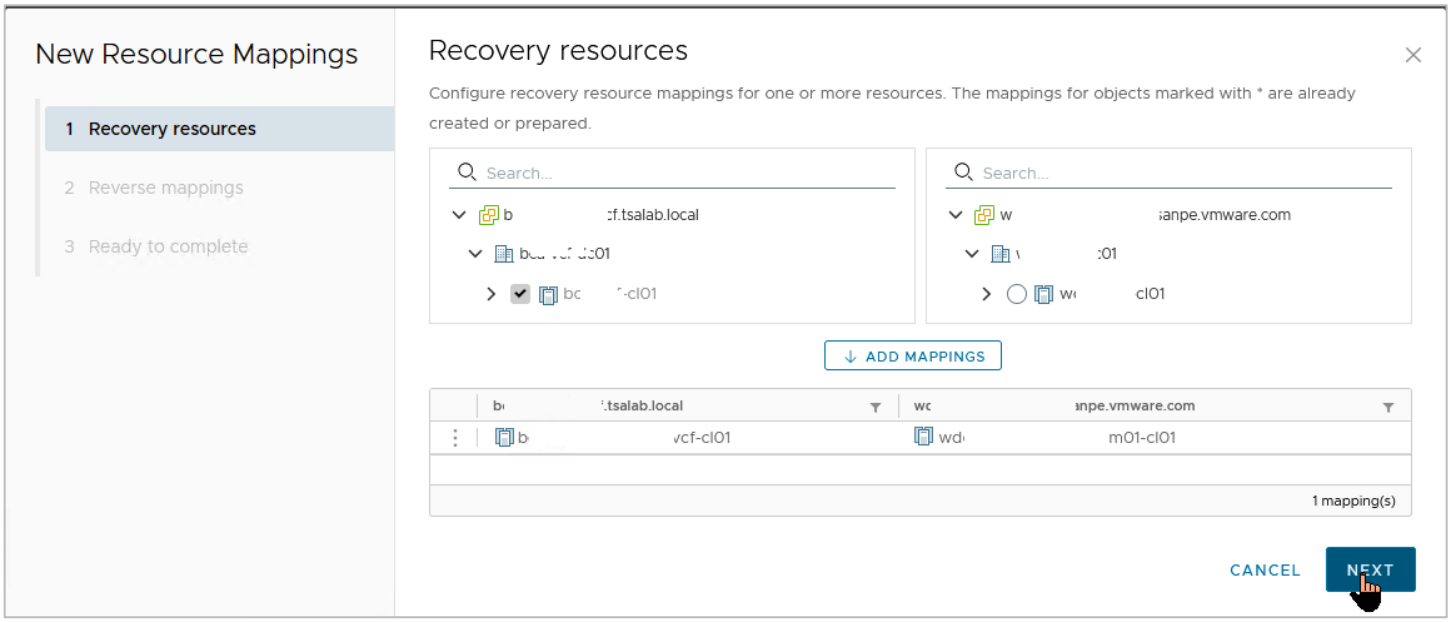

4. Accept the option to auto-configure a reverse mapping and click **Next**.

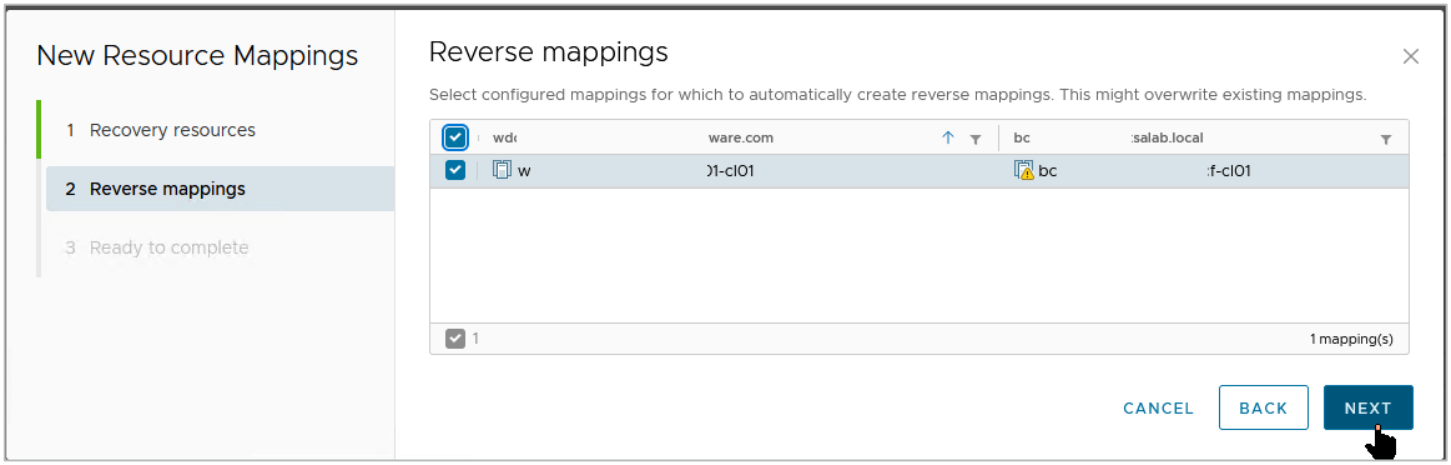

5. Click **Finish** to complete the process.

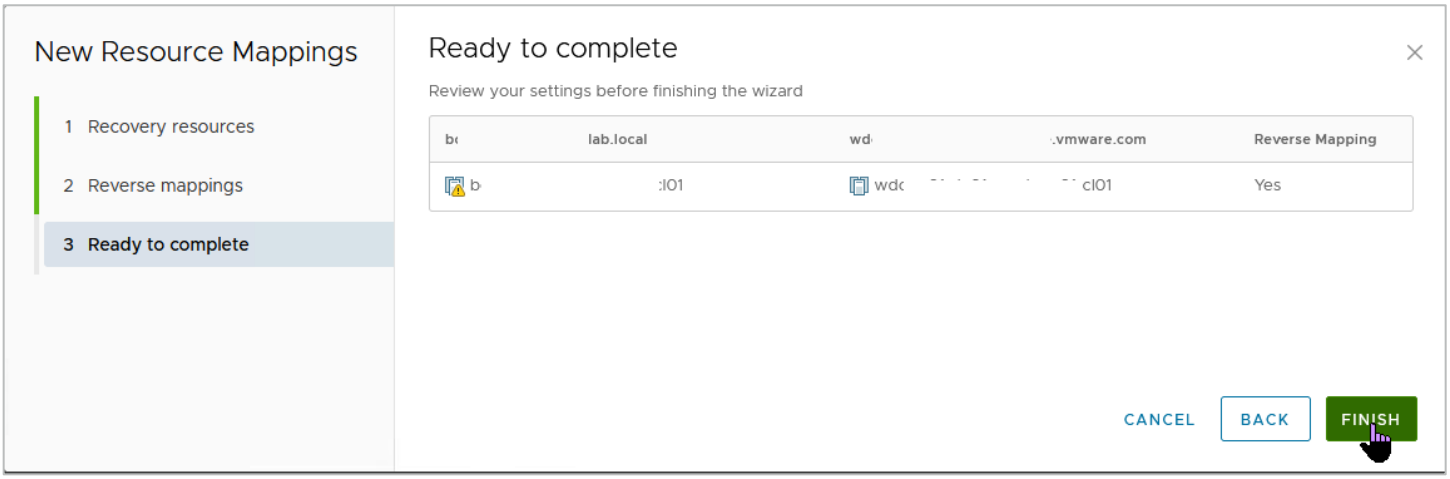

6. Click **Continue** to acknowledge and dismiss the warning about a possible impact on existing protected VMs.

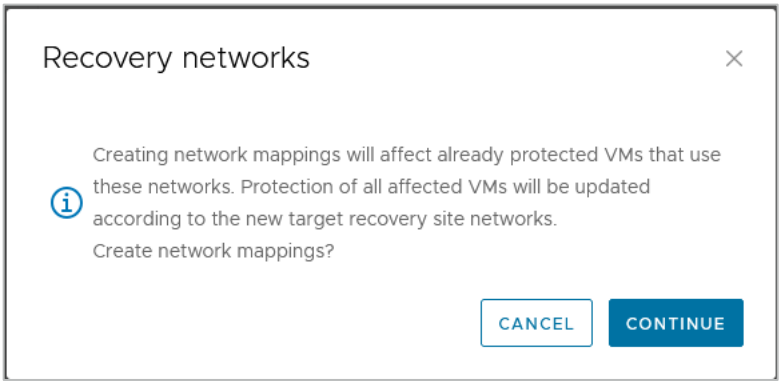

When VMware Live Site Recovery uses vSphere Replication to replicate a protected VM to the recovery site, it also creates a representation of the VM in the vCenter at the recovery site. This representation is somewhat similar to the .vmx file that describes the running VM at the protected site. The major difference is that this representation is just a placeholder (aka "stub"), which can't be powered on. This placeholder file is stored in a designated datastore, which might not necessarily be the datastore with the full replicated copy of the protected VM. The "placeholder" datastore must exist on both sides for VMware Live Site Recovery to protect workloads in either direction.

7. From the **Placeholder Datastores** menu, click **New**.

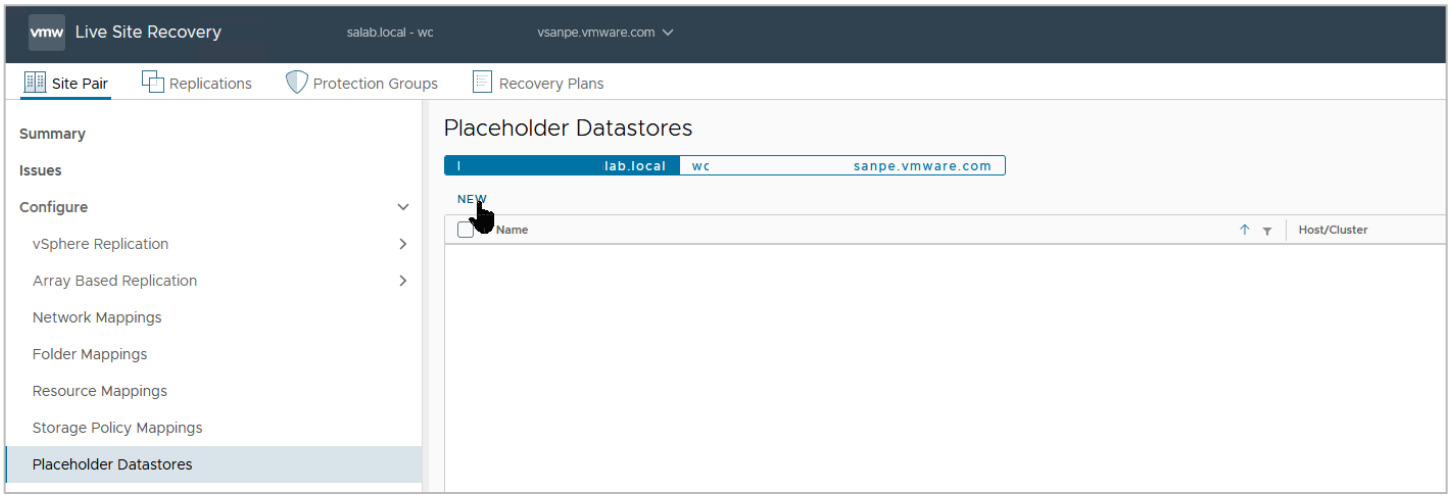

8. Select the datastore you want to use to store the placeholders and click **Add**.

Note: VMware Live Site Recovery requires the specified datastore to have a minimum of 6GB of free space.

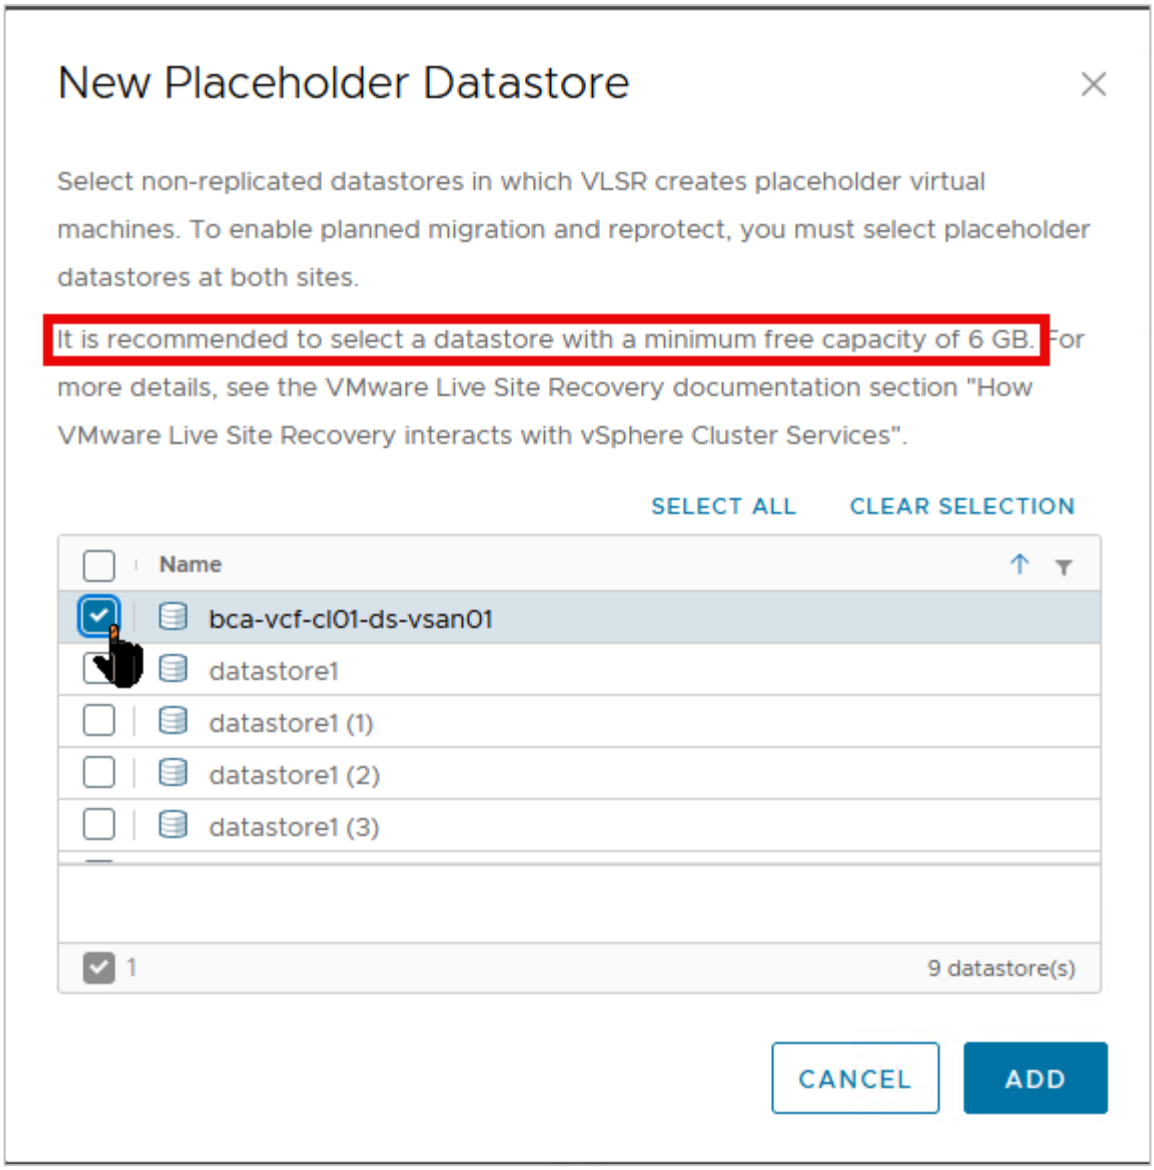

Here is the placeholder datastore at the protected site:

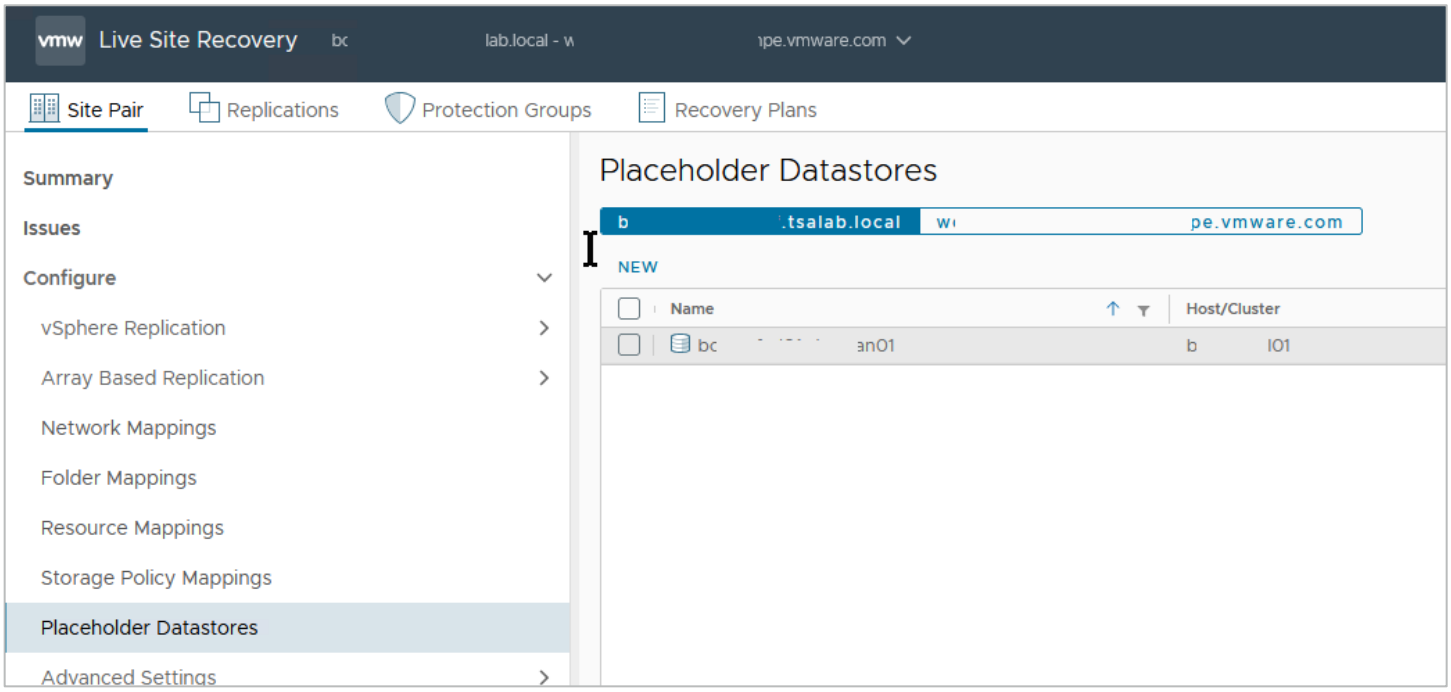

From this menu, you can click on the recovery site (see the arrow below) to specify the corresponding placeholder datastore for that site. Here's the placeholder datastore at the recovery site:

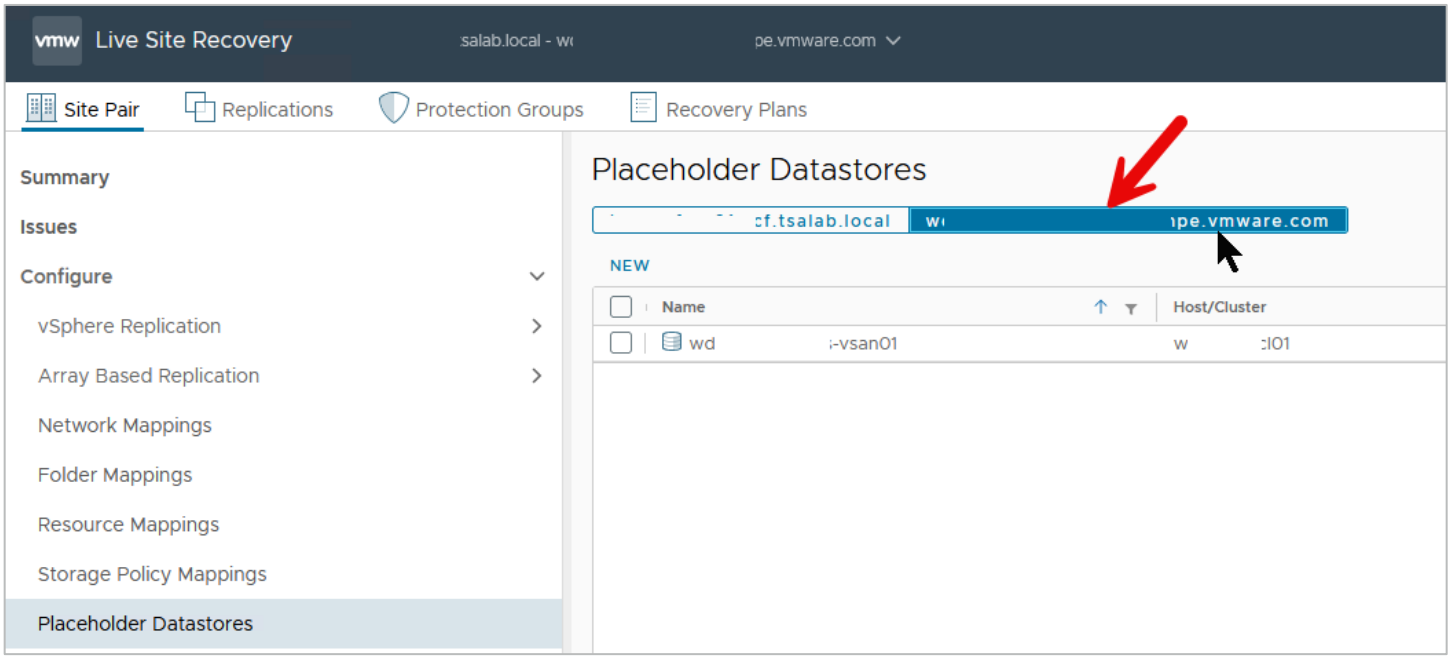

Replicate protected VMs

For VMware Live Site Recovery to protect and recover a VM, a copy of that VM must make its way from the protected site to the recovery site. Let's set up the replication part of the exercise now.

Create an outgoing replication

In this guide, the source (the protected site) is the on-premises VCF infrastructure, so let's switch to that and create an **Outgoing** replication.

1. Click **New**.

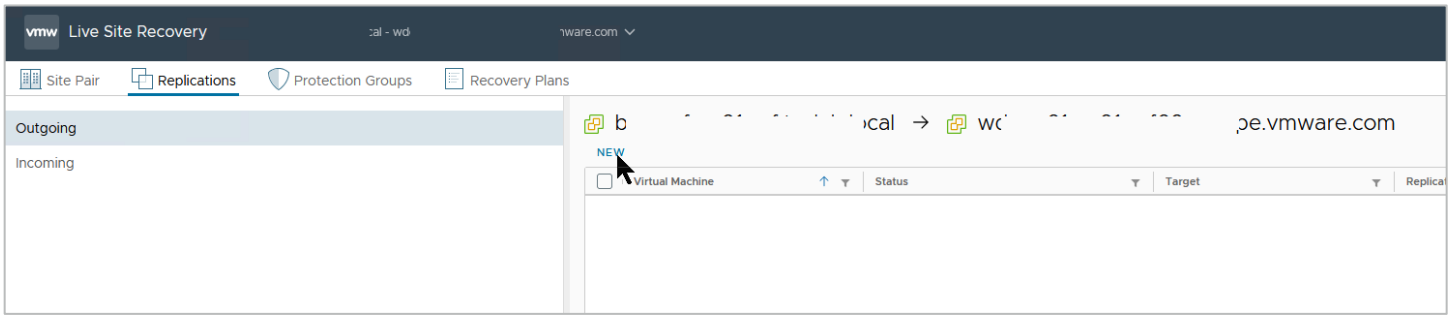

VMware Live Site Recovery introduces an enhanced replication feature, which enables automatic load-balancing of the replication engines to ensure high performance and more fine-grained replication schedules. With enhanced replication, VMware Live Site Recovery can provide up to a 1-minute recovery point objective (RPO), ensuring protected workloads are up to date.

2. If your storage arrays and infrastructure support such high replication frequencies, choose **Enhanced Replication**. Otherwise, select the **Legacy Replication** option and click **Next**.

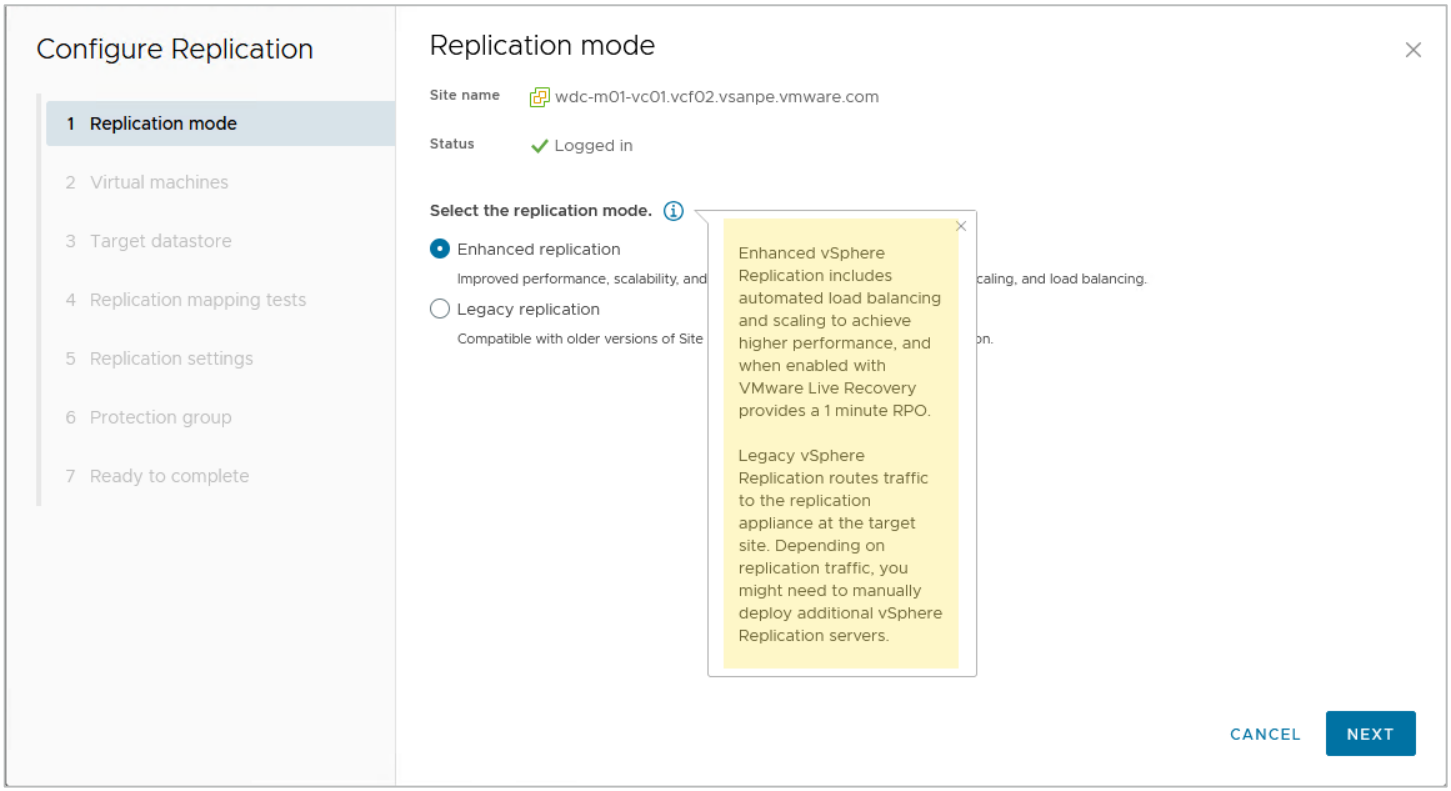

Note: If you selected the **Legacy Replication** option, and your replication engine is the VMware vSphere Replication server (or servers), you'll be prompted to select the applicable replication server or accept an automatic selection.

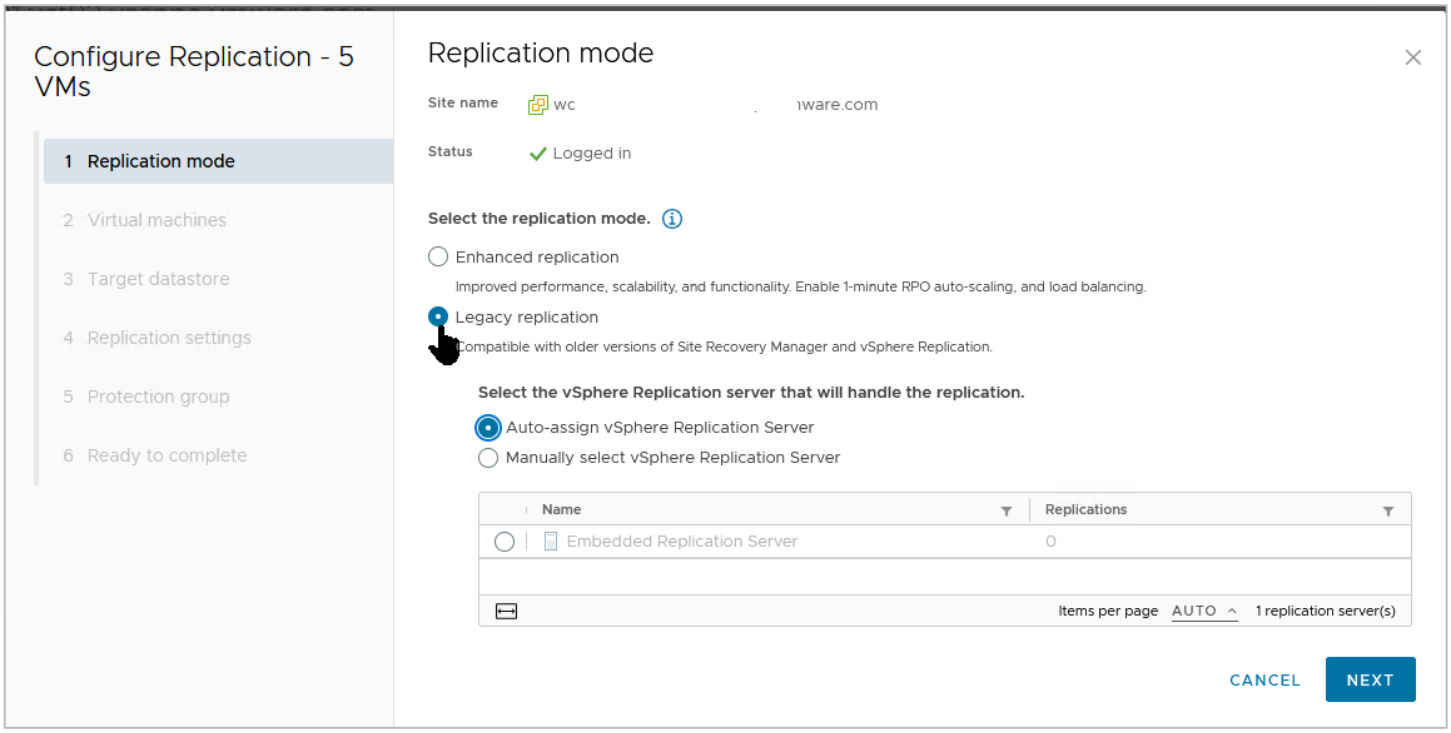

3. Select the VMs you want to protect with VMware Live Site Recovery and click **Next**.

Note: You can add or remove VMs from replication (and, consequently, protection) at will, so just select a few for now.

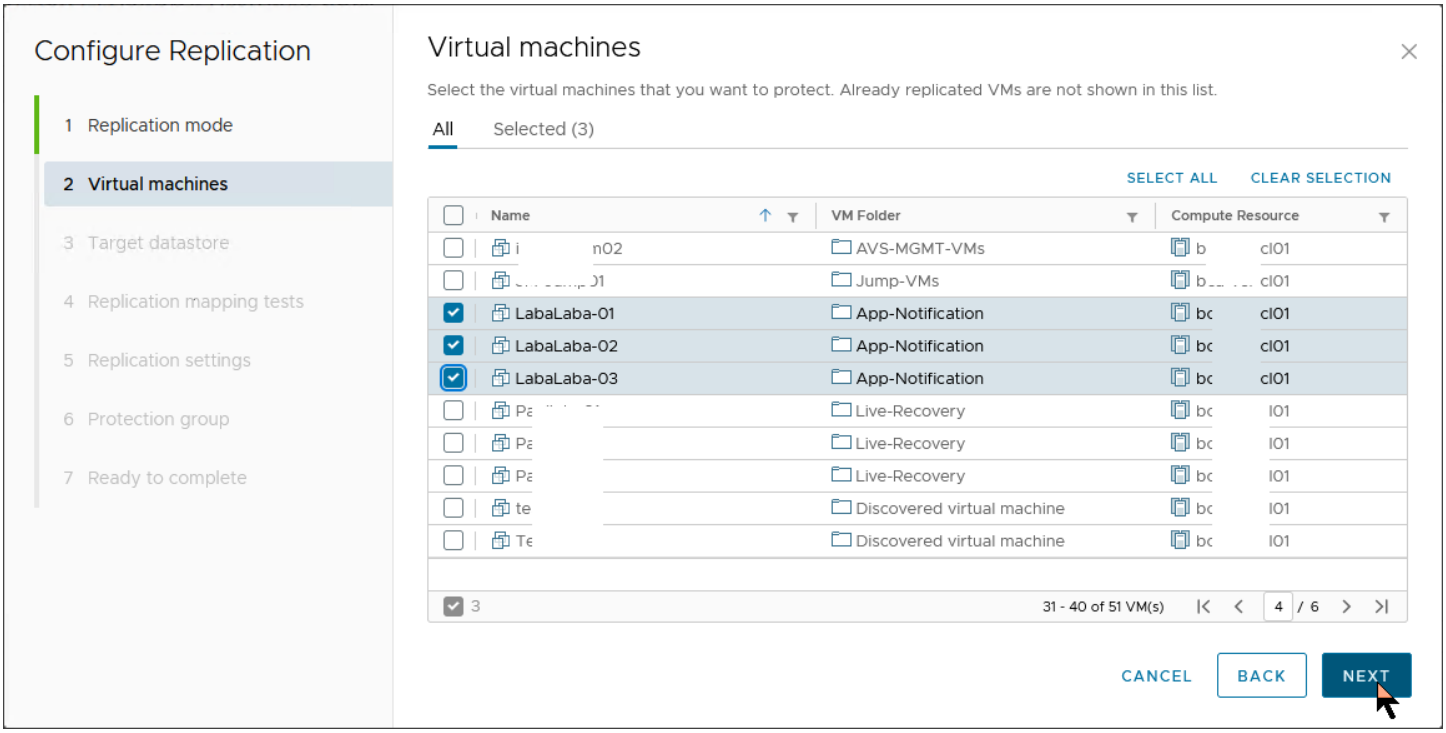

4. Select the datastore where you want to store the replicated VM copies at the target (recovery site).

Note: The target/recovery site is another VCF infrastructure that the process auto-identifies because it's already paired.

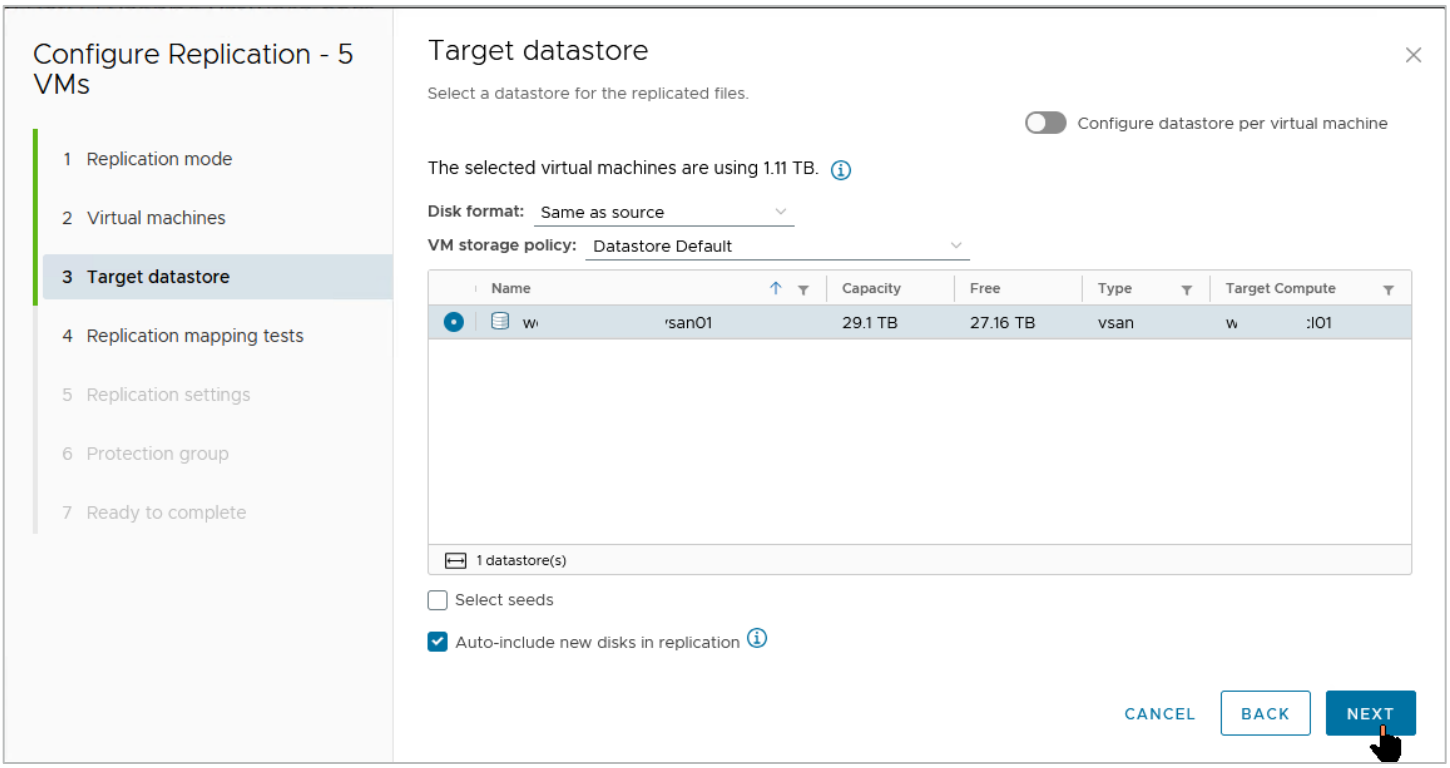

The option **Auto-include new disks in replication** is one of the amazing things about vSphere Replication. It anticipates situations where a protected VM's configuration could change after we set up the disaster recovery plans. With this option, vSphere Replication automatically incorporates the changes into the replication tasks.

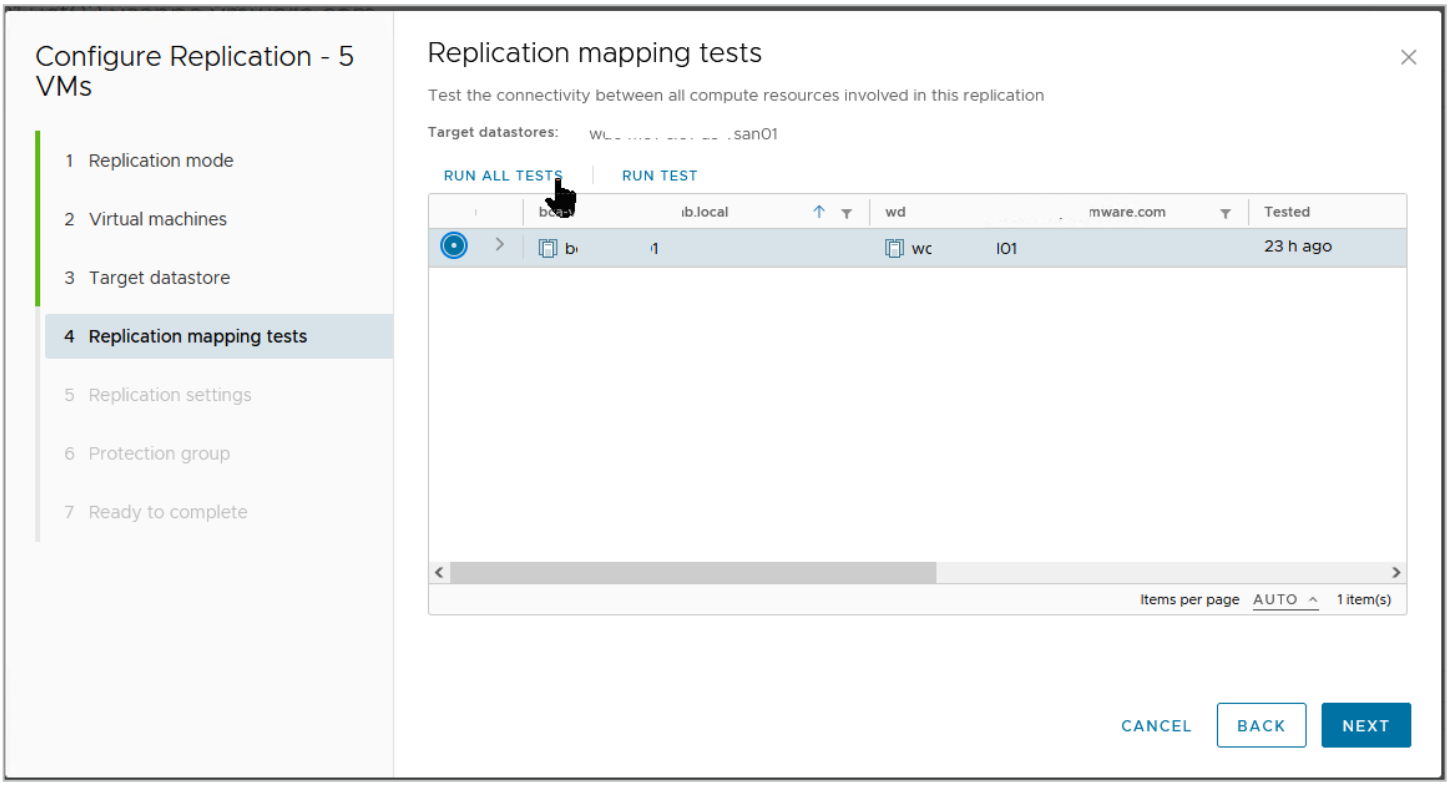

Remember the mapping tasks you performed a while back? Here's where they begin to come into play. VMware Live Site Recovery will now perform the necessary steps required to verify the connectivity and availability of the mappings.

5. If this is the first time you're doing this, select **Run All Tests** and click **Next**. If you've previously validated these, skip this or run a specific test.

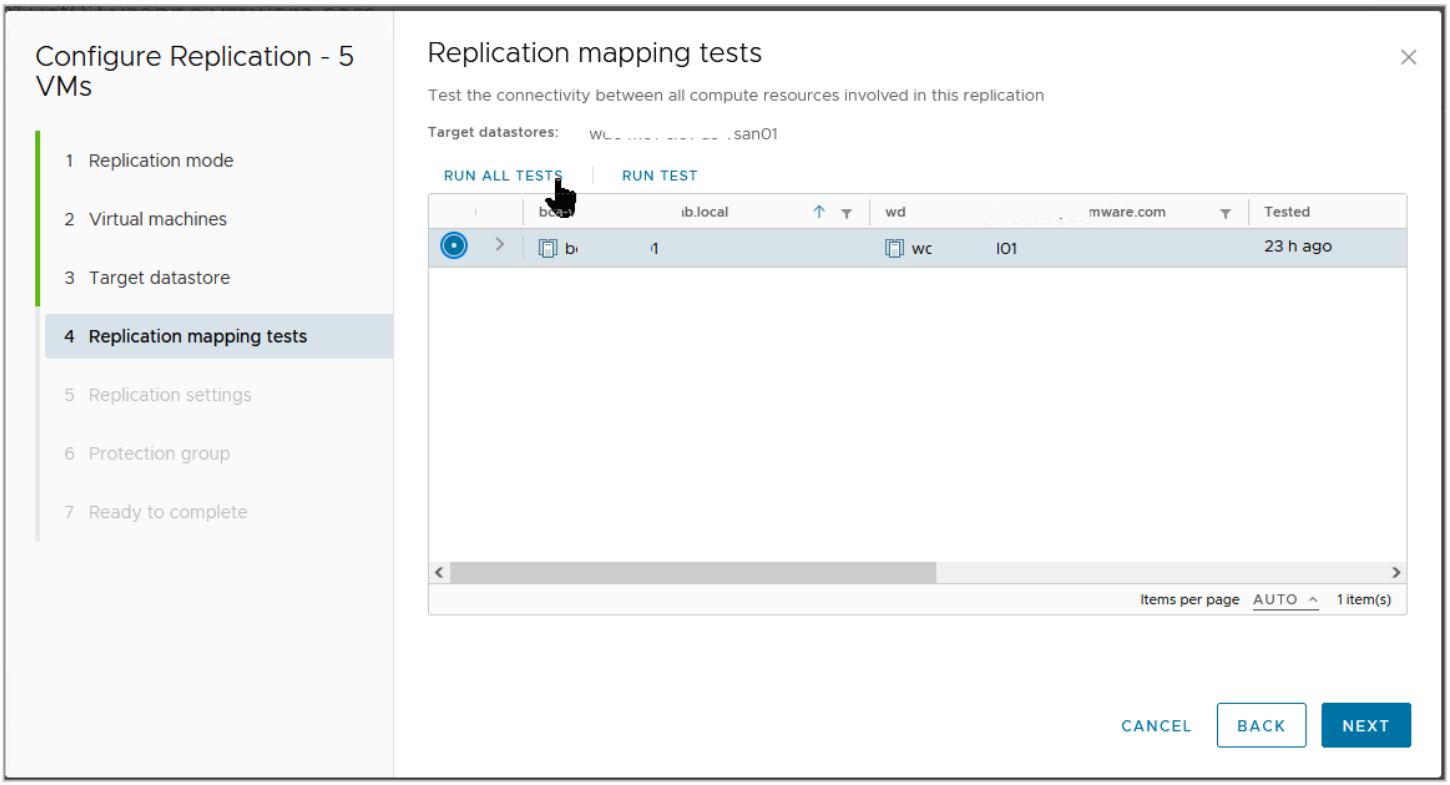

6. Click **Run Tests** on the pop-up window.

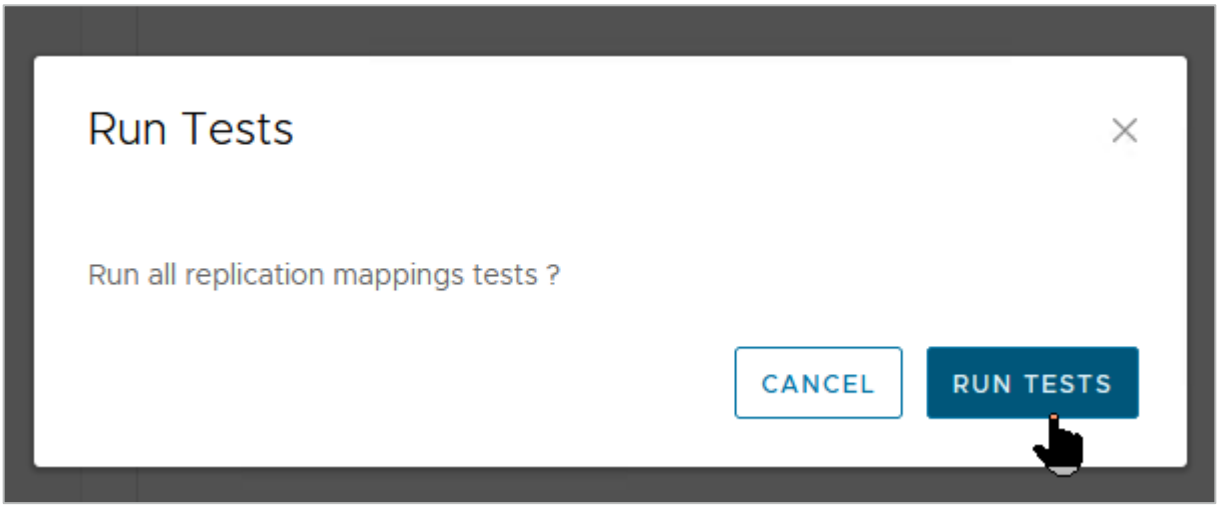

7. If all of the results come back good, click **Next**. Otherwise, review and fix any reported errors before proceeding.

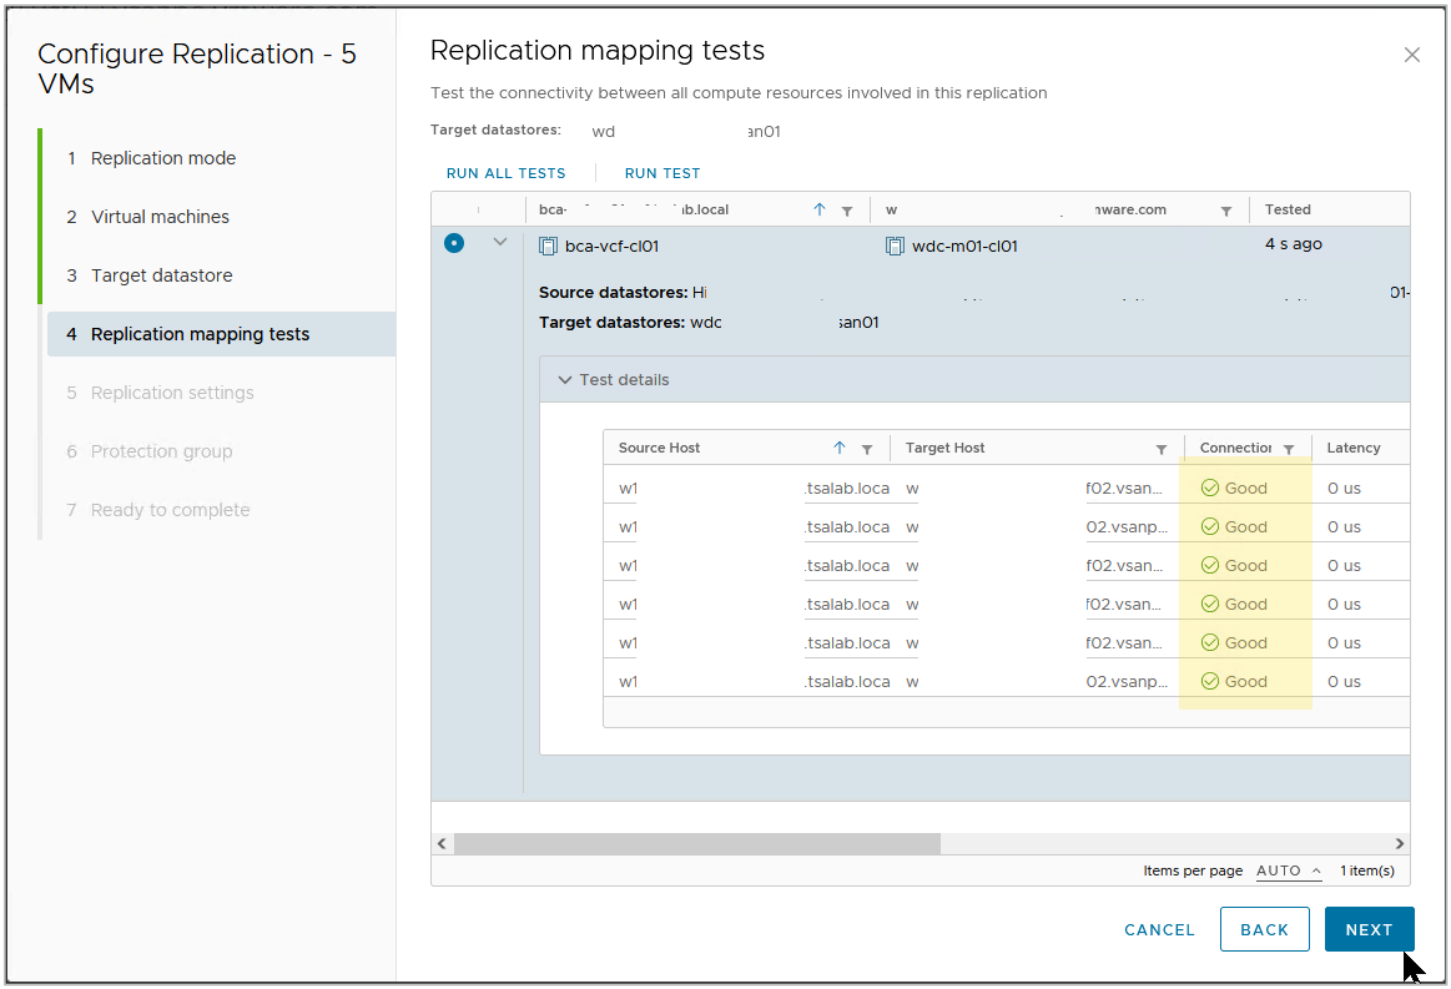

RPO/RTO, run book, protection group, and recovery plan defined

RPO and RTO might be two of the most overused acronyms when discussing disaster recovery of mission-critical applications in the enterprise. We've consciously avoided mentioning them until now because they deserve book-length attention, which we can't accommodate in this guide. Simply put:

- A **recovery time objective (RTO)** specifies the acceptable duration required for an IT infrastructure to recover from a disaster event and resume operations. The objective is to make this window as short as possible. Several external, environmental, and infrastructural factors influence an RTO, so we won't demonstrate this concept here.
- A **recovery point objective (RPO)** specifies the acceptable loss of services or data in a disaster event. It measures how up to date the enterprise data is after such an event. Admins, operators, and business owners/stakeholders want an RPO of 0, but financial, human, and technological constraints make this difficult to attain.

vSphere Replication provides RPO options from as low as 1 minute and as high as 24 hours, ensuring greater flexibility in recovery targets and allowing you to fine-tune your DR plans based on your infrastructural constraints. This means vSphere Replication attempts to synchronize and replicate every state change in the protected VM site as frequently as every minute. At any point, the copy of the VMs at the recovery site is identical to the original protected VM no more than 1 minute ago. On the extreme end of the spectrum, vSphere Replication can maintain a 24-hour RPO.

Consult your storage vendor for official guidance for RPOs supported by your non-vSAN arrays.

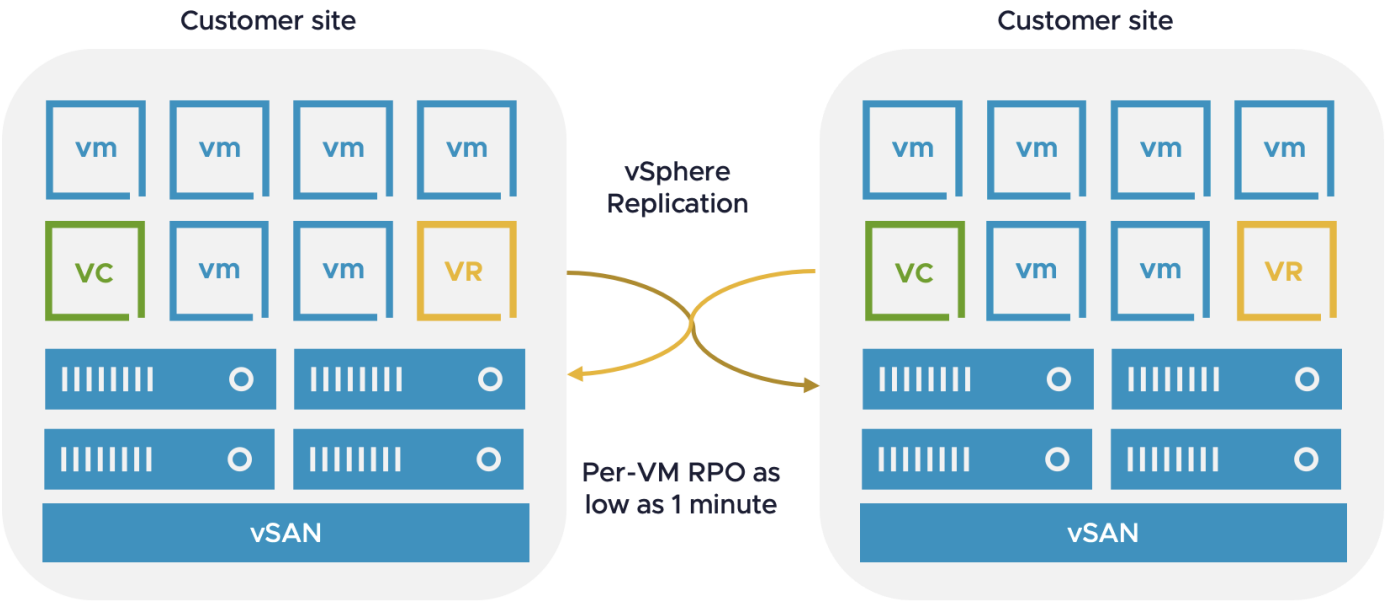

The vSphere Replication Admin Guide provides comprehensive documentation for the other capabilities and features on this screen, so we won't duplicate that information here.

8. For this exercise, select a 15-minute replication frequency and click **Next.**

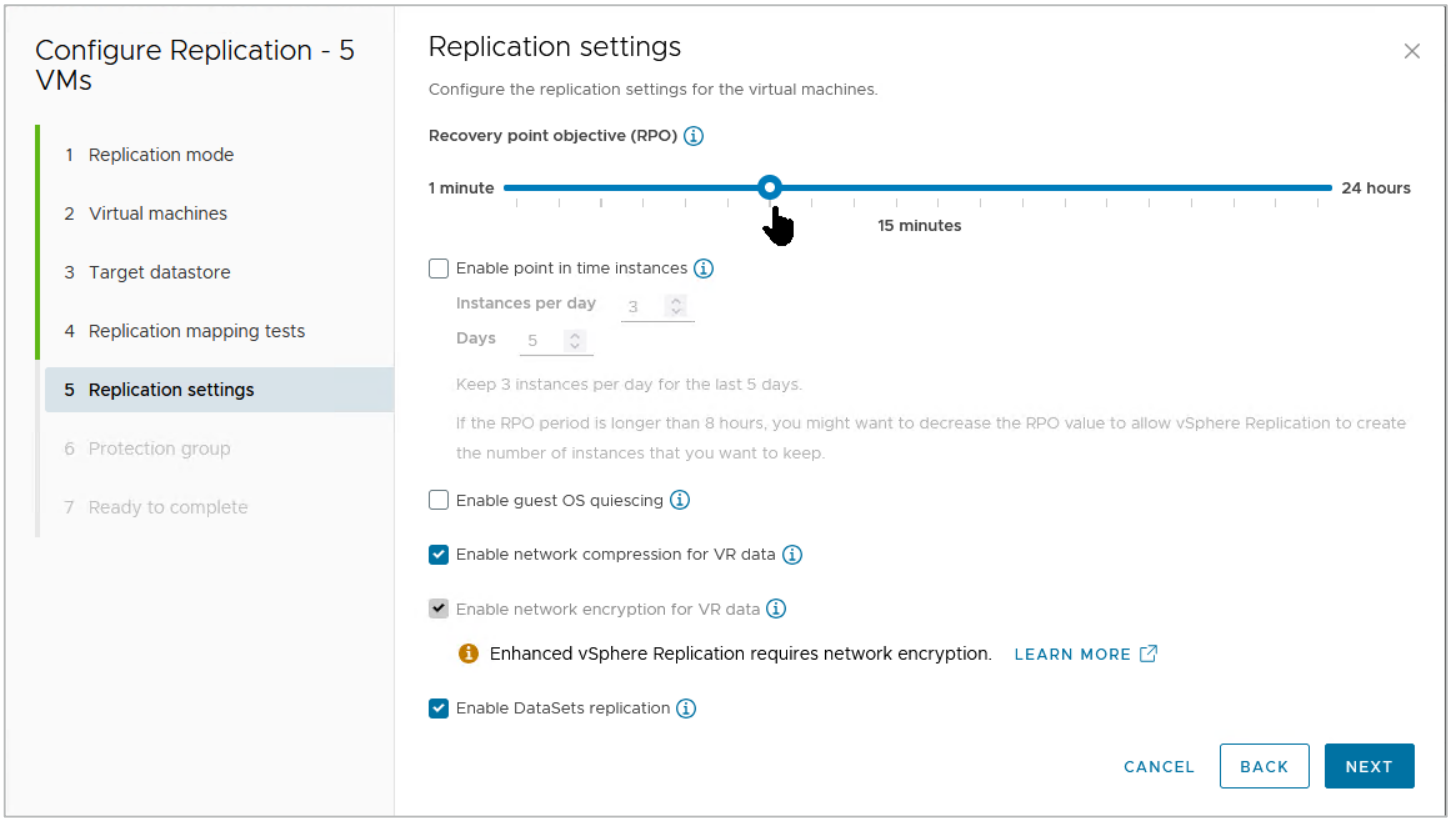

We'll talk about **Protection group** and other configurations later.

9. Select **Do not add to protection group now** and click **Next**.

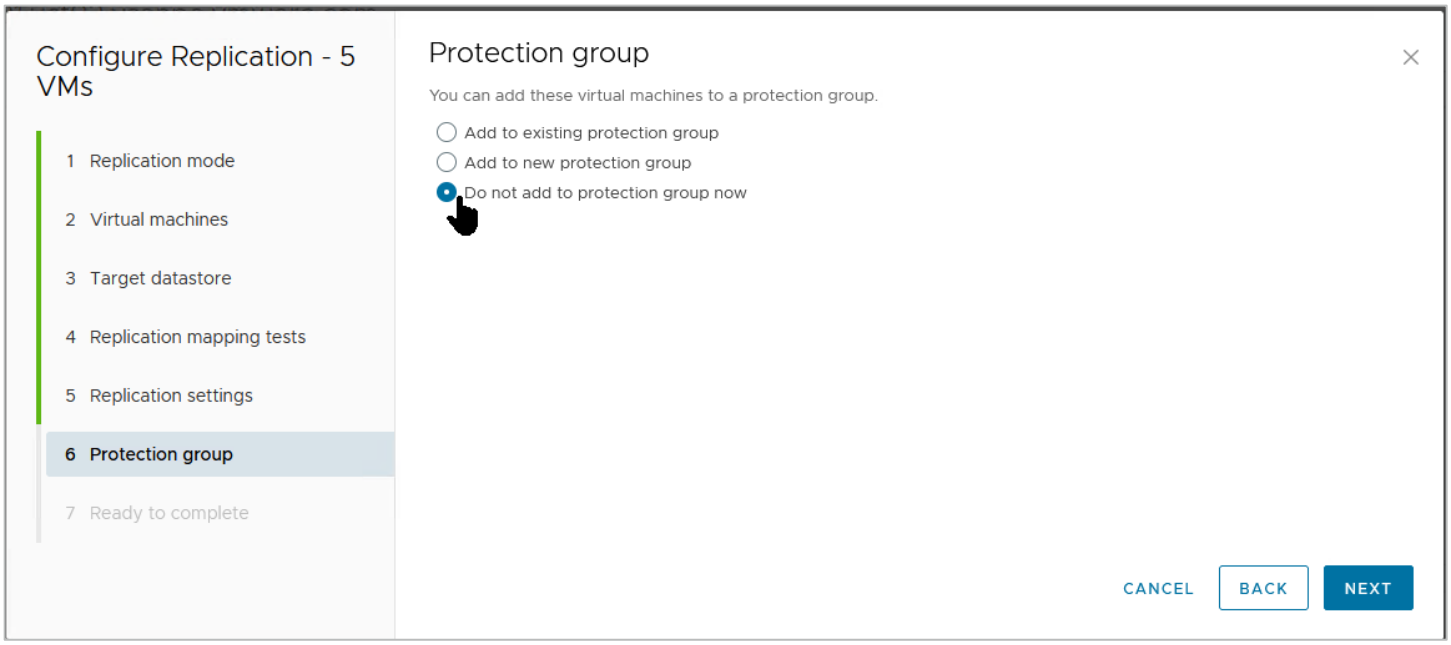

10. Review and verify the configuration settings and click **Finish** if it looks good, or click **Back** to make some modifications.

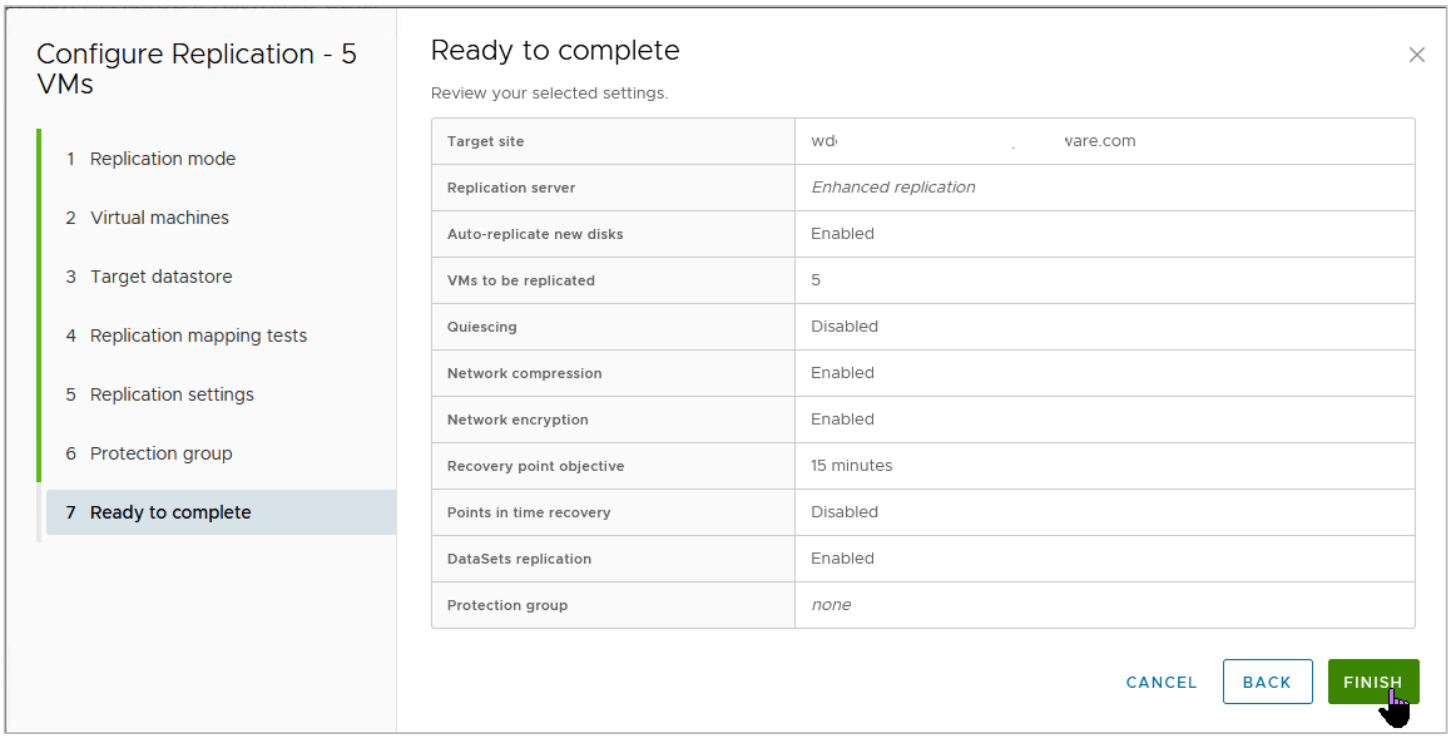

You’re finished with the replication setup.

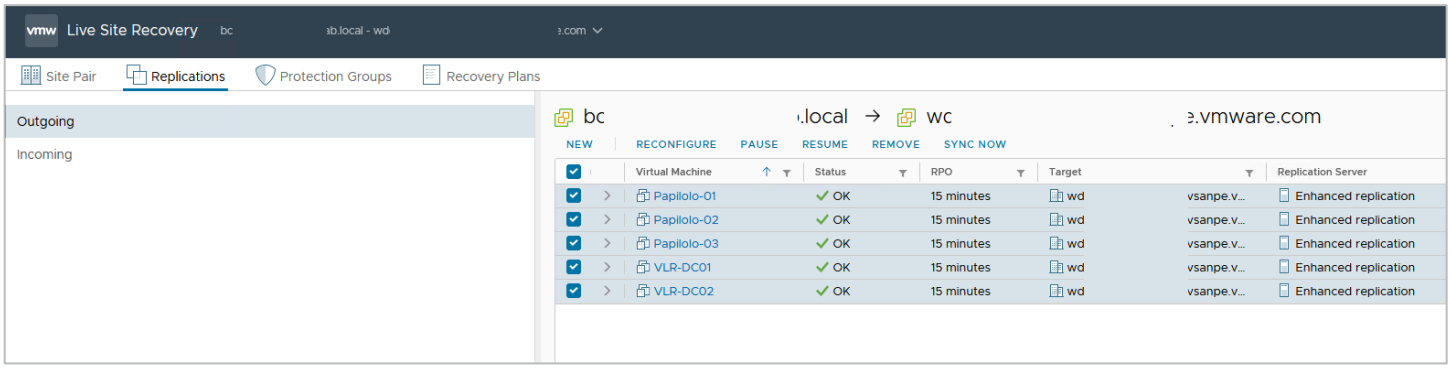

Create protection groups and recovery plan

Critical enterprise applications typically don't exist or function independently. They depend on other services and workloads, and others depend on them. When designing a BCDR plan, these dependencies influence configuration and workflow choices and options. VMware Live Site Recovery provides a feature called protection groups, which contain VMs you want to recover together as a unit. Many factors, such as the type of storage and the unit of replication, influence the decision-making processes involved in creating and using protection groups.

For this exercise, you'll create protection groups based on the services and characteristics of the VMs we want to protect and recover. There are three categories: Domain Controllers, SQL Server, and a Windows client. This is the primary influence on your configuration choice.

Create a protection group for the Domain Controller VMs

1. Go to the **Protection Groups** tab and select **New Protection Group**.

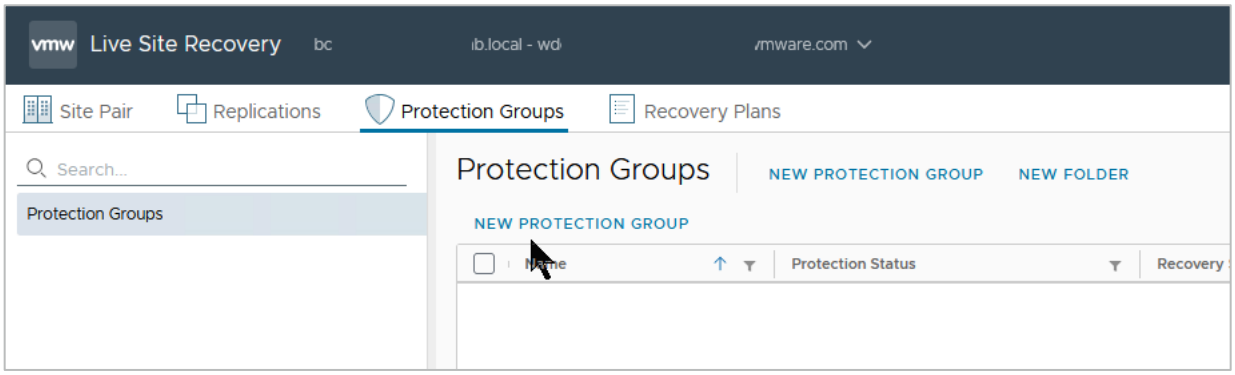

2. Give the group a descriptive name. The description is optional.

Next to **Direction**, select the option that shows your protected site → recovery site.

3. Click **Next**.

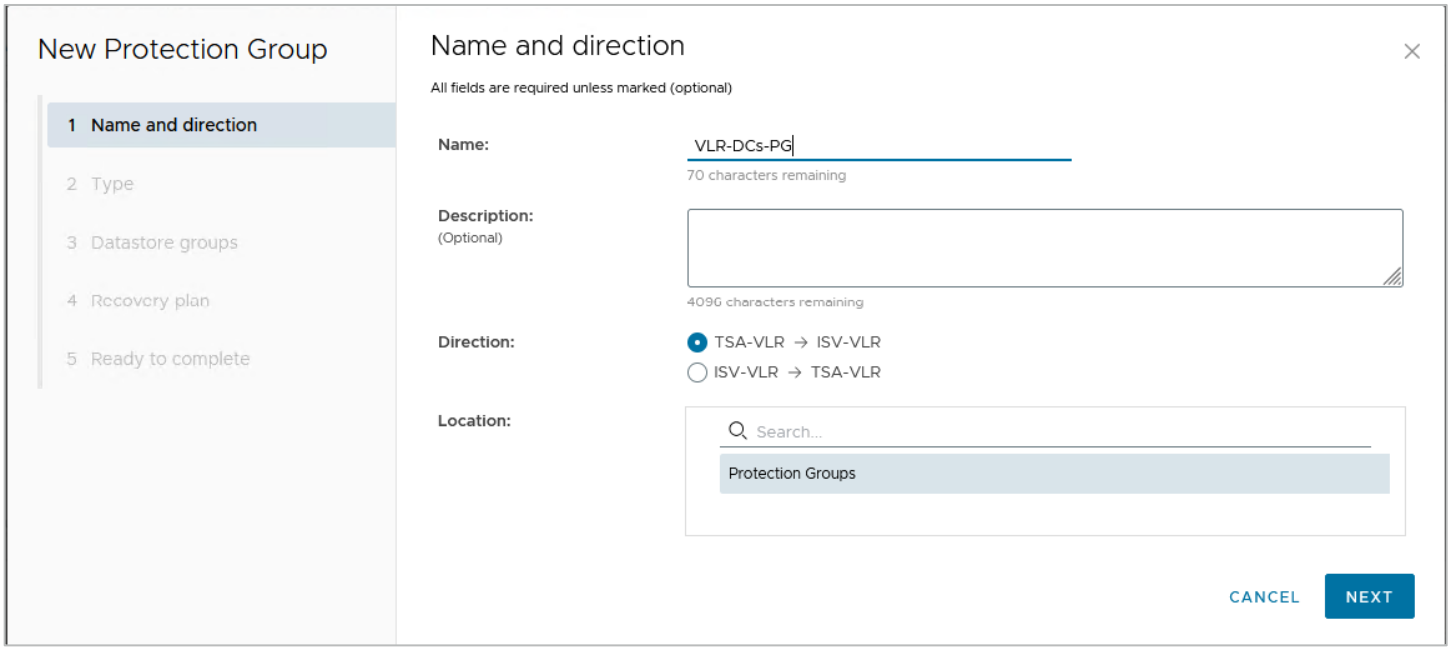

4. Because you're using vSphere Replication, select **Individual VMs (vSphere Replication)** and click **Next**.

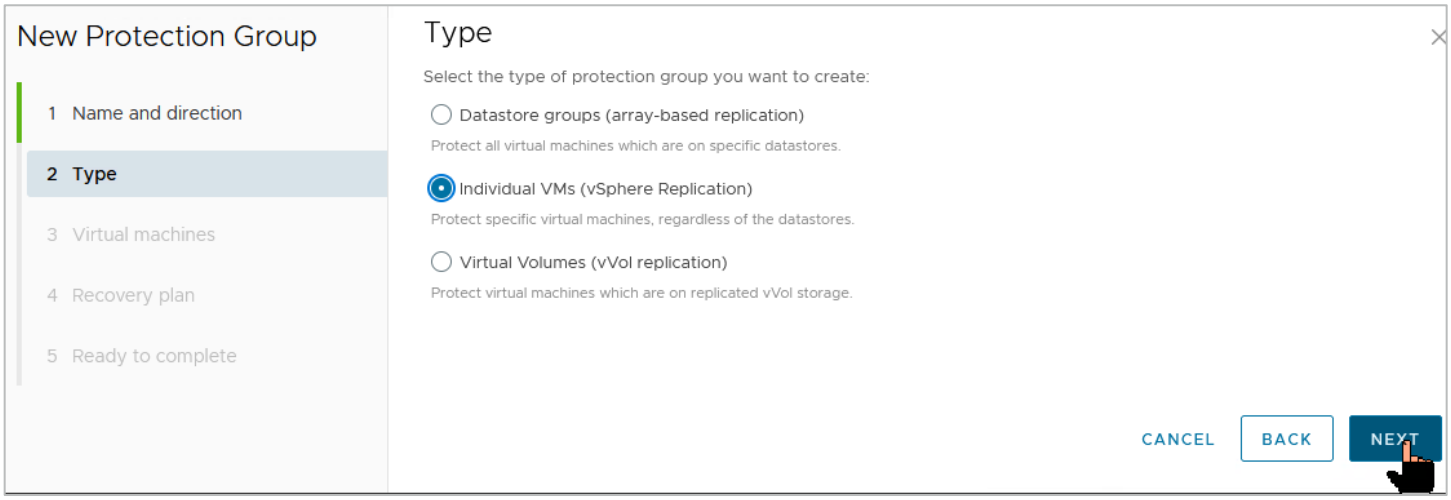

5. This protection group is for the Domain Controllers, so select those VMs and click **Next**.

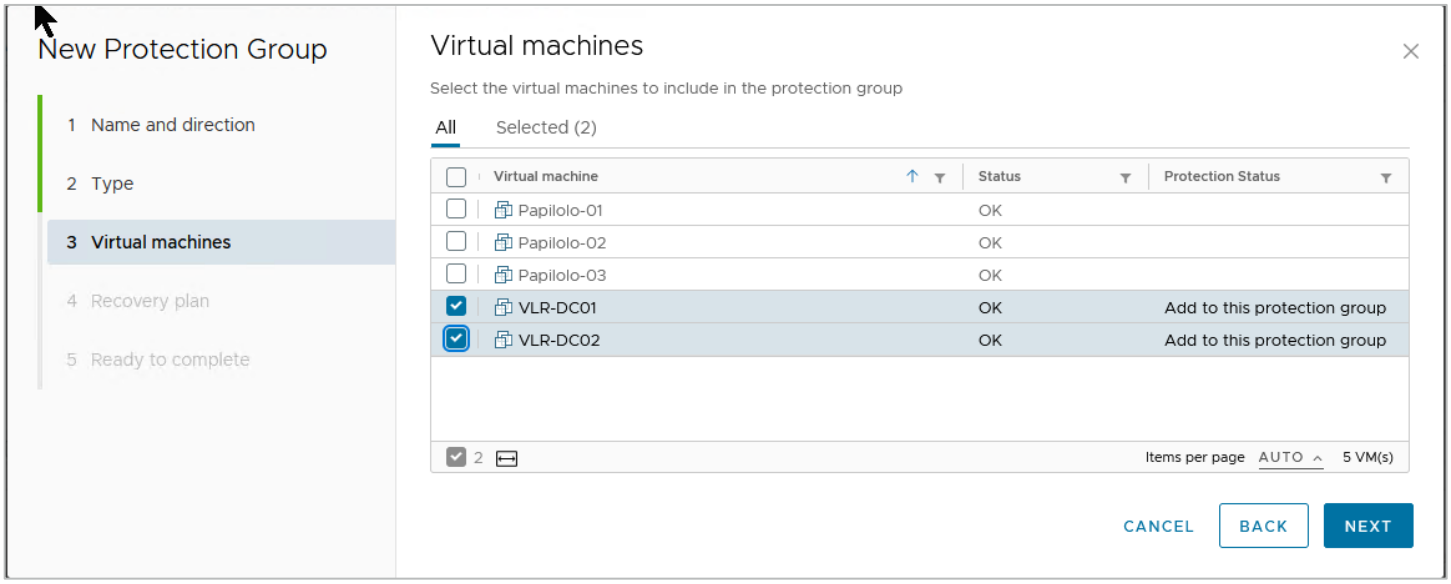

Create a recovery plan for the Domain Controller VMs

A recovery plan defines and configures the steps, plans, and actions guiding your BCDR plan. Imagine it as the run book an admin would typically refer to and follow if they were to perform a disaster recovery operation manually.

In VMware Live Site Recovery, a recovery plan contains the logic and workflow of getting the copy of the protected VMs up and running in the recovery site when a disaster is declared and the recovery is initiated. You must add at least one recovery plan to the protection group. Because you don't already have a recovery plan, you'll create one now.

6. Select **Add to a new recovery plan.**

7. Next to **Recovery plan name**, type an intuitive and descriptive name and click **Next**.

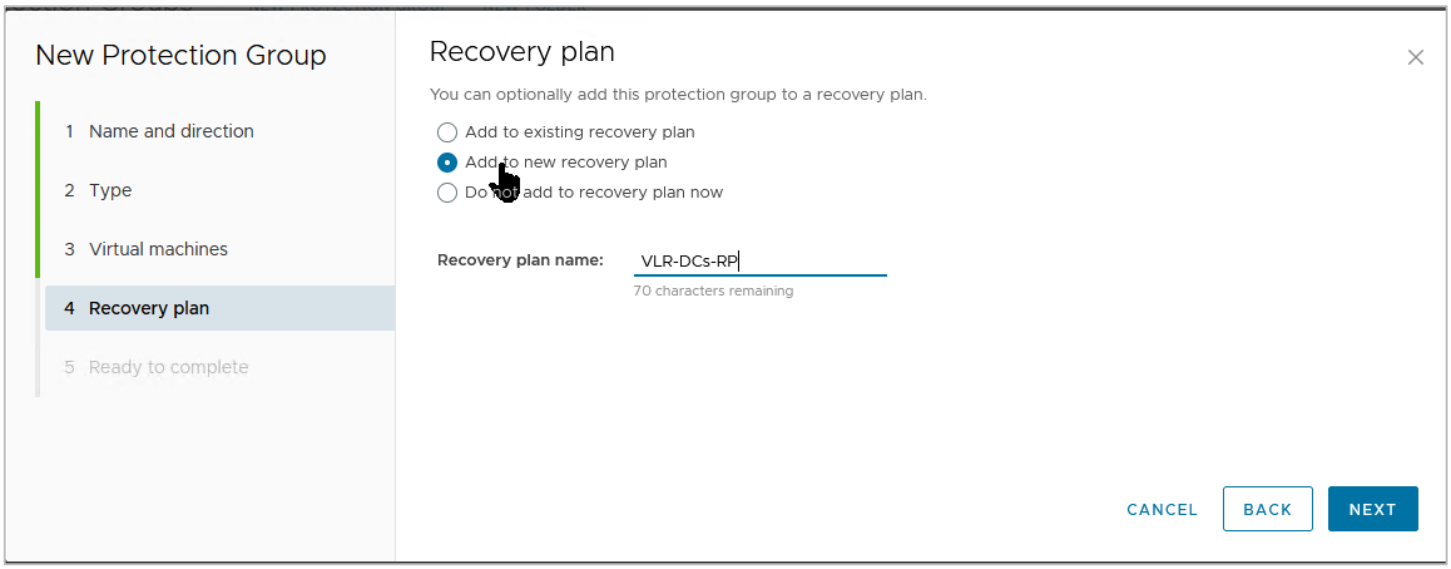

8. Review the result and click **Finish**.

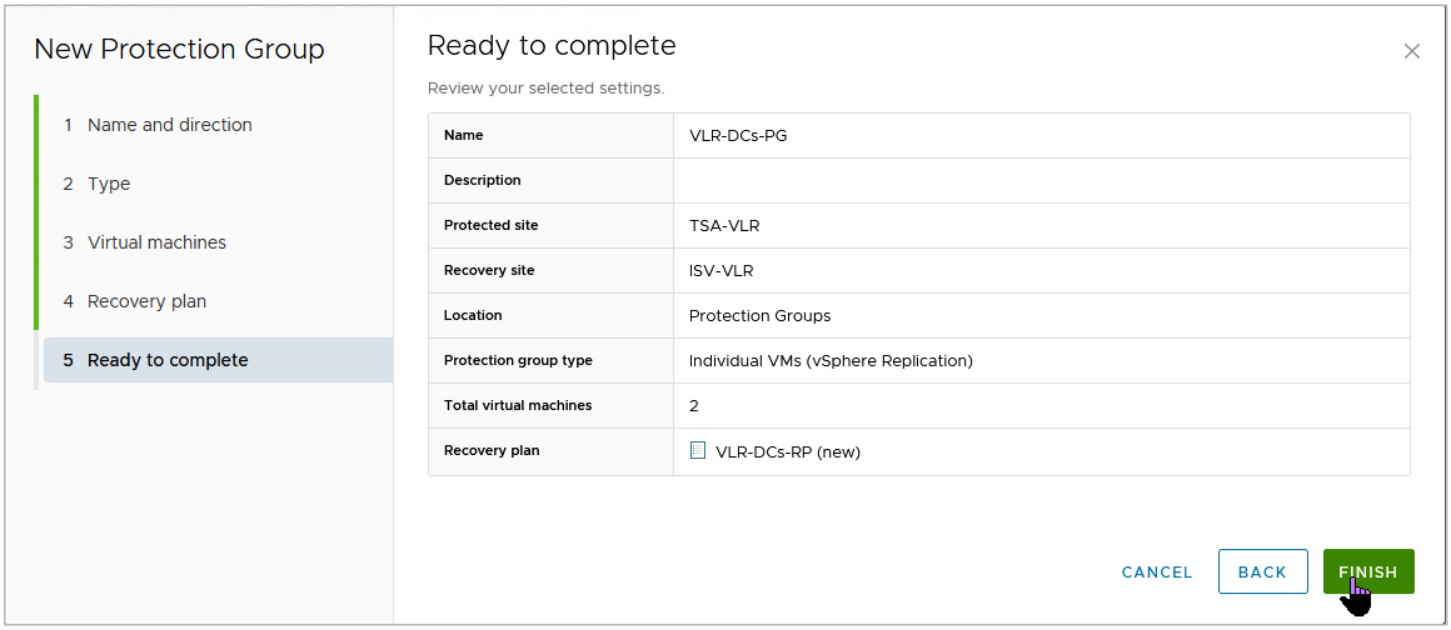

You're done creating the protection group and recovery plan for the Domain Controllers.

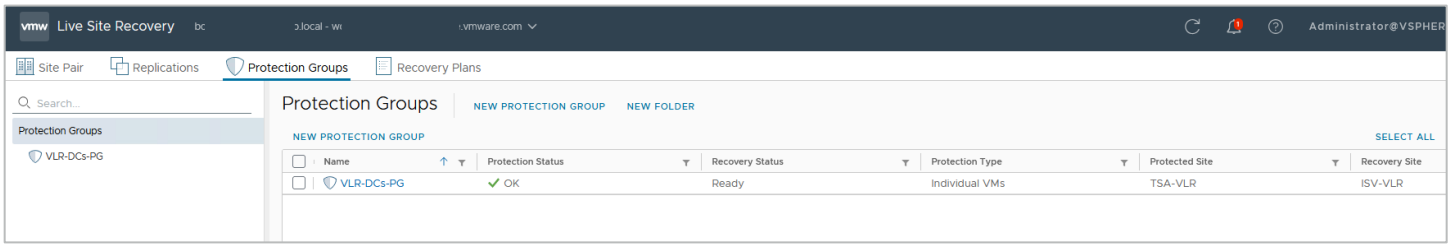

Create a protection group for the SQL Server VMs

Go ahead and create another protection group and recovery plan, this time for the SQL Server VMs.

Note: Because the process is similar, we won't describe every step here.

9. Give the protection group a descriptive name.

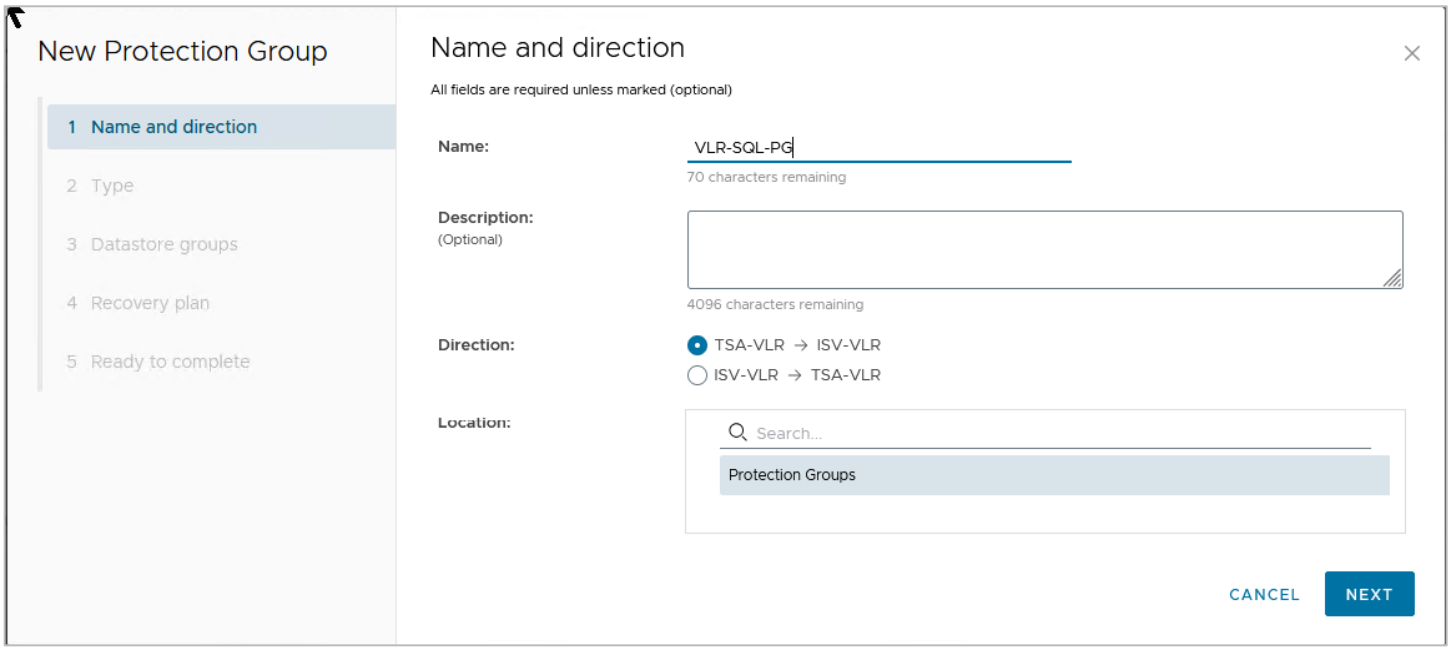

10. Select **Individual VMs (vSphere Replication)**.

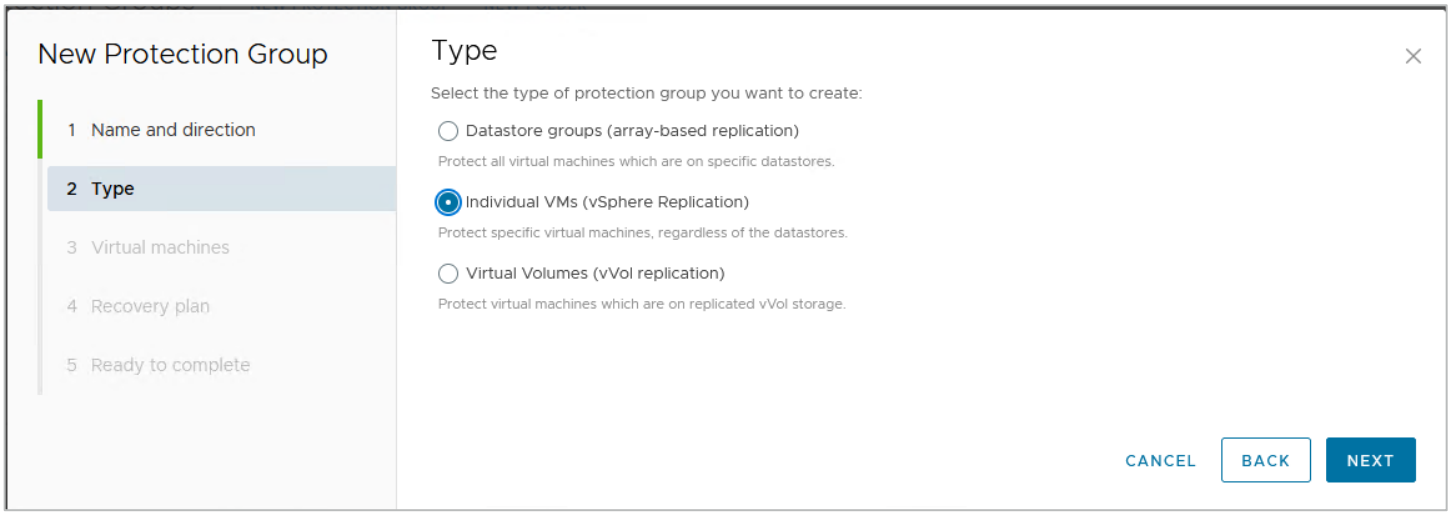

11. Select the SQL Server VMs.

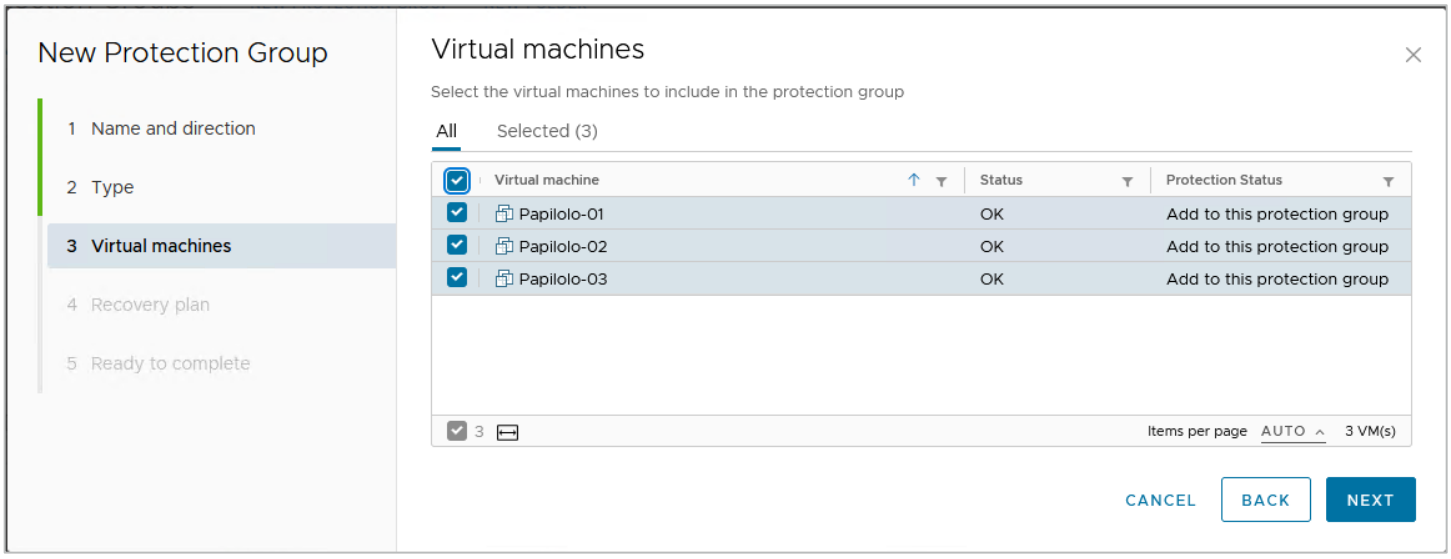

Create a protection group for the SQL Server VMs

12. Create a corresponding **Recovery Plan** for it.

That's it.

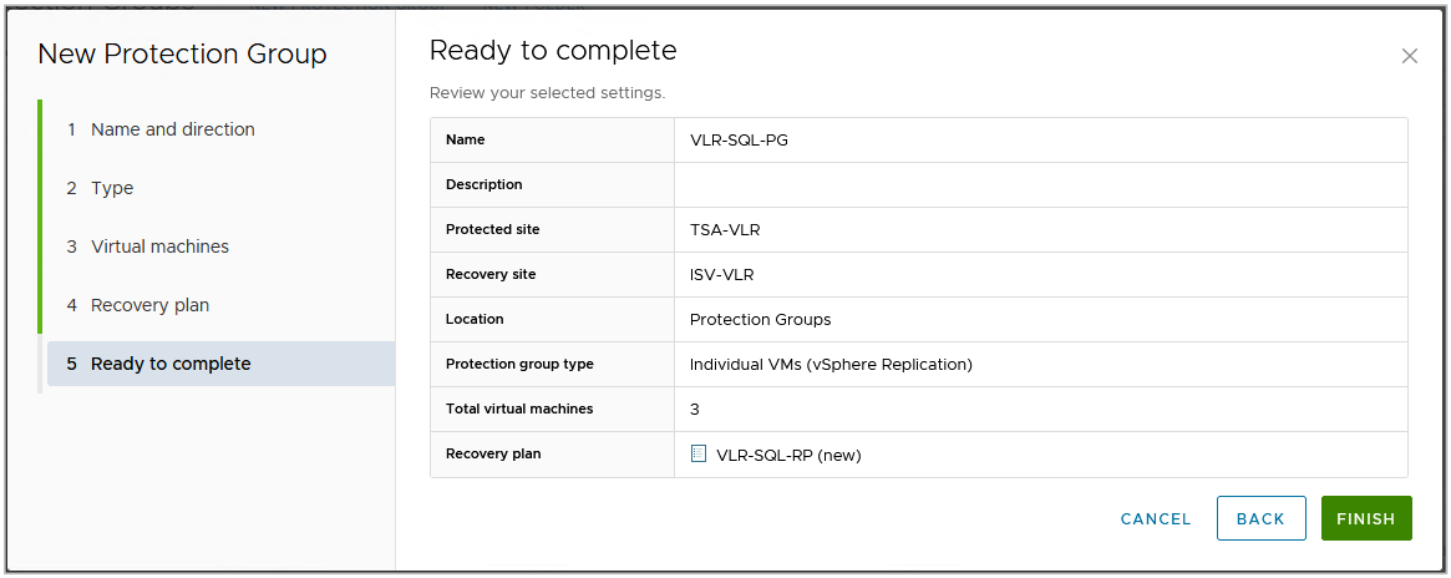

Here's what the protection group list looks like now:

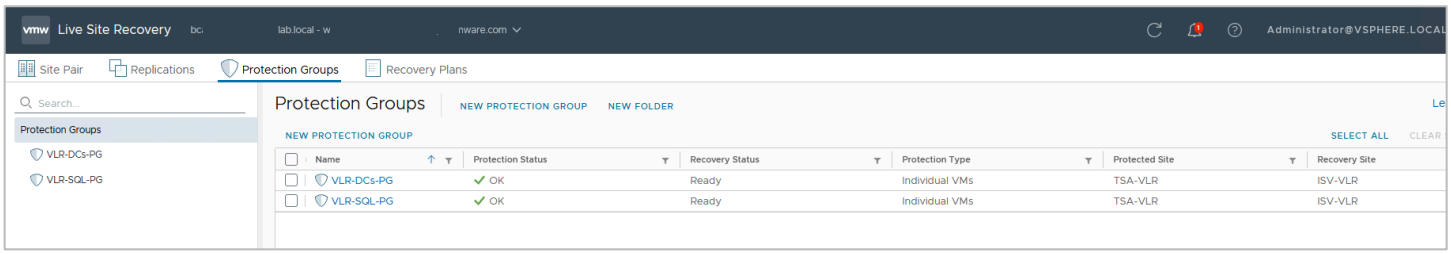

Modify the recovery plan

One of the most common tasks of a recovery plan is to configure the specific test ("bubble") network you'd like to recover VMs into during a test recovery exercise.

We previously discussed the importance and use of a bubble network. We also configured a site-wide bubble network for all test recovery operations in previous steps.

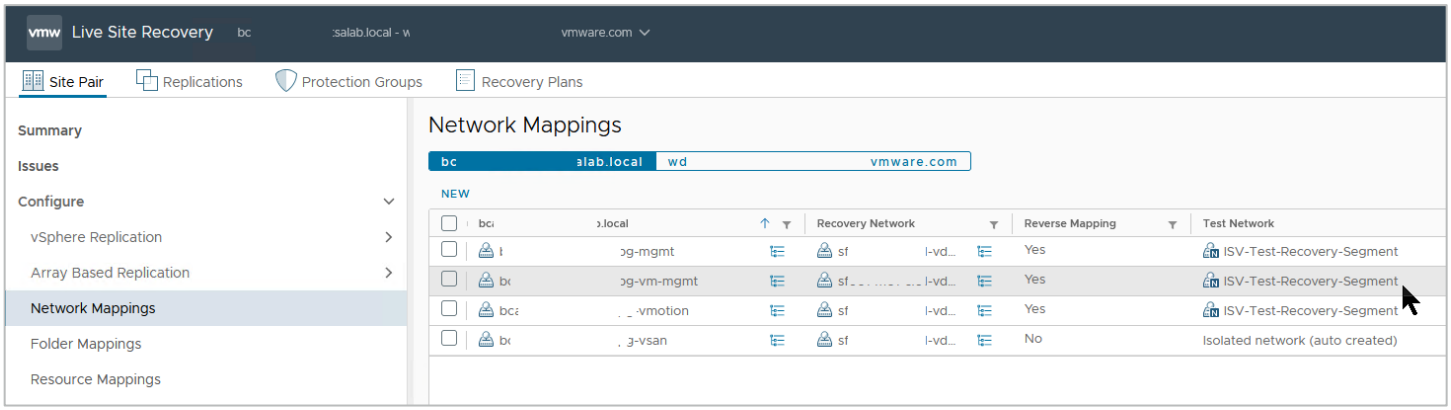

However, the flexibility of VMware Live Site Recovery also gives us the ability to configure the test recovery network at the recovery plan level. This way, you can specify different test networks for different classes of protected workloads.

You'll modify the recovery plan to see how you can control the fencing required for test recovery operations.

1. Select the recovery plan you want to modify and click **Edit**.

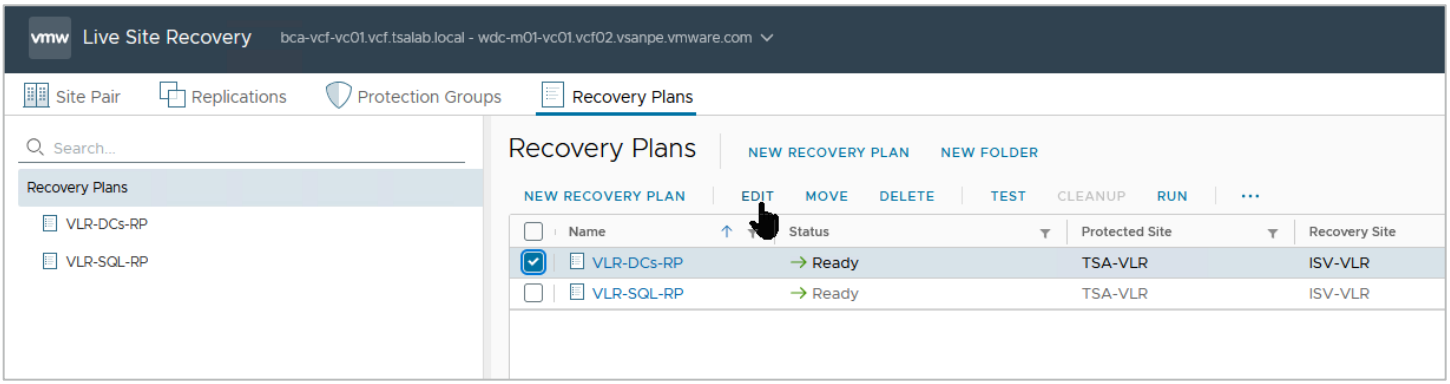

2. You'll only modify the test network settings, so click **Next** on the next two screens.

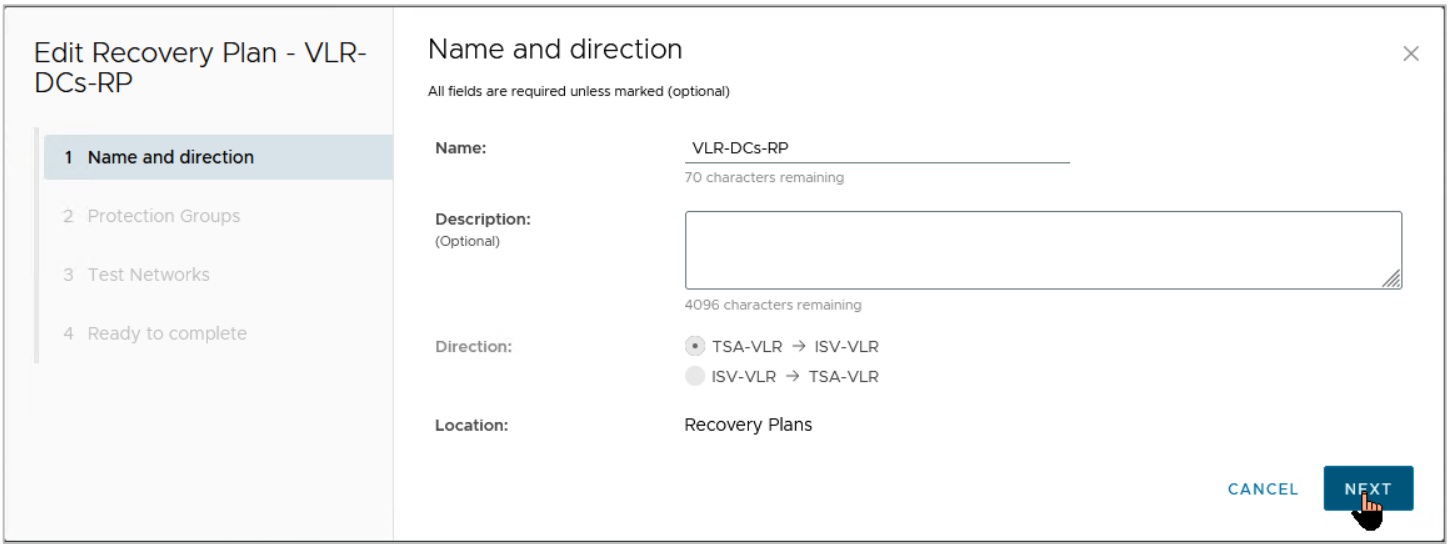

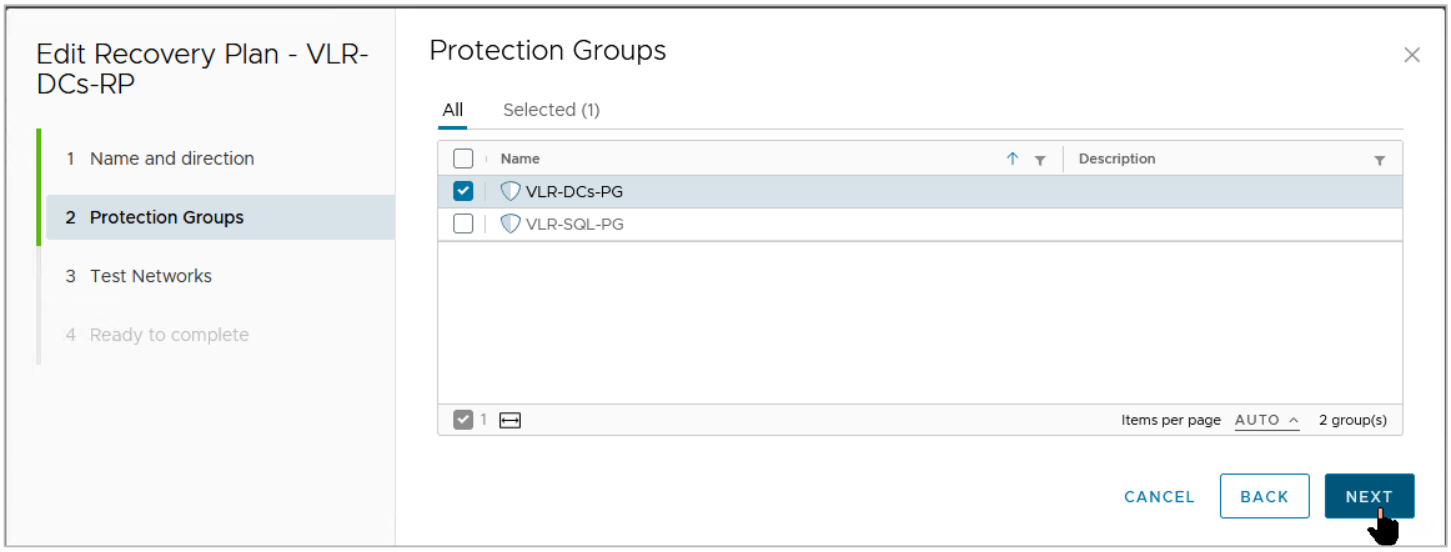

Notice how the test network options default to **Use site-level mapping**? You'll change that now.

3. Click **Change** on each of them.

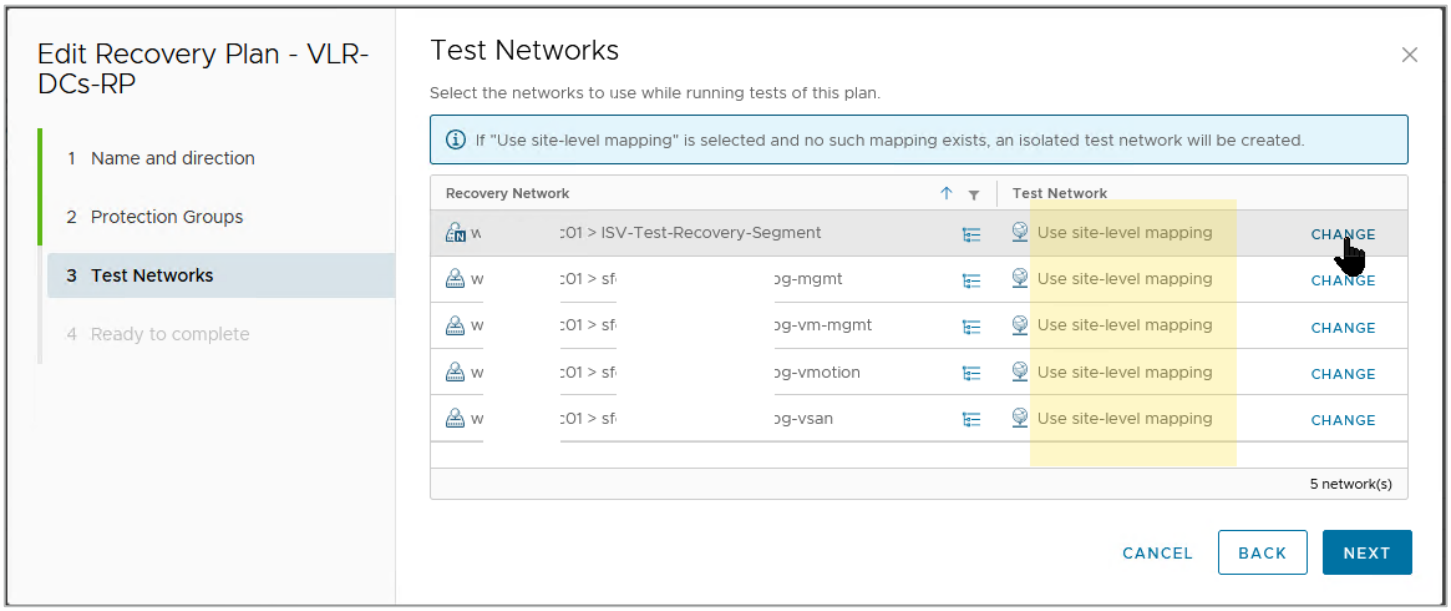

4. Select the fenced-off test network you created in prior steps and click **Save**.

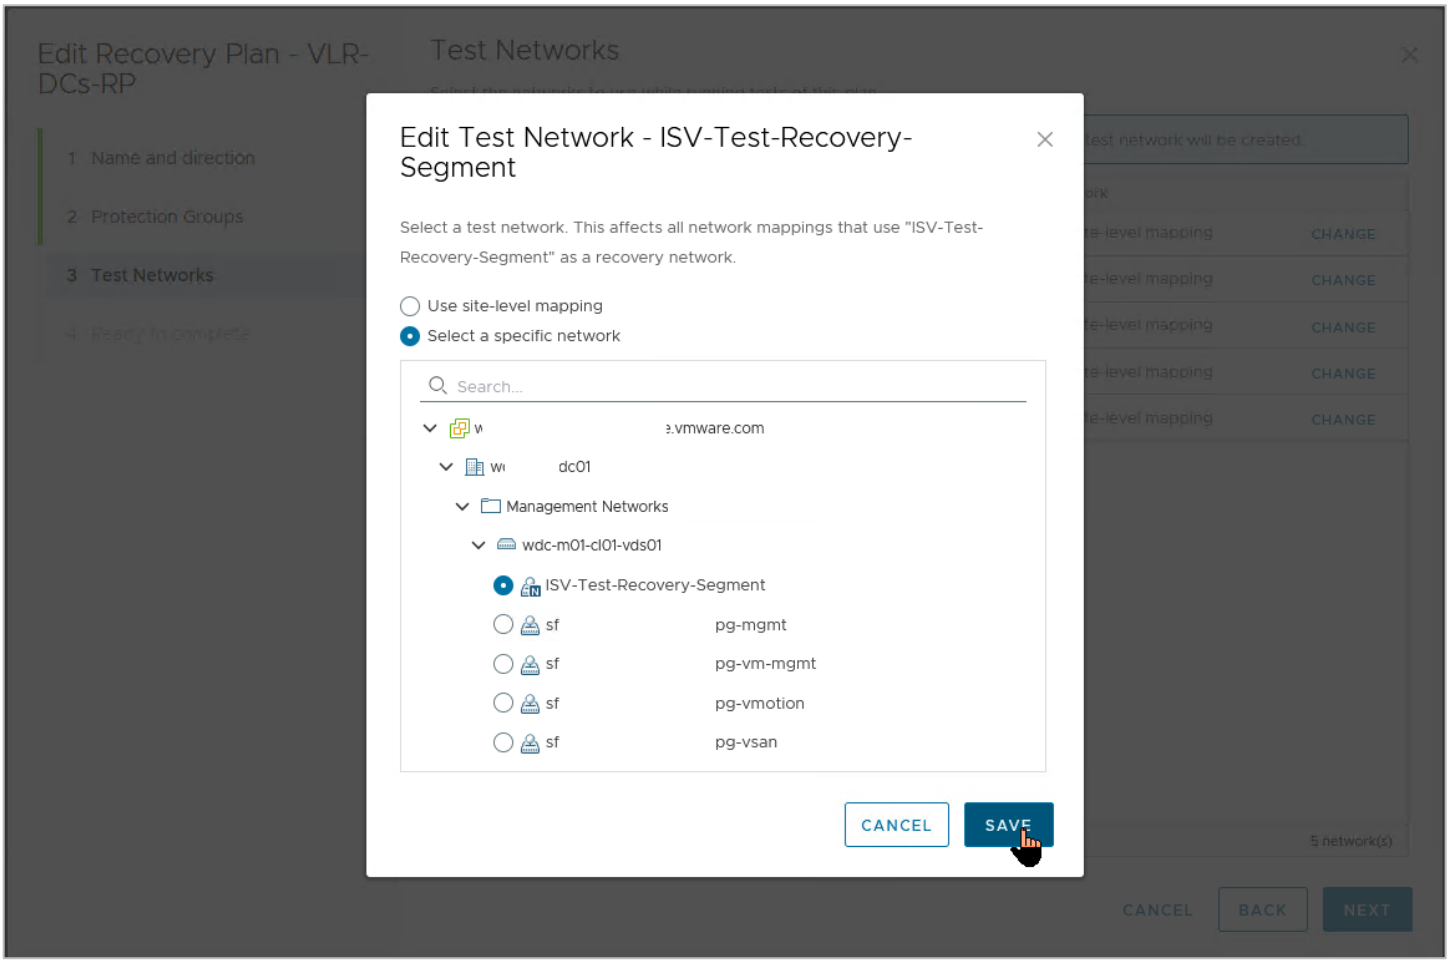

5. Repeat the process for additional mappings you want for other test networks.

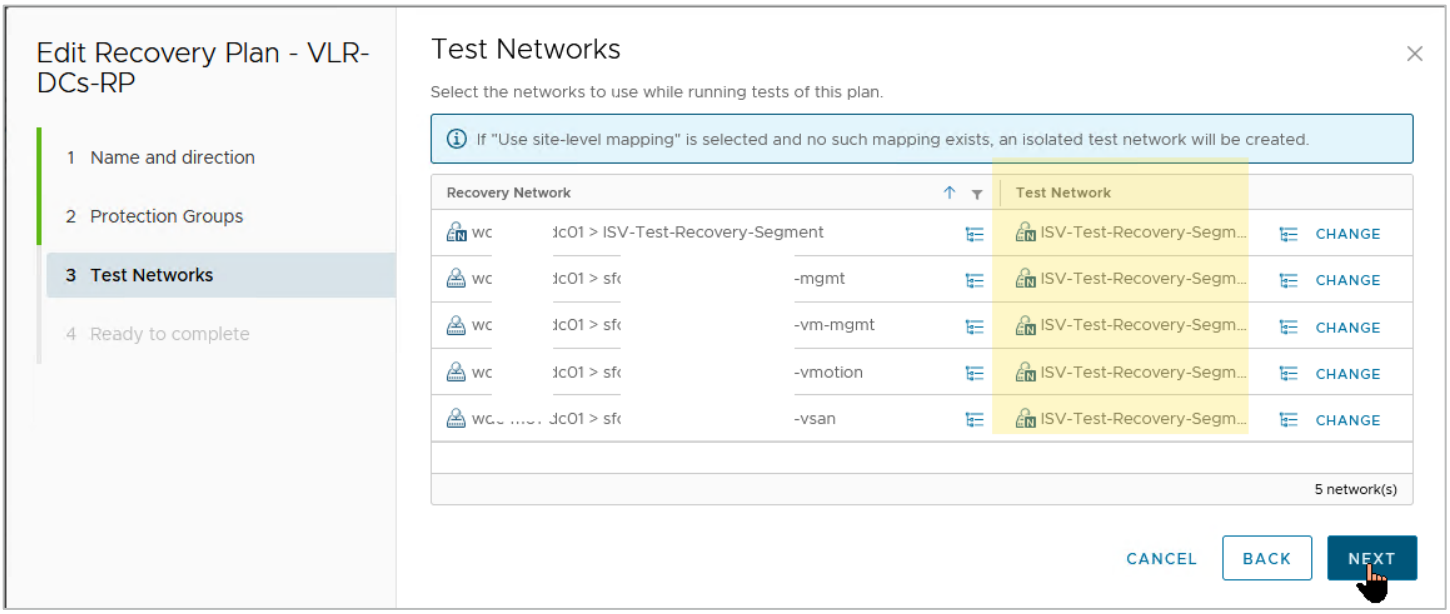

6. Click **Finish** when you're done.

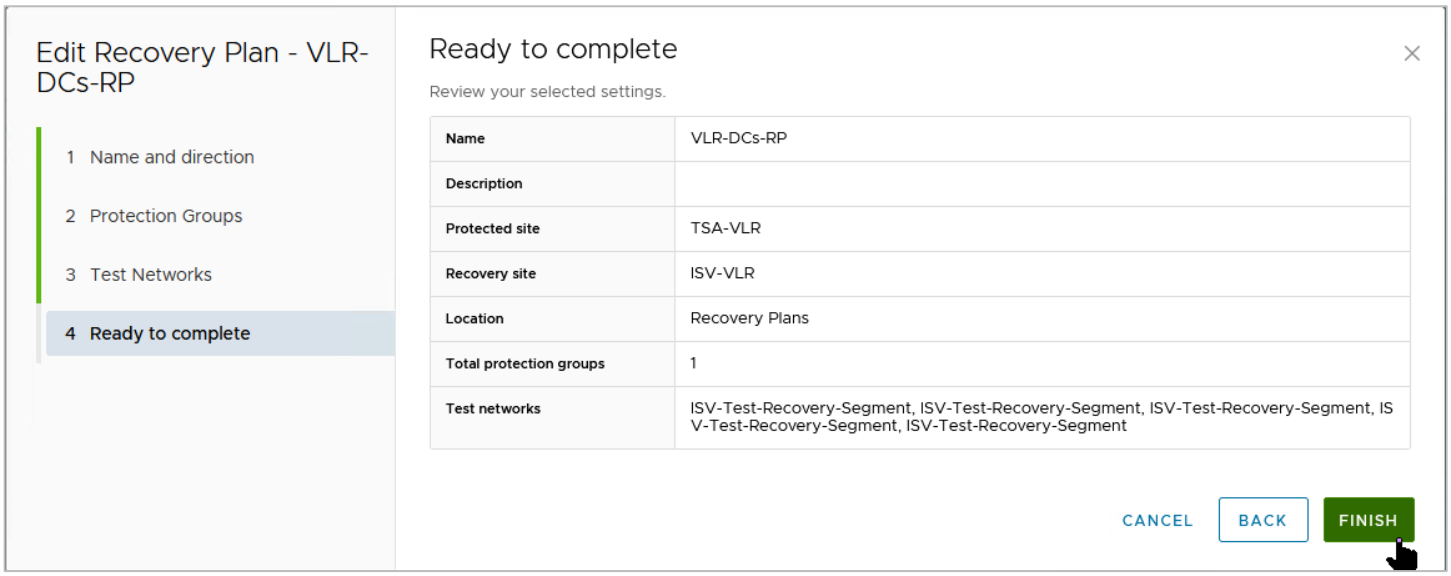

Define the actions in the recovery plans

You created a corresponding recovery plan for each of the protected groups. This was to ensure that one group of protected VMs (the Domain Controllers, for example) becomes completely available before the other VMs are brought online. You want to initiate the recovery of each group of VMs separately.

Here are the recovery plans:

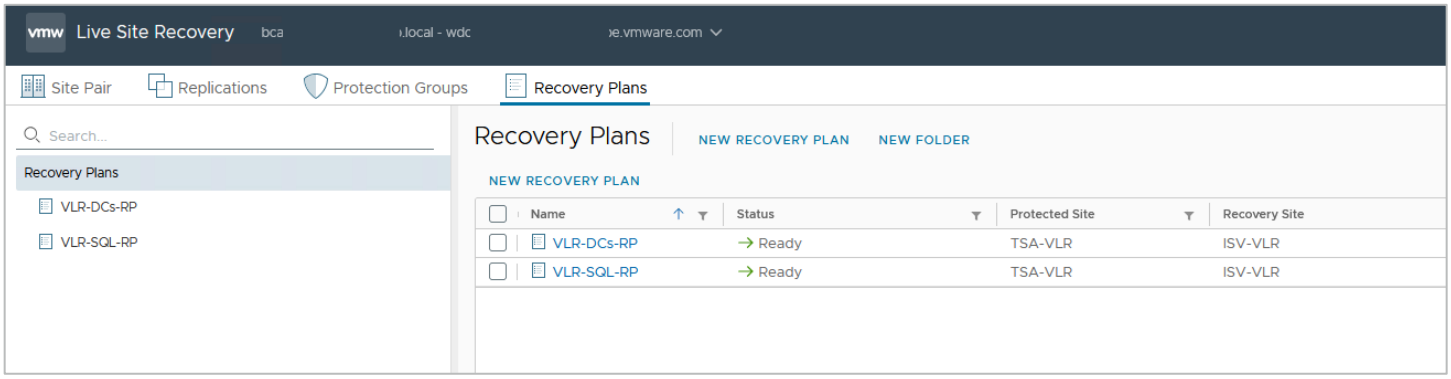

We've mentioned that recovery plans are like a run book for BCDR projects in VMware Live Site Recovery. Next, you'll define the elements of the run book in each of the recovery plans.

Configure the actions for the Domain Controller VMs

First, you'll configure the VMware Live Site Recovery Domain Controller recovery plan: **VLSR-DCs-RP**.

1. Click on the name to select it.

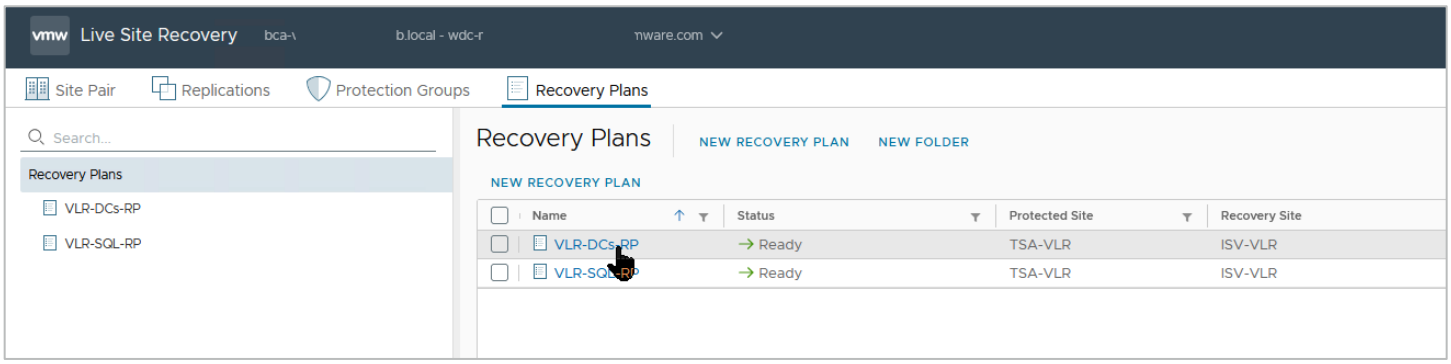

2. Select the **Virtual Machines** tab to display the VMs covered by the plan.

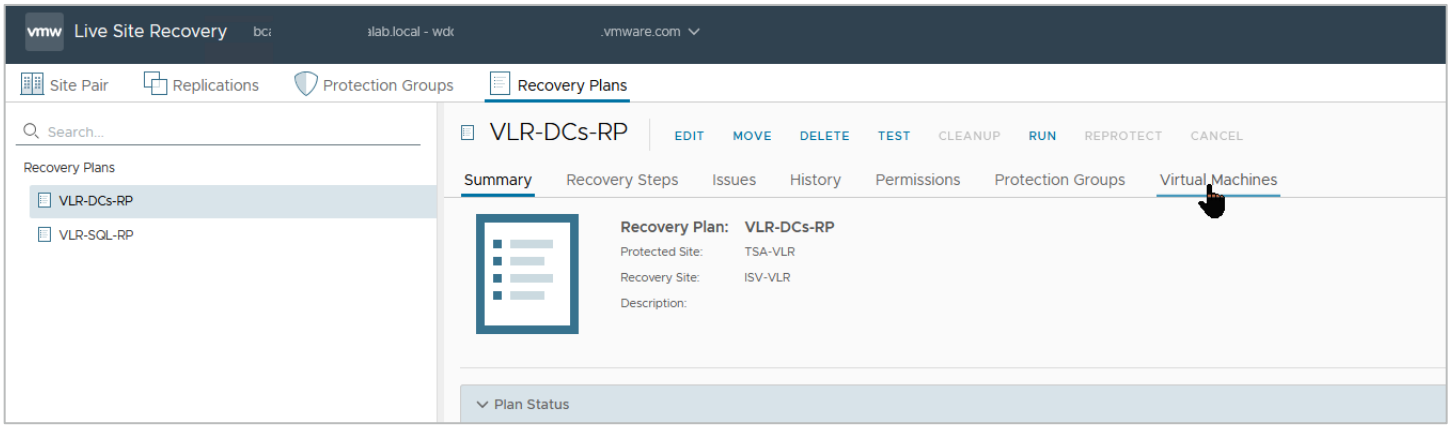

3. Select the checkbox next to the VM you want to configure and click **Configure Recovery**.

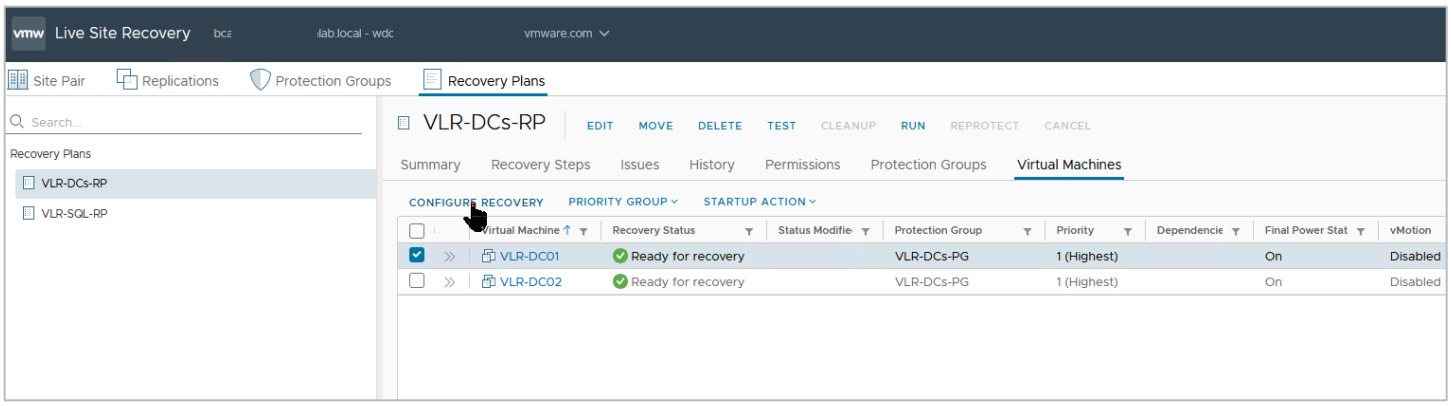

Recovery plans give you many configuration options and flexibility for controlling the desired outcomes for your DR run book. As you'll see, you can configure VMware Live Site Recovery to change the IP address and other necessary IP configurations of the recovered VM.

A warning about the virtualized Domain Controller safety feature: One of the challenges to overcome in recovering Domain Controllers is specifying the order for them to come up to ensure safeguards are correctly in

place. Since restoring a Domain Controller from a backup copy forces the Domain Controller to discard its RID pool, you are likely wondering, "Where does it get a new pool of RIDs if the RID Master is unavailable?" This is a legitimate question in a disaster event where we assume everything in the protected site (including the RID Master itself) is unavailable. Luckily, the Domain Controller safety feature accounts for this scenario by allowing the restored/recovered Domain Controller to regain services after multiple reboots or by manually forcing the Domain Controller's NTDS to run **`restart-service NTDS-force`** if it can communicate with another Domain Controller.

Start by ensuring that you're recovering the Domain Controller holding the FSMO roles first—the other Domain Controllers shouldn't be recovered until this one has been fully recovered. This is done in VMware Live Site Recovery by using the **VM Dependencies** option, which we'll talk about later in this guide).

Here is how you do this for **VLSR-DC02**, which depends on **VLSR-DC01** (the FSMO role holder).

4. Select **VLSR-DC02** and click **Configure Recovery**.

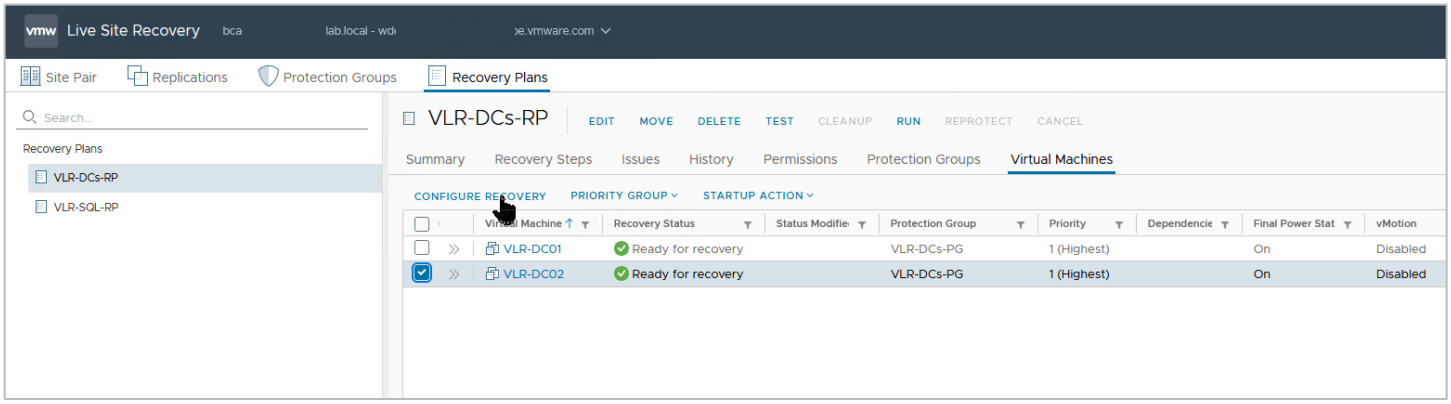

5. Expand **VM Dependencies** and select **View all**. This will show you all the VMs in the recovery plan.
6. Select the other VM you want this VM to depend (or wait) on and click **OK**.

Next, you'll add a **Post Power On Step** task to **VLSR-DC01** which calls a script to reboot the VM after it has been fully recovered. This is a `shutdown -r -t 0` command—nothing fancy. This reboot allows **VLSR-DC01** to self-heal and start its relevant services, which allows it to be available to heal **VLSR-DC02,** which depends on it, and to also provide Active Directory domain services to all the other domain-joined VMs to be recovered.

7. Select **VLSR-DC01** and click **Configure Recovery**.

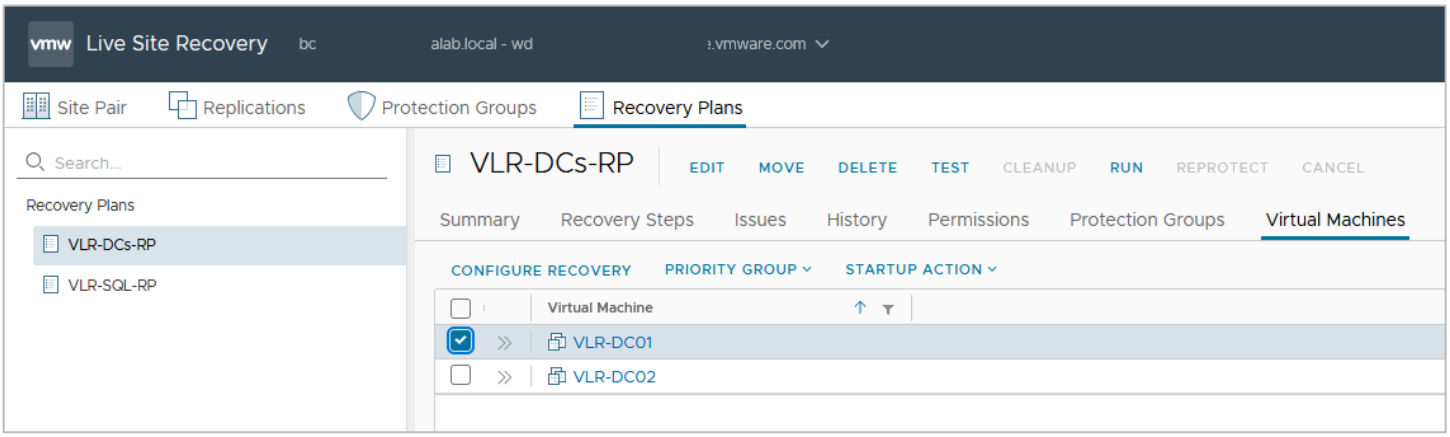

8. Expand **Post Power On Steps** and click **New**.

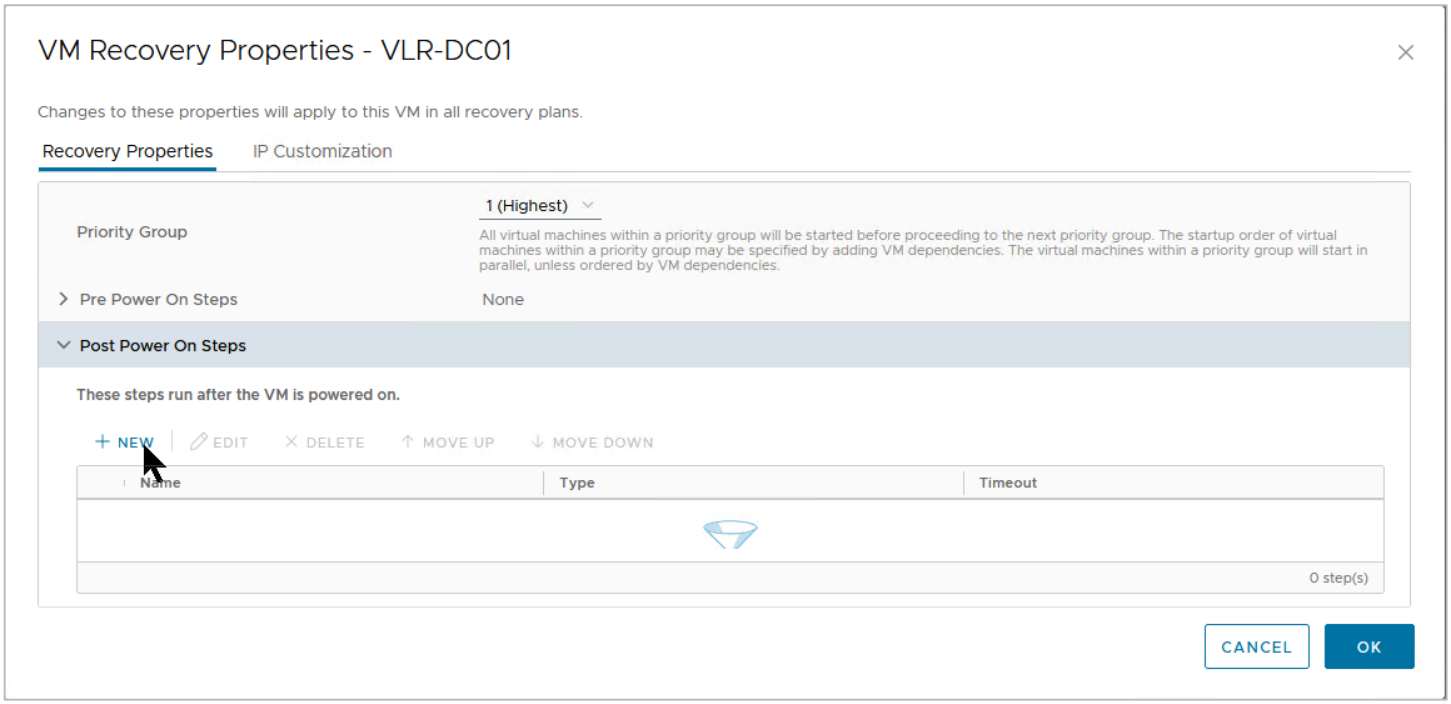

9. Select **Command on Recovered VM**.
10. Give it a descriptive name.
11. Type in the command to run (in our case, we're calling a PowerShell Script named `Run-Post-Script.ps1`, located in the `C:\Install-Files` folder.
12. Click **Add**.

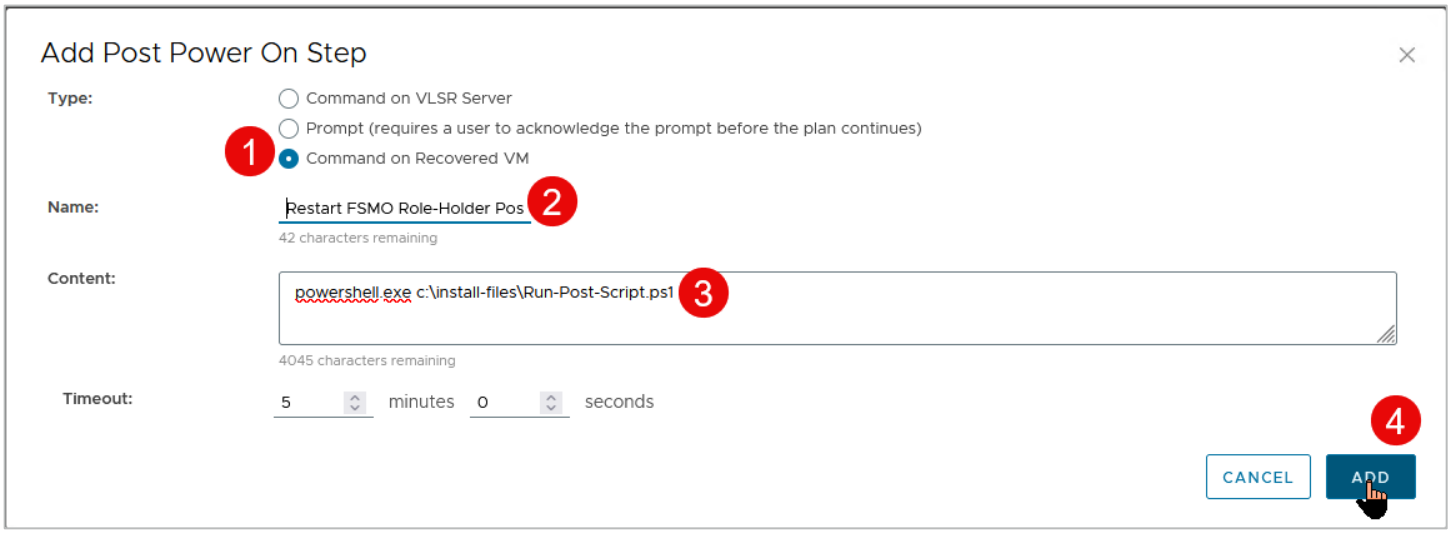

This brings you back to the **VM Recovery Properties** menu.

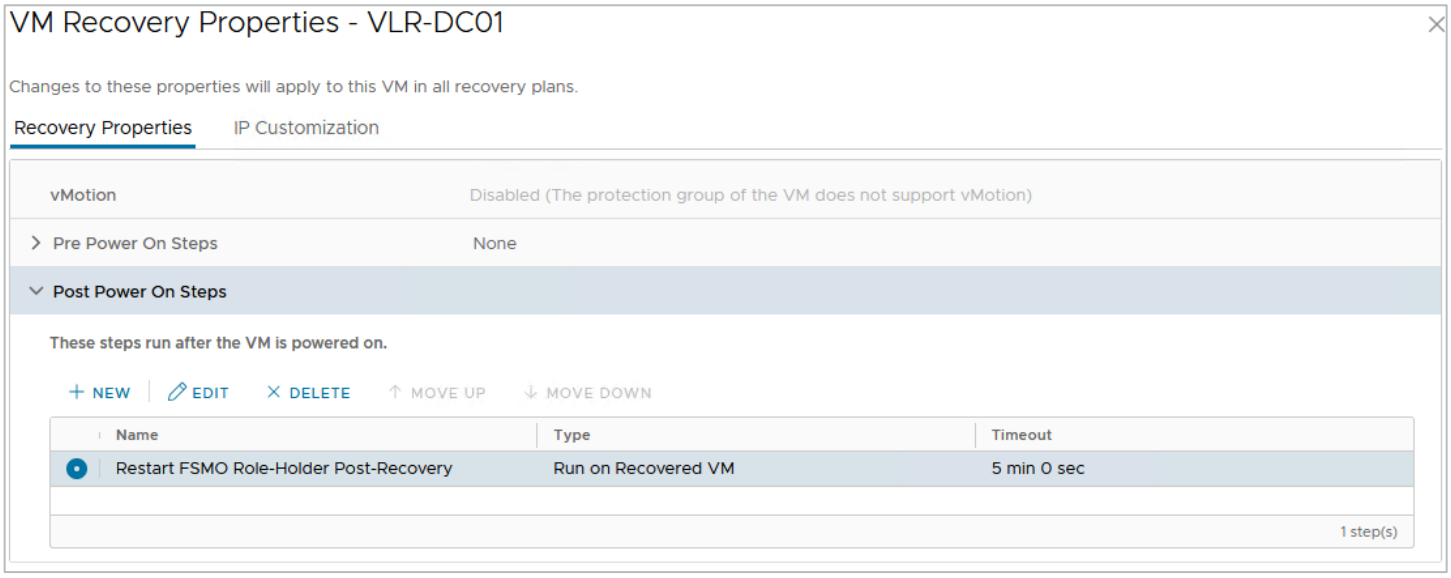

Here is the content **`Run-Post-Script.ps1`**:

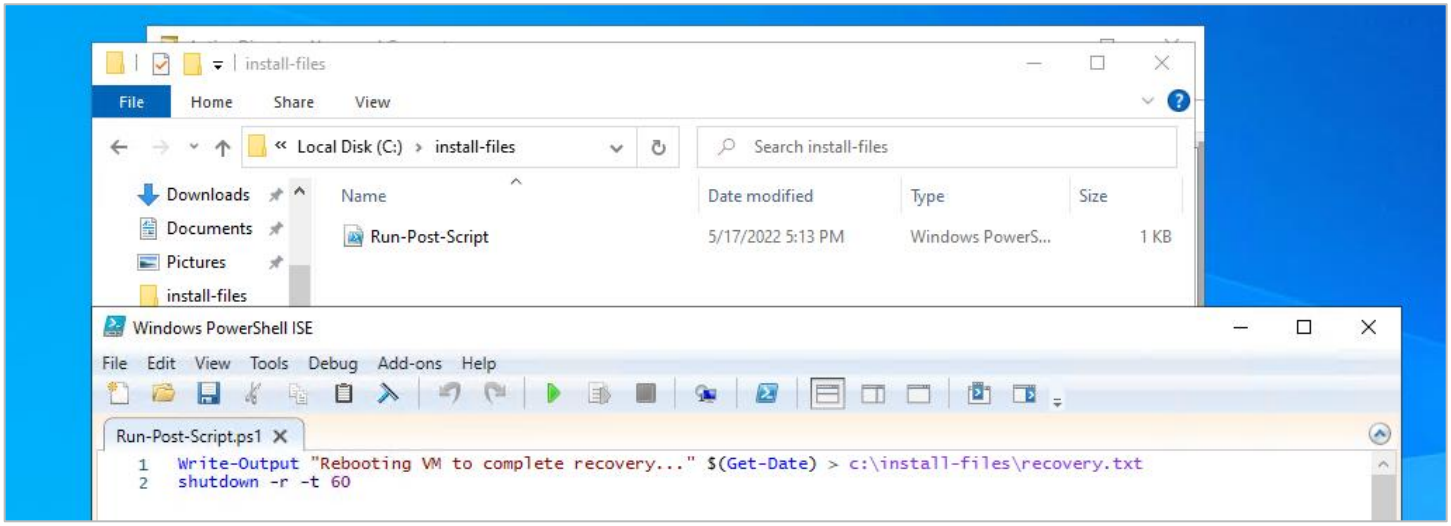

Change the recovered VMs' IP settings

Now, you'll configure the TCP/IP settings for your protected VMs.

1. Go back to the **VM Recovery Properties** menu and click on the **IP Customization** tab.
2. Select the drop-down button in **Select IP customization mode**.
3. Select **Manual IP customization**.

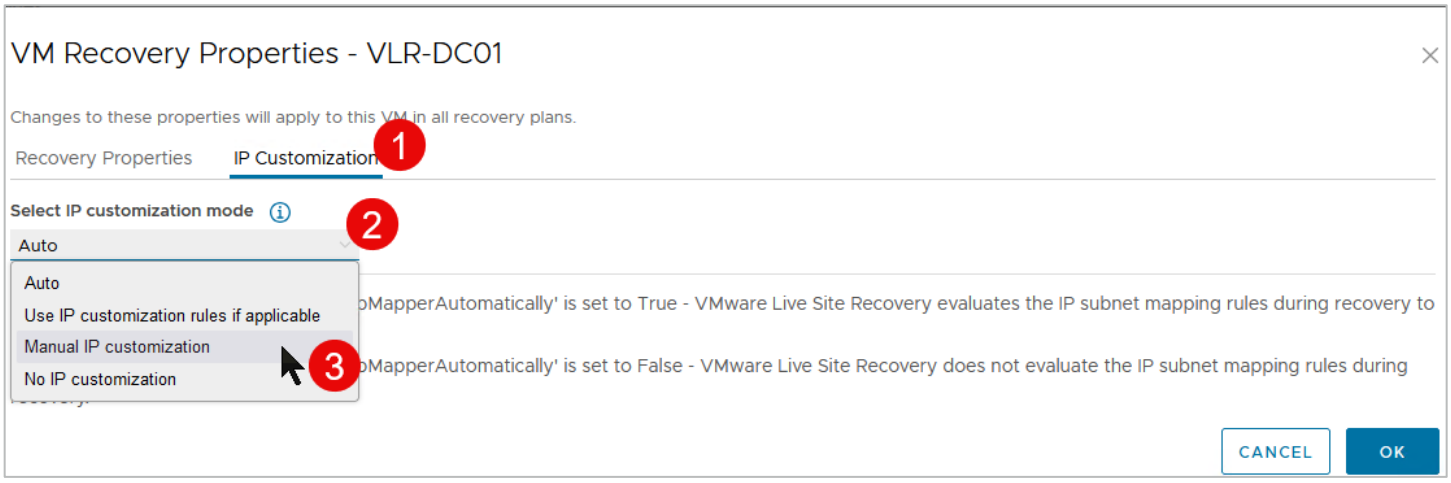

4. Click on **IP Settings - NIC 1** and then **Configure** next to **Protected Site**.

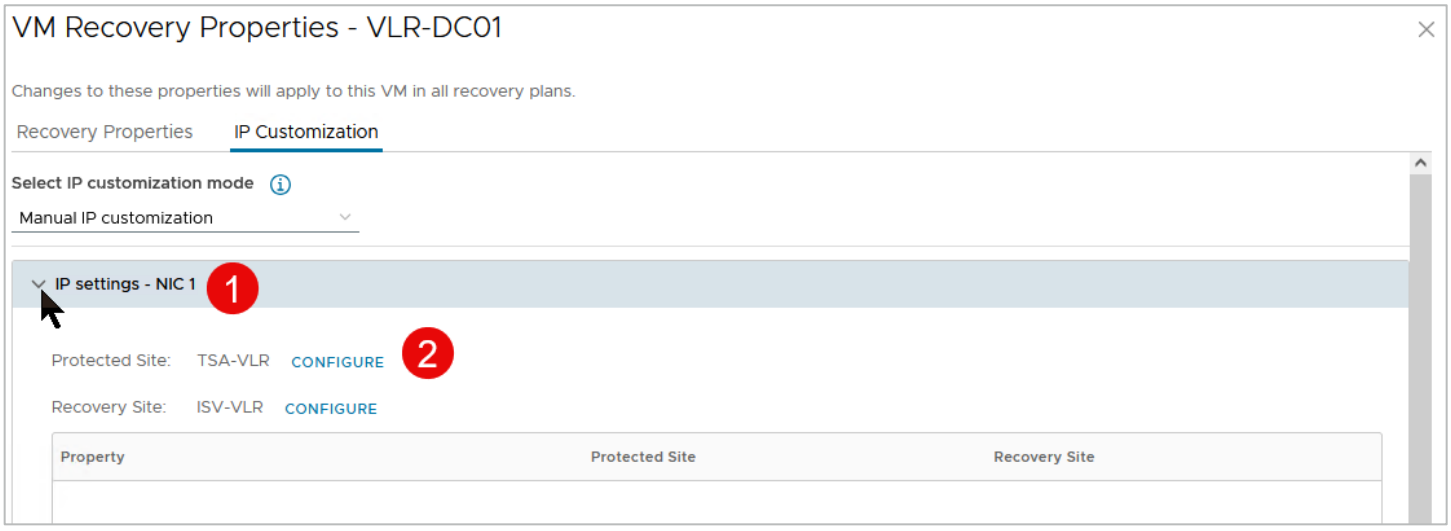

5. Click on **Use the following IPv4 address** and click **Retrieve**. This auto-populates the fields with the VM's current IP address.

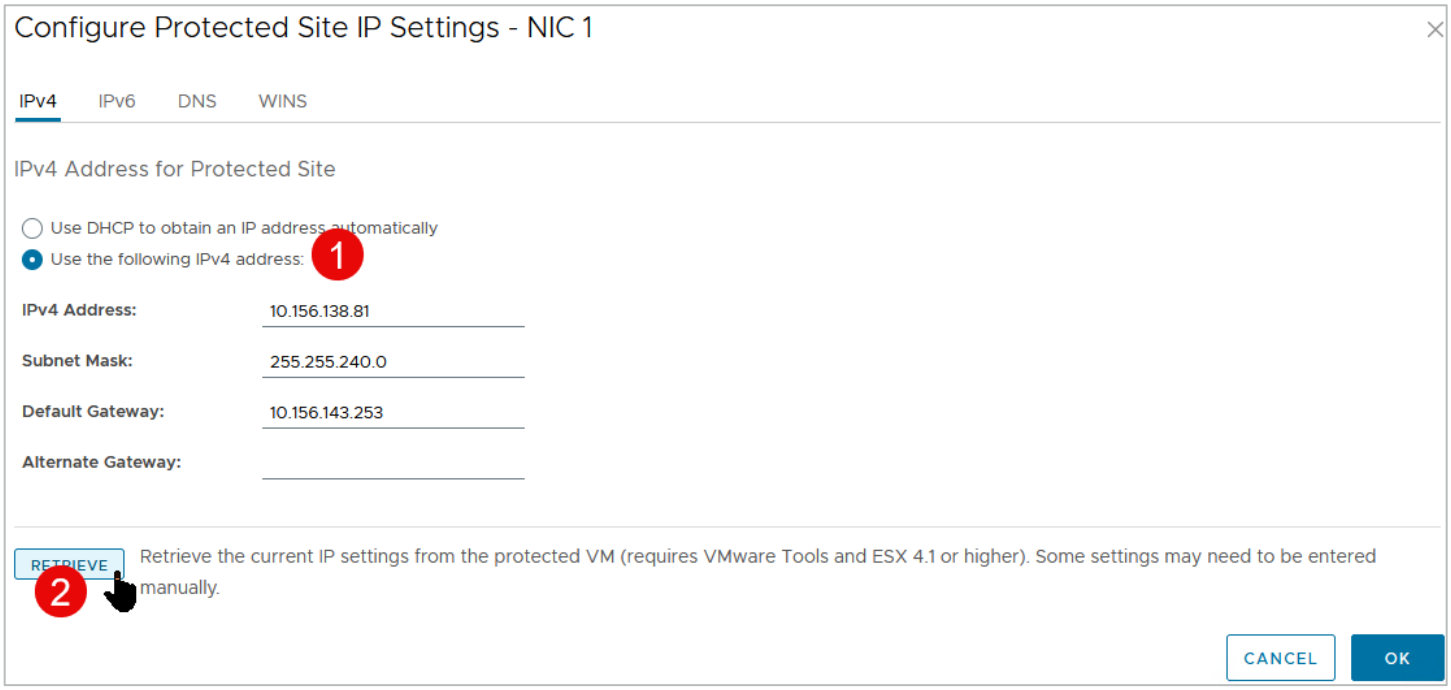

6. Repeat the process for the DNS information. Skip IPv6 and WINS for this exercise.
7. Click **OK** to complete the configuration.

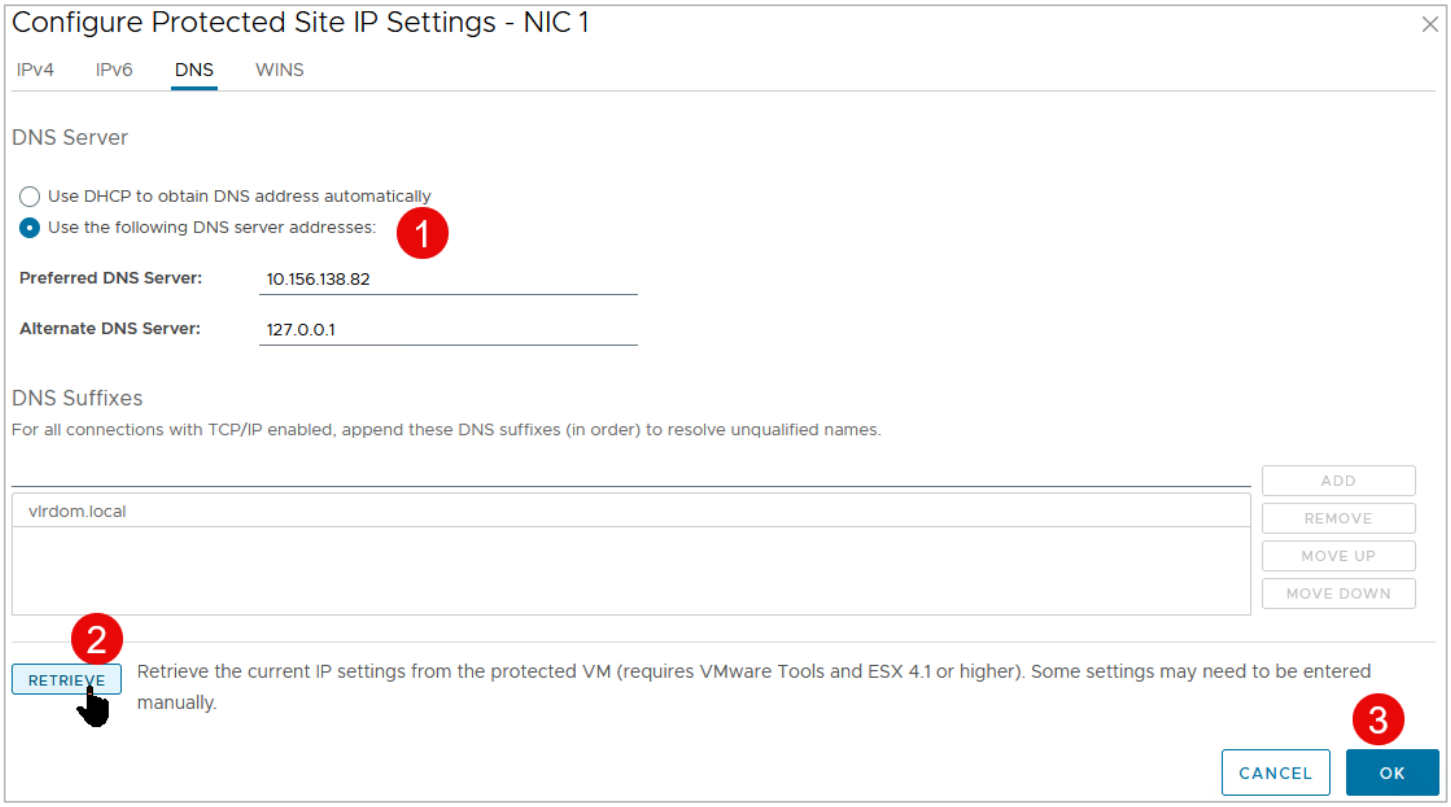

This brings you back to the **VM Recovery Properties → IP Customization** screen.

8. Next to **Recovery Site**, click on **Configure** to specify the IP address information you want to apply to the VM upon recovery.

You'll notice that the **Retrieve** option isn't available on this screen—the values don't currently exist on the VM.

9. Go through the same steps you did for the **Protected Site** values and click **OK** to complete the configuration.

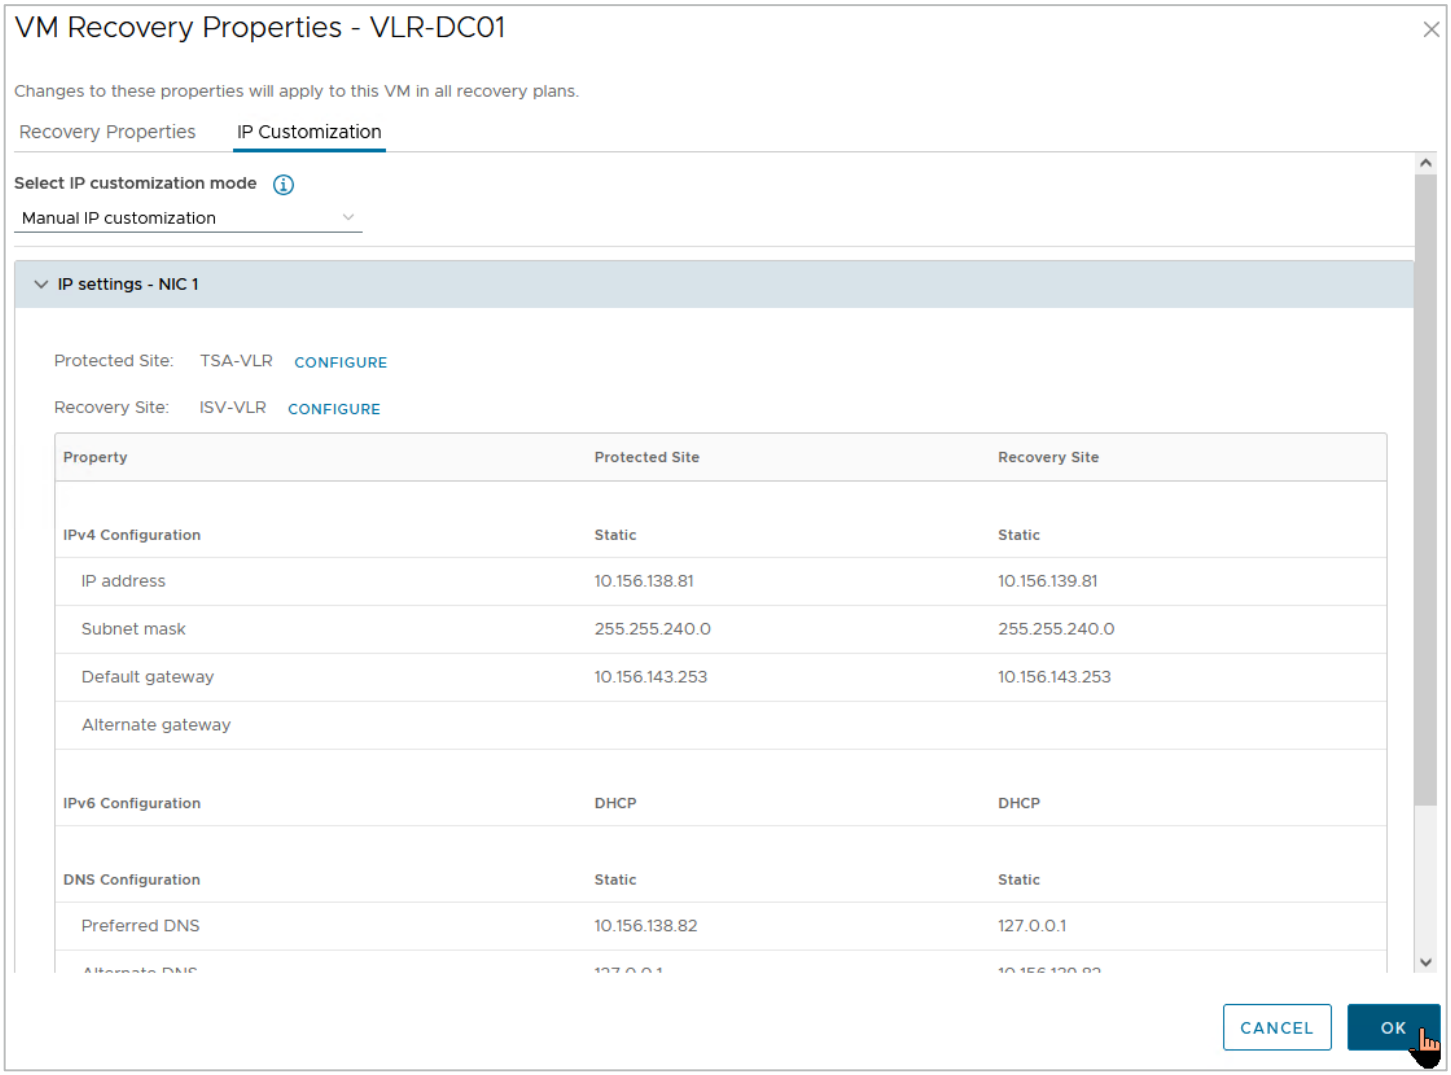

10. Complete this process for all the VMs in all recovery groups unless you want them to:

- Get their IP address configuration information from a DHCP server/IPAM at the recovery site.
- Keep the same IP address because you've stretched the protected site's network segmentsto the recovery site.

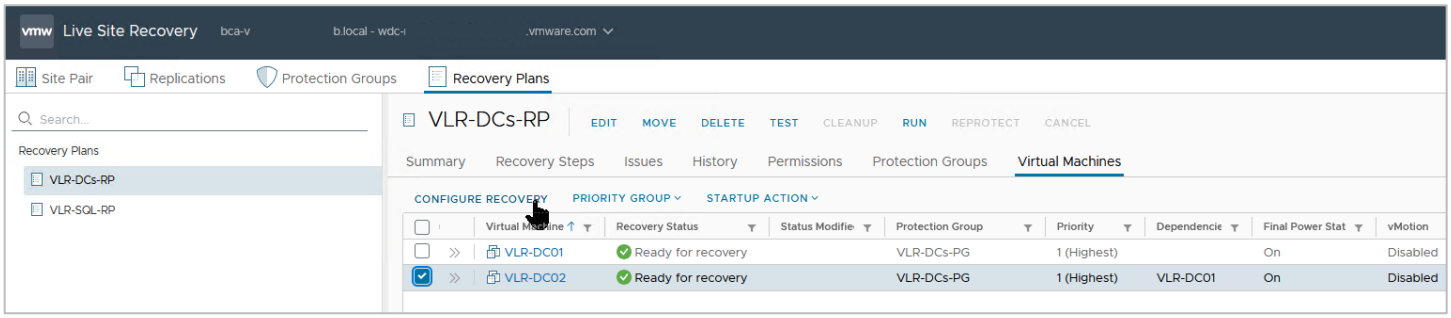

In this example, we create a recovery group for three SQL servers in a protection group. We create a dependency among them such that Papilolo-02 and Papilolo-03 aren't recovered and powered on before Papilolo-01 is fully recovered.

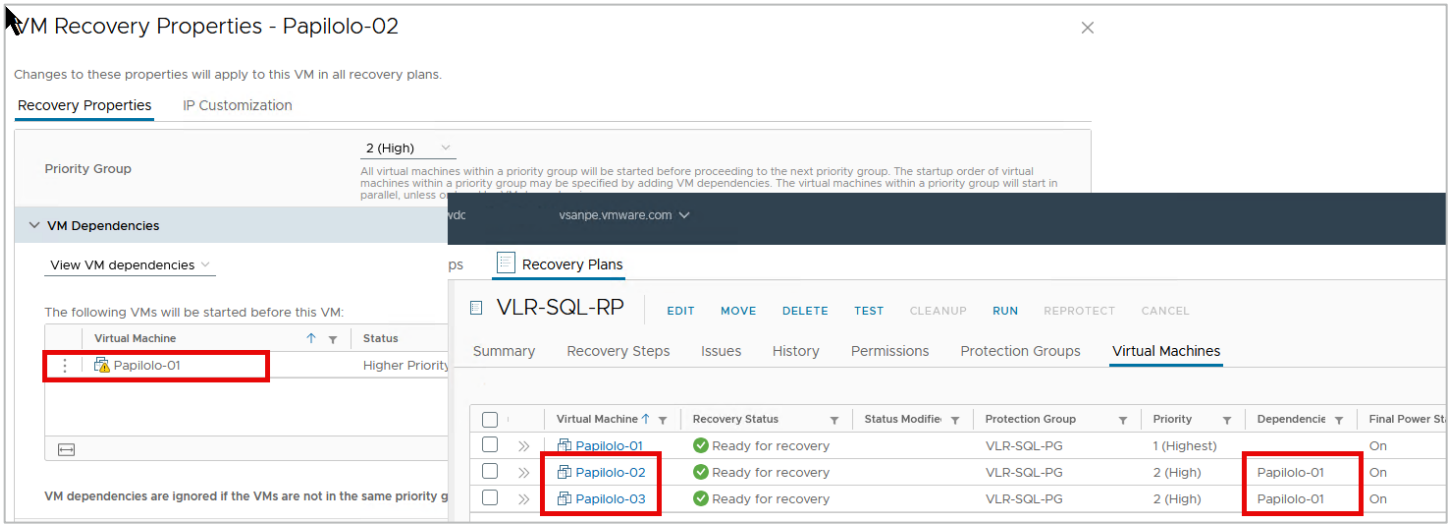

Why create a dependency? This is so the other 2 SQL Server VMs aren't recovered until the listener and cluster virtual IP configurations are normalized after recovering the first SQL Server VM. The parameters for these two Windows/SQL Server clustering configuration settings must be correct and available for the cluster and its resources to become available after recovery.

The recovery process changes the IP address of the recovered VMs and connects them to a different network segment in the recovery site. Consequently, the listener's and cluster virtual IP's IP addresses must also change. This is something that VMware Live Site Recovery can't do natively because it's application-agnostic.

11. Use VMware Live Site Recovery's in-guest script initiation capability to make the changes, just as you did for the operations master Domain Controller. You only need to do this once for the cluster, so only place the script inside Papilolo-01. This ensures the Papilolo-01 VM is recovered first and its configuration changes are completed before Papilolo-02 and Papilolo-03 are recovered.

Here's what that configuration looks like on Papilolo-01:

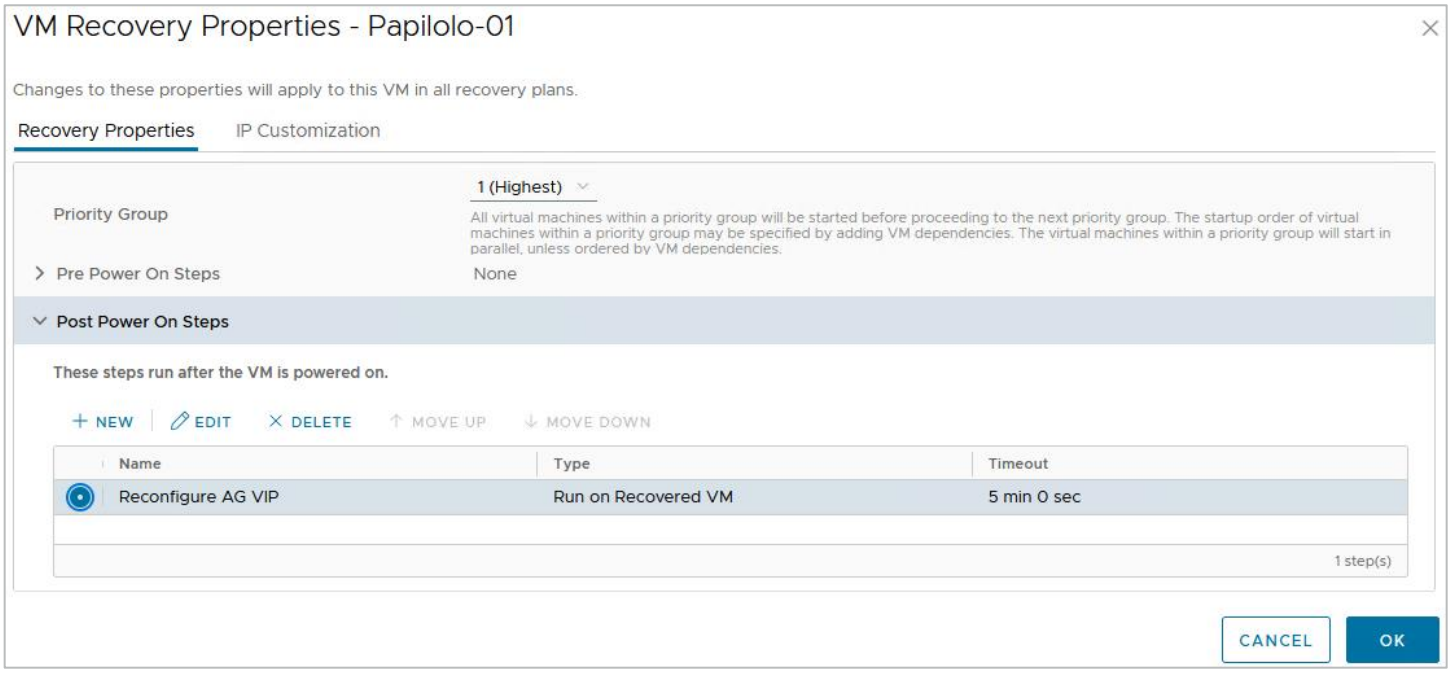

Here's the guest-side command that calls the in-guest PowerShell Script **`Change-Cluster-AG-VIP.ps1`**, located in the **`E:\Install-Files`** folder on the Papilolo-01 VM.

Note: Please follow your internal corporate security practices for storing and running in-guest scripts when deciding where to place these sample scripts.

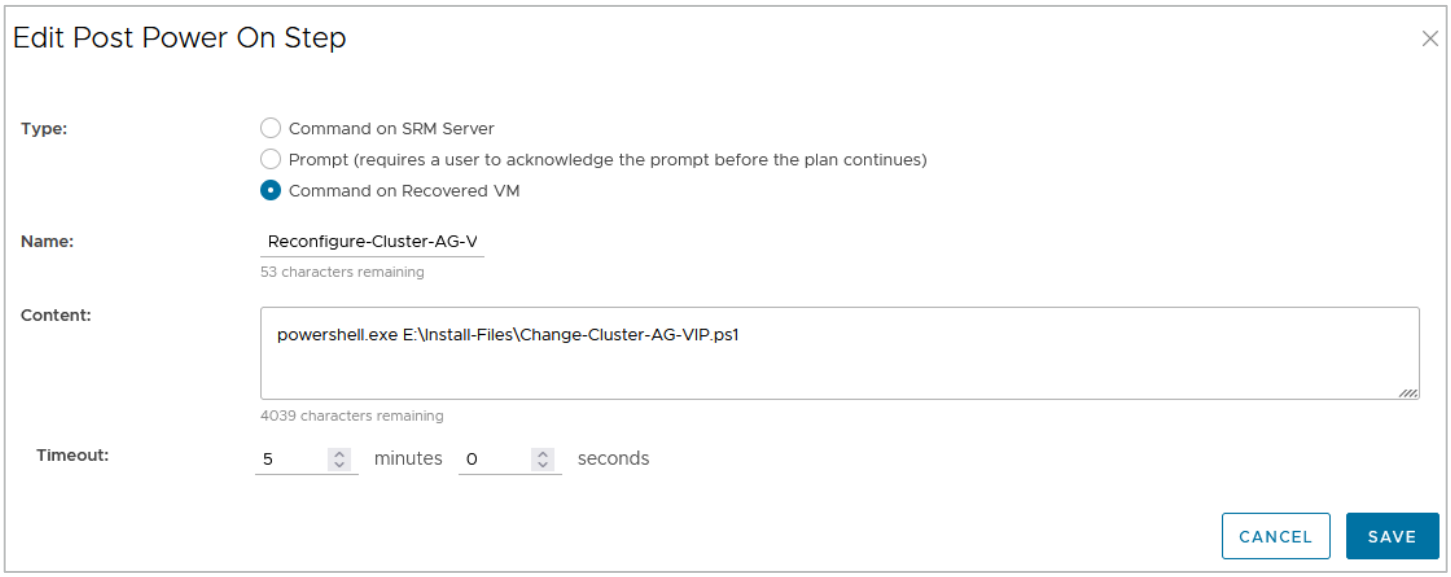

Here's a screenshot of the script itself. You can find this sample script in the appendix: "Change-Cluster-AG-VIP.ps1."

```
# Change-Cluster-AG-VIP.ps1 (For reconfiguring recovered MS SQL Server cluster properties)
Import-Module FailoverClusters

# Let's Force-Start our Cluster first

# Immediately post-recovery, the whole Cluster is down

Start-ClusterNode -FQ

# Let's define our new IP address and subnet mask for the Cluster IP Address
$newClusIP = "10.156.139.87"  # Replace with your new IP address
$newClusMask = "255.255.240.0"    # Replace with your subnet mask

# Get the IP Address of the Cluster resource
$setNewClusIP = Get-ClusterResource -Name "SRM-AG01_Clus_IP"

# Set the new IP address and subnet mask for the Cluster resource
$setNewClusIP | Set-ClusterParameter -Name Address -Value $newClusIP
$setNewClusIP | Set-ClusterParameter -Name SubnetMask -Value $newClusMask

########## Next, we modify the AG VIP
# Let's define our new IP address and subnet mask for the AG VIP Address
$newAGIP = "10.156.139.88"  # Replace with your new IP address
$newAGMask = "255.255.240.0"    # Replace with your subnet mask

# Get the IP Address of the AG resource
$setNewAGIP = Get-ClusterResource -Name "SRM-AG01-IP"

# Set the new IP address and subnet mask for the AG resource
$setNewAGIP | Set-ClusterParameter -Name Address -Value $newAGIP
$setNewAGIP | Set-ClusterParameter -Name SubnetMask -Value $newAGMask

# Bring the resources offline
Stop-ClusterResource "SRM-AG01_Clus_IP"
Stop-ClusterResource "SRM-AG01_SRM-AG-List"
Stop-ClusterResource "Cluster Name"
Stop-ClusterResource "SRM-AG01"
Stop-ClusterResource "SRM-AG01-IP"

# We now start up everything
Start-ClusterResource "SRM-AG01"
Start-ClusterResource "SRM-AG01-IP"
Start-ClusterResource "SRM-AG01_SRM-AG-List"
Start-ClusterResource "SRM-AG01_Clus_IP"
Start-ClusterResource "Cluster Name"
```

Test the disaster recovery plan

You have all the configuration pieces in place. Now you're ready to test your disaster recovery plan. Testing is essential because you don't want to discover that your DR plan doesn't do what you want during an actual disaster.

1. From the **Recovery Plans** tab, click on **Test**.
2. Notice that the **Plan Status** shows **Ready**. This indicates the recovery plan is ready to run.

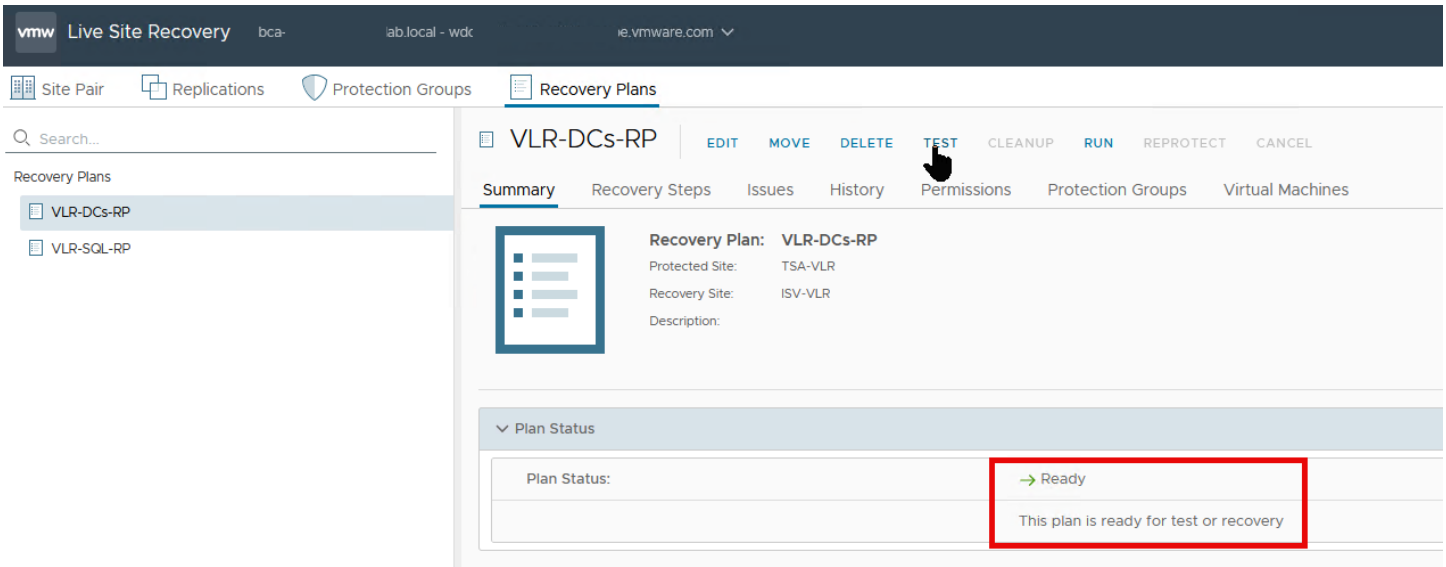

3. Confirm that "Replicate recent changes to recovery site" is checked, then click "Next"

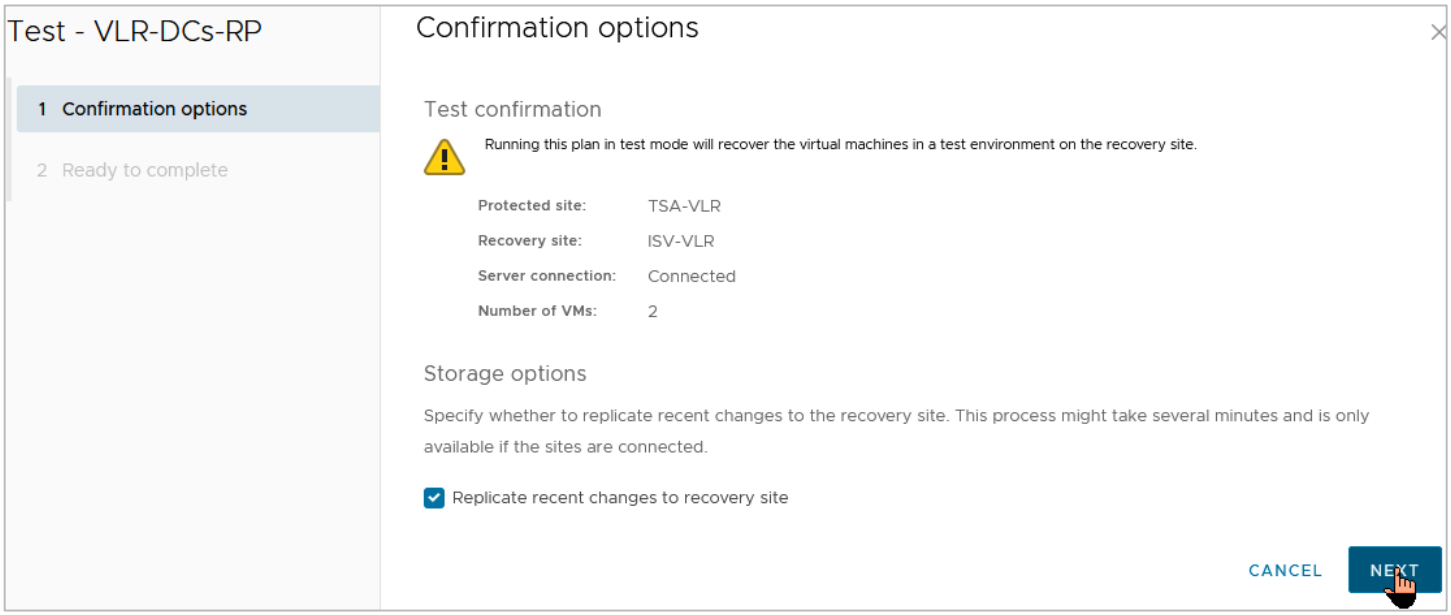

4. Click **Finish** to begin the test recovery process.

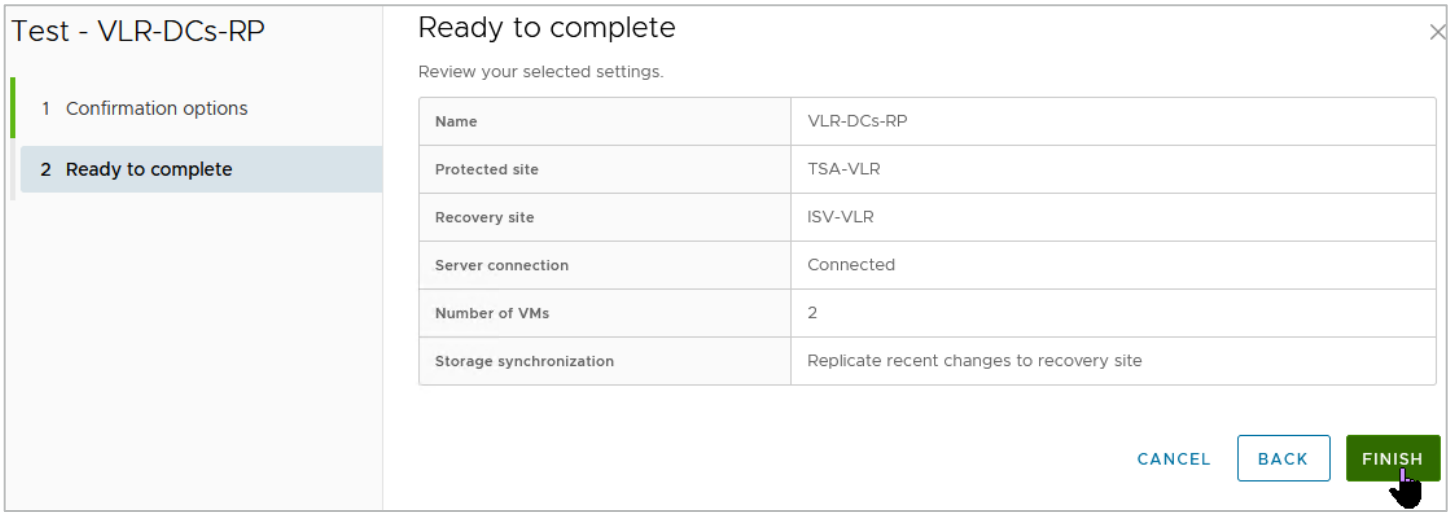

Recovery Steps shows detailed information about actions taken during recovery. Also, notice that DC02 was powered on only after DC01 was fully recovered and the in-guest script was run. This is the dependency you previously configured in the recovery plan.

Recovery Plans

VLR-DCs-RP EDIT MOVE DELETE TEST CLEANUP RUN REPROTECT CANCEL

Summary Recovery Steps Issues History Permissions Protection Groups Virtual Machines

EXPORT STEPS TEST CLEANUP RUN REPROTECT CANCEL

Plan status: Test complete

Description: The virtual machines have been recovered in a test environment at the recovery site. Review the plan history to

Recovery Step	Status	Step Started	Step Completed
1. Synchronize storage	Success	Wednesday, September 11, 2024 2:01:01 PM	Wednesday, September 11, 2024 2:01:01 PM
2. Restore recovery site hosts from standby	Success	Wednesday, September 11, 2024 2:01:01 PM	Wednesday, September 11, 2024 2:01:01 PM
3. Suspend non-critical VMs at recovery site			
4. Create writable storage snapshot	Success	Wednesday, September 11, 2024 2:01:01 PM	Wednesday, September 11, 2024 2:01:05 PM
5. Configure test networks	Success	Wednesday, September 11, 2024 2:01:04 PM	Wednesday, September 11, 2024 2:01:05 PM
6. Power on priority 1 VMs	Success	Wednesday, September 11, 2024 2:01:04 PM	Wednesday, September 11, 2024 2:04:50 PM
6.1. VLR-DC02	Success	Wednesday, September 11, 2024 2:01:04 PM	Wednesday, September 11, 2024 2:04:50 PM
6.1.1. Guest startup	Success	Wednesday, September 11, 2024 2:01:04 PM	Wednesday, September 11, 2024 2:02:00 PM
6.1.2. Customize IP	Success	Wednesday, September 11, 2024 2:02:00 PM	Wednesday, September 11, 2024 2:02:16 PM
6.1.3. Guest shutdown	Success	Wednesday, September 11, 2024 2:02:16 PM	Wednesday, September 11, 2024 2:02:30 PM
6.1.4. Power on	Success	Wednesday, September 11, 2024 2:03:52 PM	Wednesday, September 11, 2024 2:03:54 PM
6.1.5. Wait for VMware tools	Success	Wednesday, September 11, 2024 2:03:54 PM	Wednesday, September 11, 2024 2:04:50 PM
6.2. VLR-DC01	Success	Wednesday, September 11, 2024 2:01:05 PM	Wednesday, September 11, 2024 2:03:51 PM
6.2.1. Guest startup	Success	Wednesday, September 11, 2024 2:01:05 PM	Wednesday, September 11, 2024 2:01:59 PM
6.2.2. Customize IP	Success	Wednesday, September 11, 2024 2:01:59 PM	Wednesday, September 11, 2024 2:02:28 PM
6.2.3. Guest shutdown	Success	Wednesday, September 11, 2024 2:02:28 PM	Wednesday, September 11, 2024 2:02:49 PM
6.2.4. Power on	Success	Wednesday, September 11, 2024 2:02:49 PM	Wednesday, September 11, 2024 2:02:53 PM
6.2.5. Wait for VMware tools	Success	Wednesday, September 11, 2024 2:02:53 PM	Wednesday, September 11, 2024 2:03:48 PM
6.2.6. Command: Restart FSMO Role-Holder Post-Recovery	Success	Wednesday, September 11, 2024 2:03:48 PM	Wednesday, September 11, 2024 2:03:51 PM
7. Power on priority 2 VMs			
8. Power on priority 3 VMs			
9. Power on priority 4 VMs			
10. Power on priority 5 VMs			

Notice that the recovered Domain Controller VMs are powered on and running in the vCenter on the recovery site.

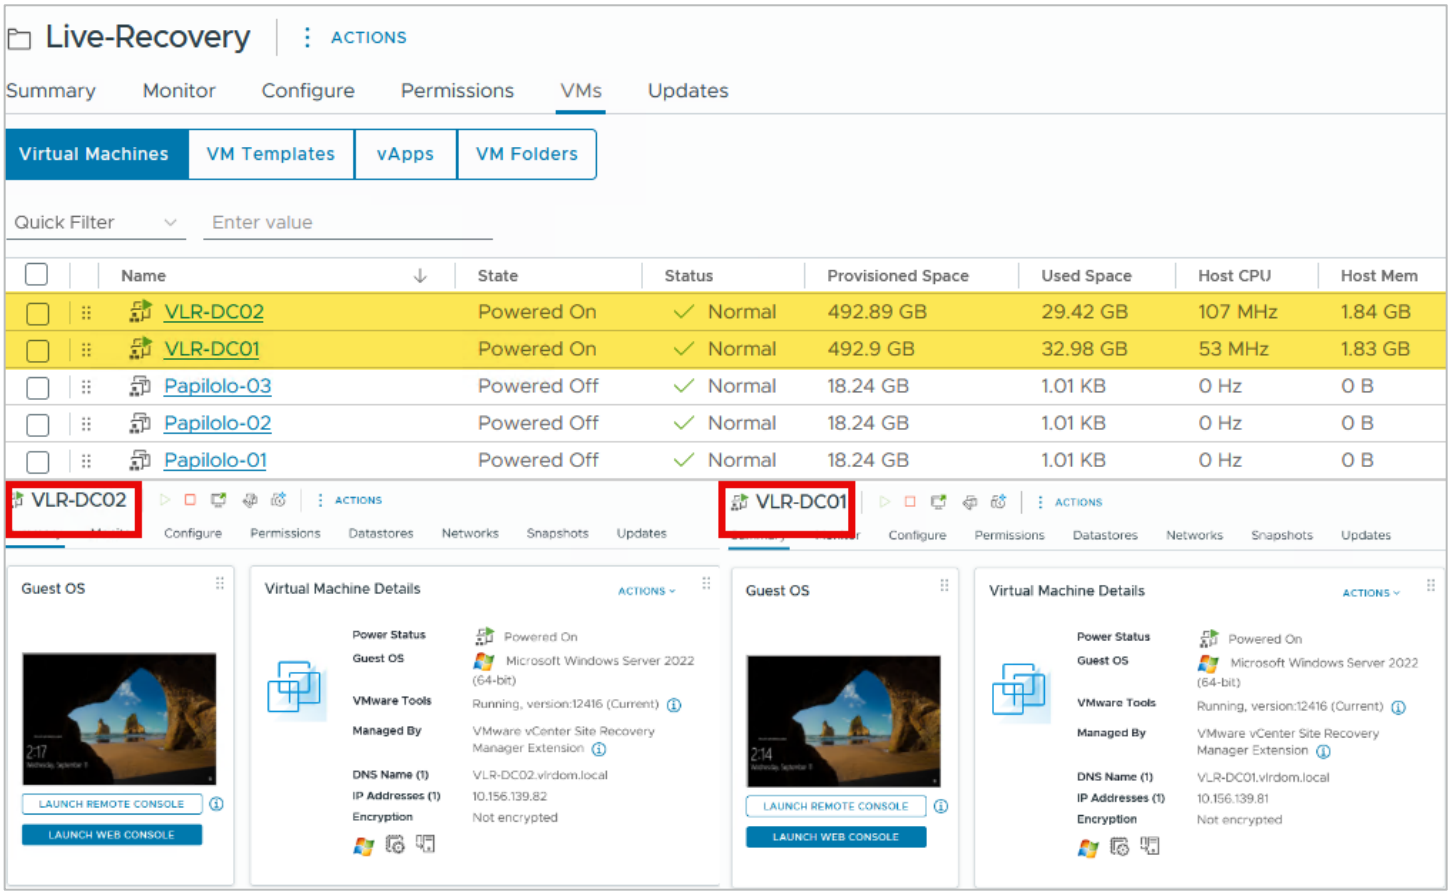

The same Domain Controller and SQL Server VMs are still running uninterrupted at the protected site.

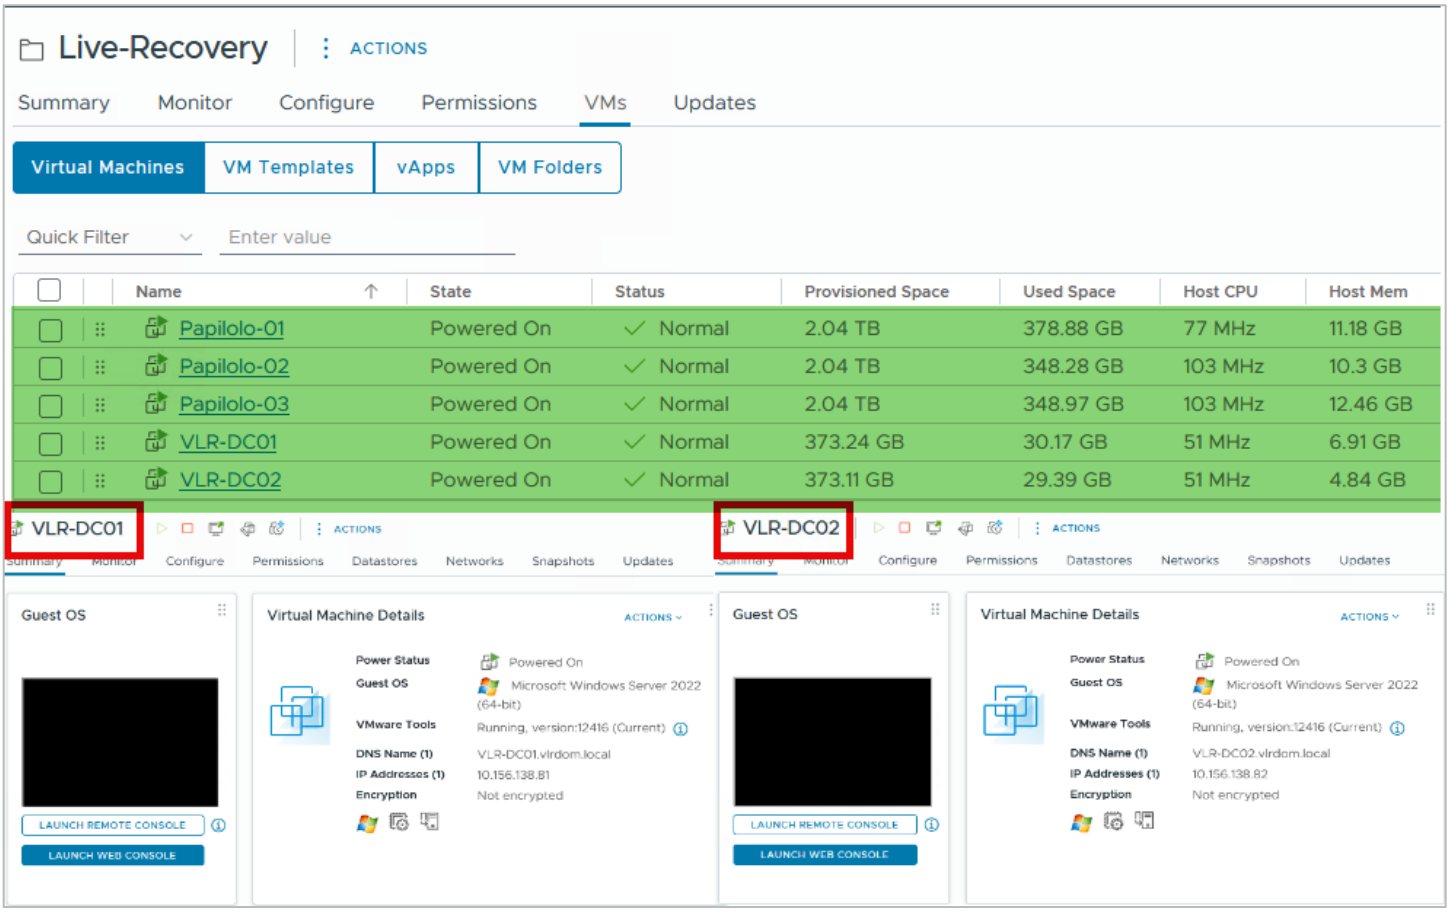

5. Log into the protected and recovered VMs to verify they are both accessible.

This demonstrates the unparalleled on-demand DR plan verification capabilities of VMware Live Site Recovery. You can leverage the test failover feature to satisfy internal and external regulatory compliance or SLA conformance requirements without interrupting their production infrastructure or scheduling an outage.

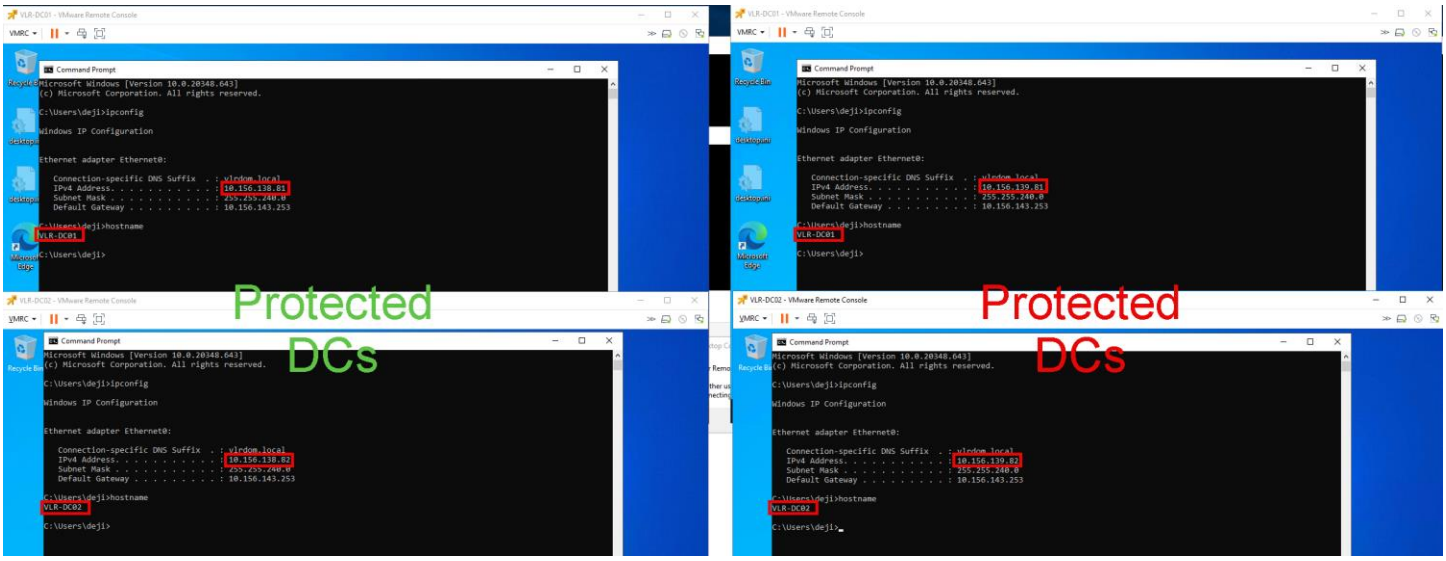

Safe Active Directory Domain Controller recovery in action

Let's take a look at what happened to our protected Active Directory infrastructure after a simulated disaster recovery event completed using VMware Live Site Recover.

The first time the recovered Domain Controllers boot up, Windows automatically detects the change in their VM-Generation ID.

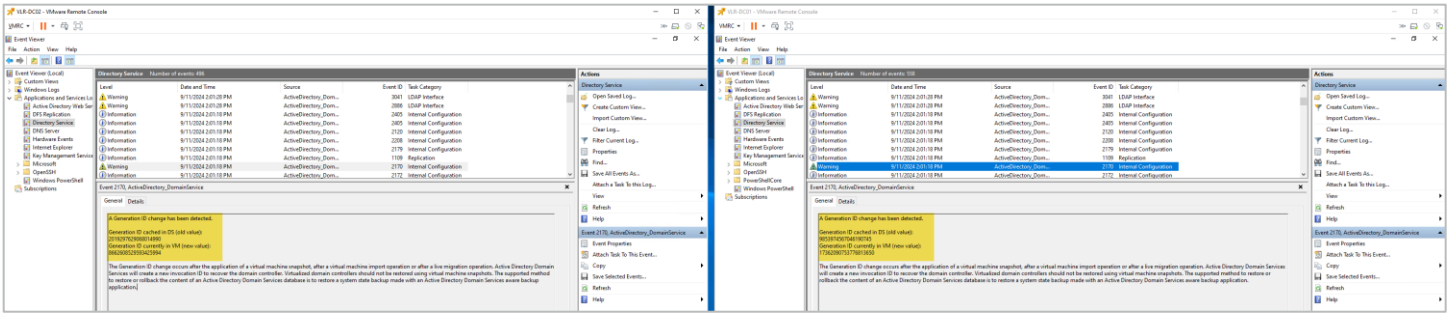

The Windows Domain Controller safety feature immediately kicks in, and the recovered Domain Controller VMs are taken through the remediation process. Among other effects discussed in previous sections, Netlogon, DNS, and other services cannot start during this remediation process.

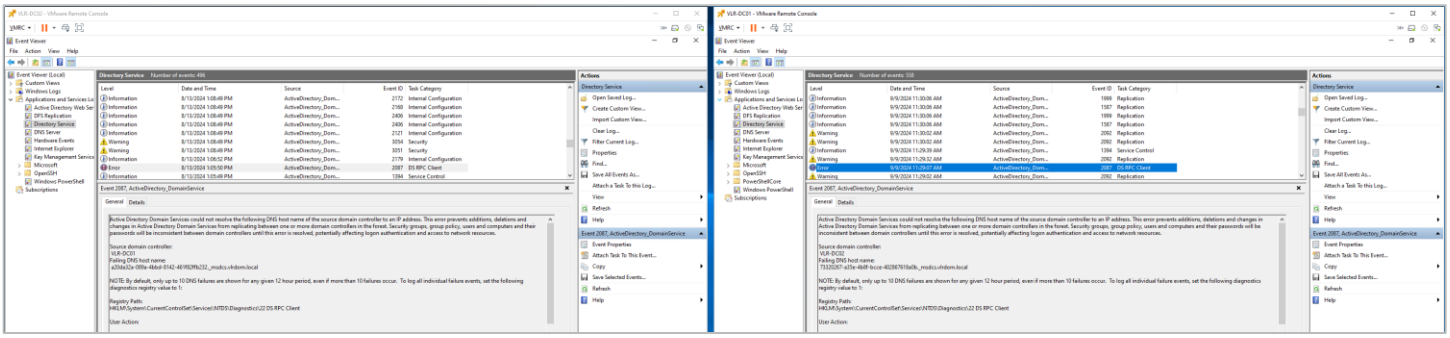

The FSMO role holder (DC01) isn't considered a [working?] Domain Controller now.

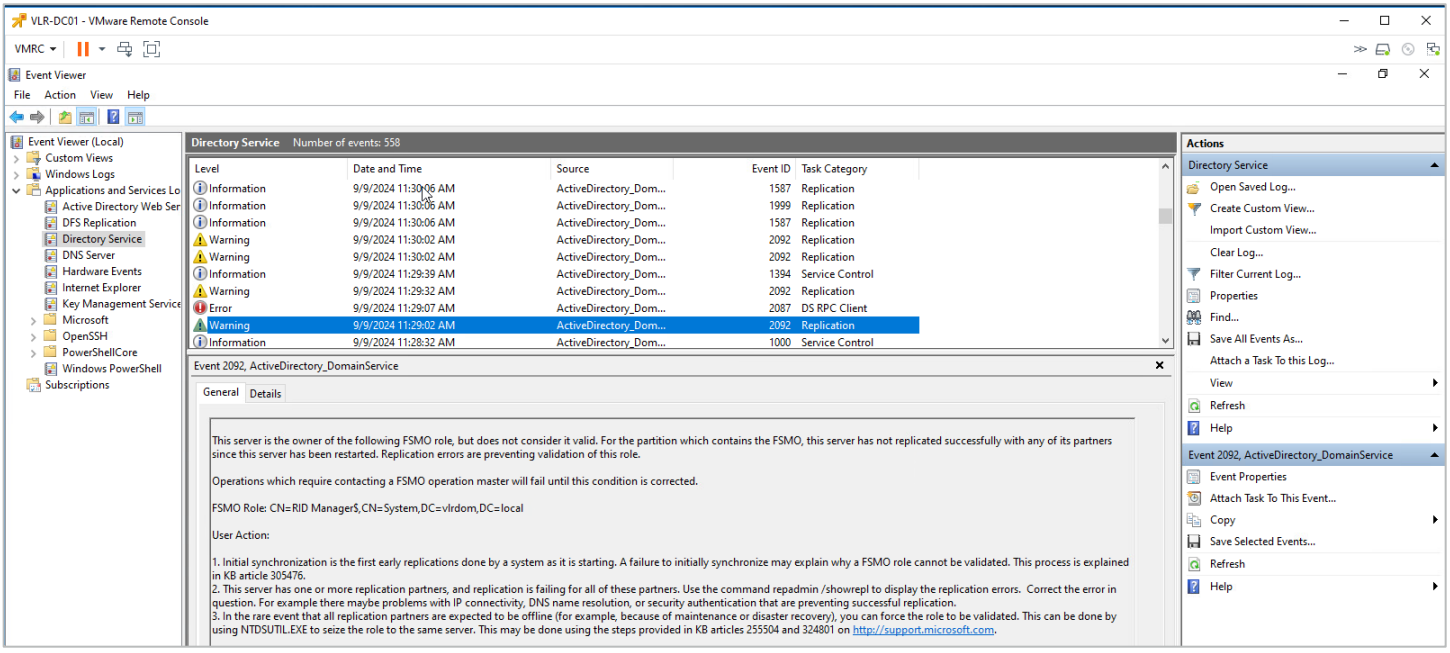

After rebooting the FSMO role holder (DC01) the second time with our in-guest script, things begin to look better.

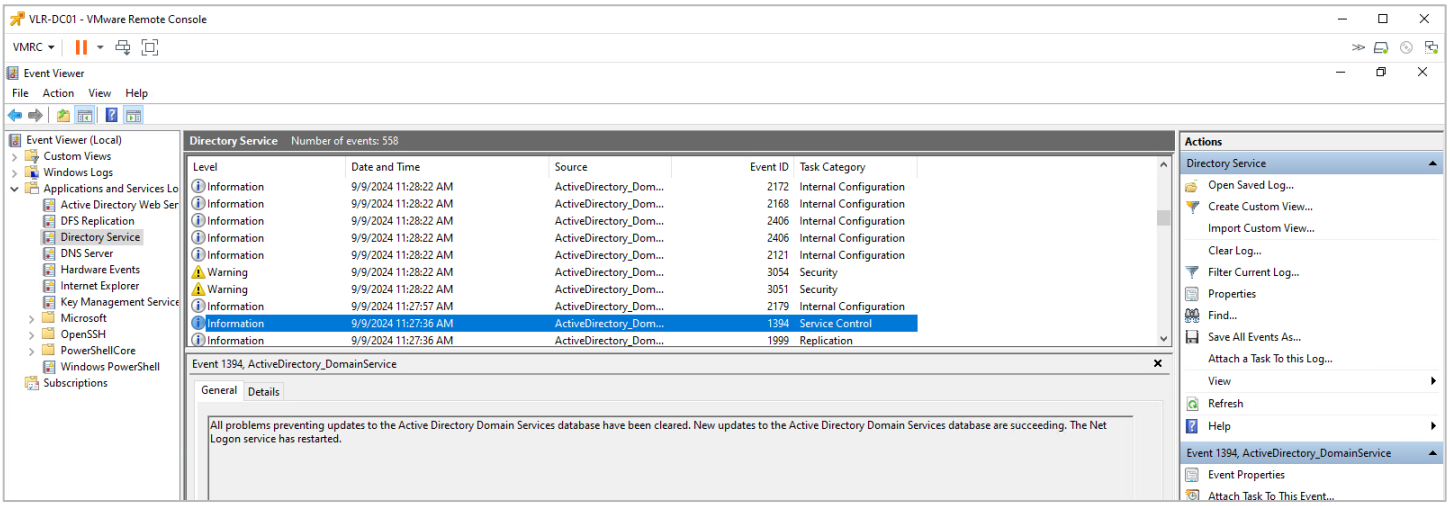

At this point, the Domain Controllers have discarded their old RID pools and obtained a new set, have a new Invocation ID, and can begin to use the new batch of USNs.

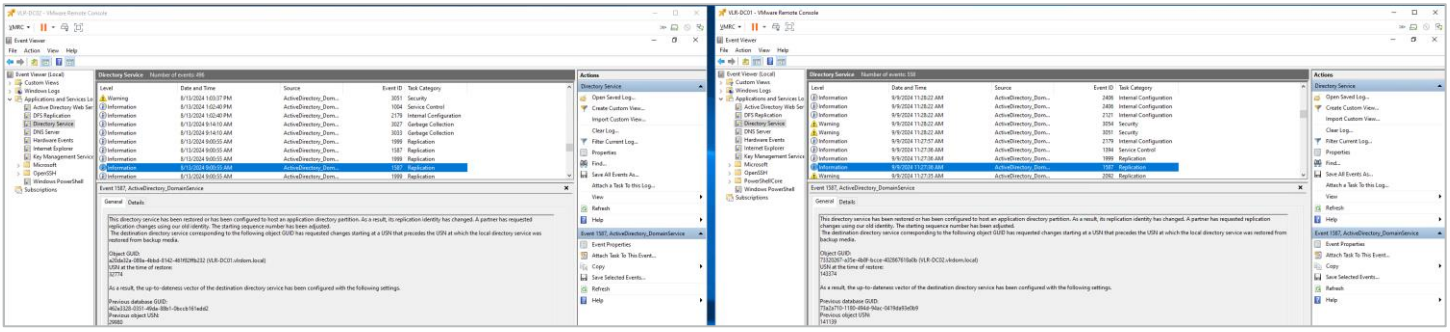

Windows has also accepted the new VM-Generation ID created for the VM by our VMware Live Site Recovery Recovery exercise. Windows will now store this for subsequent comparison next time the VM is rebooted.

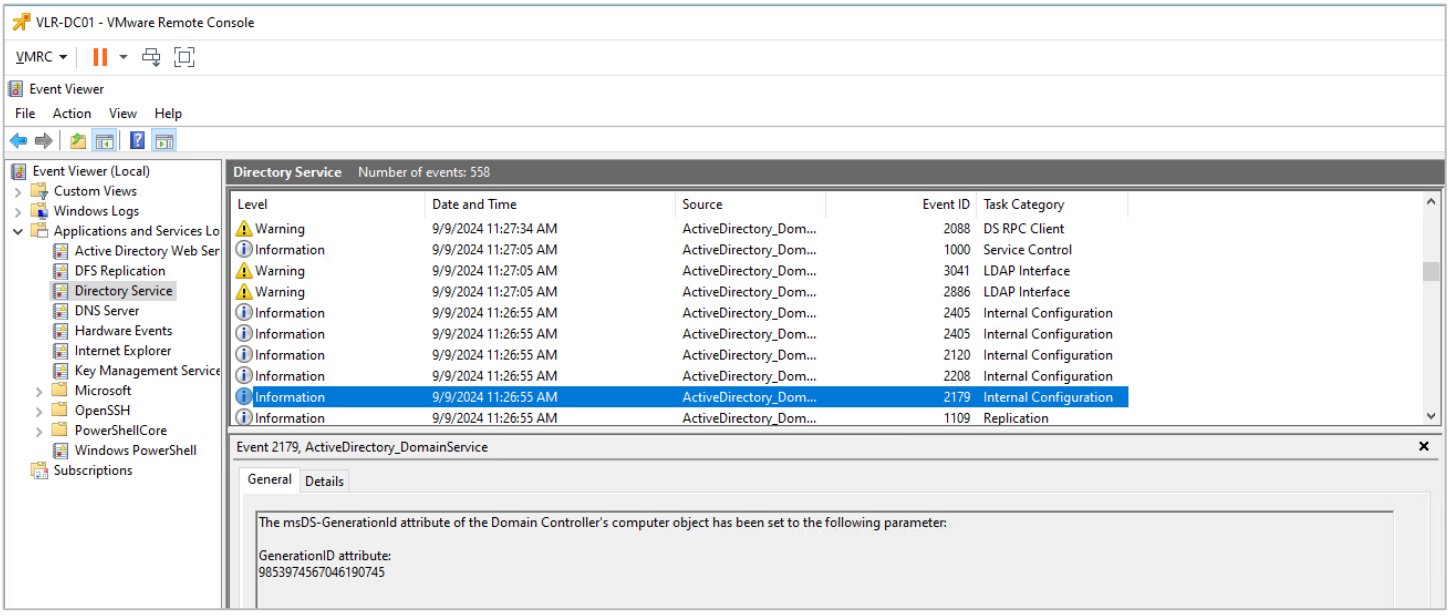

DC02 has also been successfully remediated. Because it's not the FSMO role holder, part of its healing process (for example, obtaining a new RID pool) was supported by the availability of the Role holder.

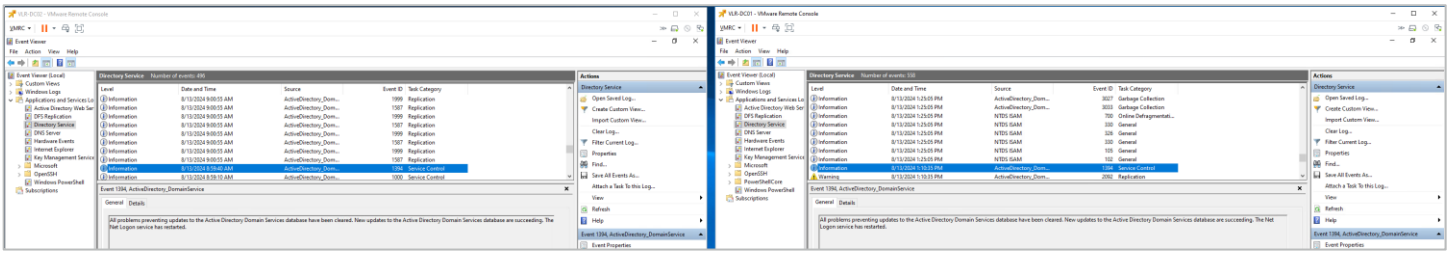

Recover the SQL Server availability group

Now that the Domain Controllers have been recovered, you're ready to recover the SQL Server availability group cluster. Remember that our objective here is to ensure that we don't just recover the individual VMs. We also want to recover the services they provide. This means that, upon recovery, the cluster service and resources (databases, jobs, scripts) also must be available, accessible, and operational.

1. Start by following the same process you did above for the Domain Controller recovery plan.

Notice the startup sequence of the two VMs in your recovery plan. VMware Live Site Recovery doesn't begin to power on Papilolo-02 and Papilolo-03 until Papilolo-01 has completed bootup and the in-guest script has been called. This is a combination of dependency and recovery priority at work.

VLR-SQL-RP EDIT MOVE DELETE TEST CLEANUP RUN REPROTECT CANCEL

Summary | Recovery Steps | Issues | History | Permissions | Protection Groups | Virtual Machines

Recovery Step	Status	Step Started	Step Completed
1. Synchronize storage	Success	Wednesday, September 11, 2024 3:56:48 PM	Wednesday, September 11, 2024 3:56:48 PM
2. Restore recovery site hosts from standby	Success	Wednesday, September 11, 2024 3:56:48 PM	Wednesday, September 11, 2024 3:56:48 PM
3. Suspend non-critical VMs at recovery site			
4. Create writable storage snapshot	Success	Wednesday, September 11, 2024 3:56:48 PM	Wednesday, September 11, 2024 3:56:52 PM
5. Configure test networks	Success	Wednesday, September 11, 2024 3:56:52 PM	Wednesday, September 11, 2024 3:56:53 PM
6. Power on priority 1 VMs	Success	Wednesday, September 11, 2024 3:56:53 PM	Wednesday, September 11, 2024 3:59:47 PM
6.1. Papilolo-01	Success	Wednesday, September 11, 2024 3:56:53 PM	Wednesday, September 11, 2024 3:59:47 PM
6.1.1. Guest startup	Success	Wednesday, September 11, 2024 3:56:53 PM	Wednesday, September 11, 2024 3:57:39 PM
6.1.2. Customize IP	Success	Wednesday, September 11, 2024 3:57:39 PM	Wednesday, September 11, 2024 3:57:43 PM
6.1.3. Guest shutdown	Success	Wednesday, September 11, 2024 3:57:43 PM	Wednesday, September 11, 2024 3:58:20 PM
6.1.4. Power on	Success	Wednesday, September 11, 2024 3:58:20 PM	Wednesday, September 11, 2024 3:58:22 PM
6.1.5. Wait for VMware tools	Success	Wednesday, September 11, 2024 3:58:22 PM	Wednesday, September 11, 2024 3:59:09 PM
6.1.6. Command: Reconfigure AG VIP	Success	Wednesday, September 11, 2024 3:59:09 PM	Wednesday, September 11, 2024 3:59:47 PM
7. Power on priority 2 VMs	Success	Wednesday, September 11, 2024 3:56:53 PM	Wednesday, September 11, 2024 4:00:41 PM
7.1. Papilolo-02	Success	Wednesday, September 11, 2024 3:56:53 PM	Wednesday, September 11, 2024 4:00:30 PM
7.1.1. Guest startup	Success	Wednesday, September 11, 2024 3:56:53 PM	Wednesday, September 11, 2024 3:57:40 PM
7.1.2. Customize IP	Success	Wednesday, September 11, 2024 3:57:40 PM	Wednesday, September 11, 2024 3:57:44 PM
7.1.3. Guest shutdown	Success	Wednesday, September 11, 2024 3:57:44 PM	Wednesday, September 11, 2024 3:58:30 PM
7.1.4. Power on	Success	Wednesday, September 11, 2024 3:59:47 PM	Wednesday, September 11, 2024 3:59:50 PM
7.1.5. Wait for VMware tools	Success	Wednesday, September 11, 2024 3:59:50 PM	Wednesday, September 11, 2024 4:00:30 PM
7.2. Papilolo-03	Success	Wednesday, September 11, 2024 3:56:53 PM	Wednesday, September 11, 2024 4:00:41 PM
7.2.1. Guest startup	Success	Wednesday, September 11, 2024 3:56:53 PM	Wednesday, September 11, 2024 3:57:40 PM
7.2.2. Customize IP	Success	Wednesday, September 11, 2024 3:57:40 PM	Wednesday, September 11, 2024 3:57:44 PM
7.2.3. Guest shutdown	Success	Wednesday, September 11, 2024 3:57:44 PM	Wednesday, September 11, 2024 3:58:30 PM
7.2.4. Power on	Success	Wednesday, September 11, 2024 3:59:47 PM	Wednesday, September 11, 2024 3:59:50 PM
7.2.5. Wait for VMware tools	Success	Wednesday, September 11, 2024 3:59:50 PM	Wednesday, September 11, 2024 4:00:41 PM
8. Power on priority 3 VMs			
9. Power on priority 4 VMs			
10. Power on priority 5 VMs			

The recovery is complete.

Live Site Recovery

Site Pair | Replications | Protection Groups | Recovery Plans

Search...

Recovery Plans
- VLR-DCs-RP
- VLR-SQL-RP

VLR-SQL-RP EDIT MOVE DELETE TEST CLEANUP RUN REPROTECT CANCEL

Summary | Recovery Steps | Issues | History | Permissions | Protection Groups | Virtual Machines

EXPORT STEPS | TEST | CLEANUP | RUN | REPROTECT | CANCEL

Plan status: Test complete

Description: The virtual machines have been recovered in a test environment at the recovery site. Revie

Recovery Step	Status	Step Started	Step Completed
1. Synchronize storage	Success	Wednesday, September 11, 2024 3:56:48 PM	Wednesday, September 11, 2024 3:56:48 PM
2. Restore recovery site hosts from standby	Success	Wednesday, September 11, 2024 3:56:48 PM	Wednesday, September 11, 2024 3:56:48 PM
3. Suspend non-critical VMs at recovery site			
4. Create writable storage snapshot	Success	Wednesday, September 11, 2024 3:56:48 PM	Wednesday, September 11, 2024 3:56:52 PM
5. Configure test networks	Success	Wednesday, September 11, 2024 3:56:52 PM	Wednesday, September 11, 2024 3:56:53 PM
6. Power on priority 1 VMs	Success	Wednesday, September 11, 2024 3:56:53 PM	Wednesday, September 11, 2024 3:59:47 PM
7. Power on priority 2 VMs	Success	Wednesday, September 11, 2024 3:56:53 PM	Wednesday, September 11, 2024 4:00:41 PM
8. Power on priority 3 VMs			

2. Log into Windows and check the SQL Server availability group cluster.

The Windows Server Failover Cluster supporting the SQL Server availability group should be fully functional.

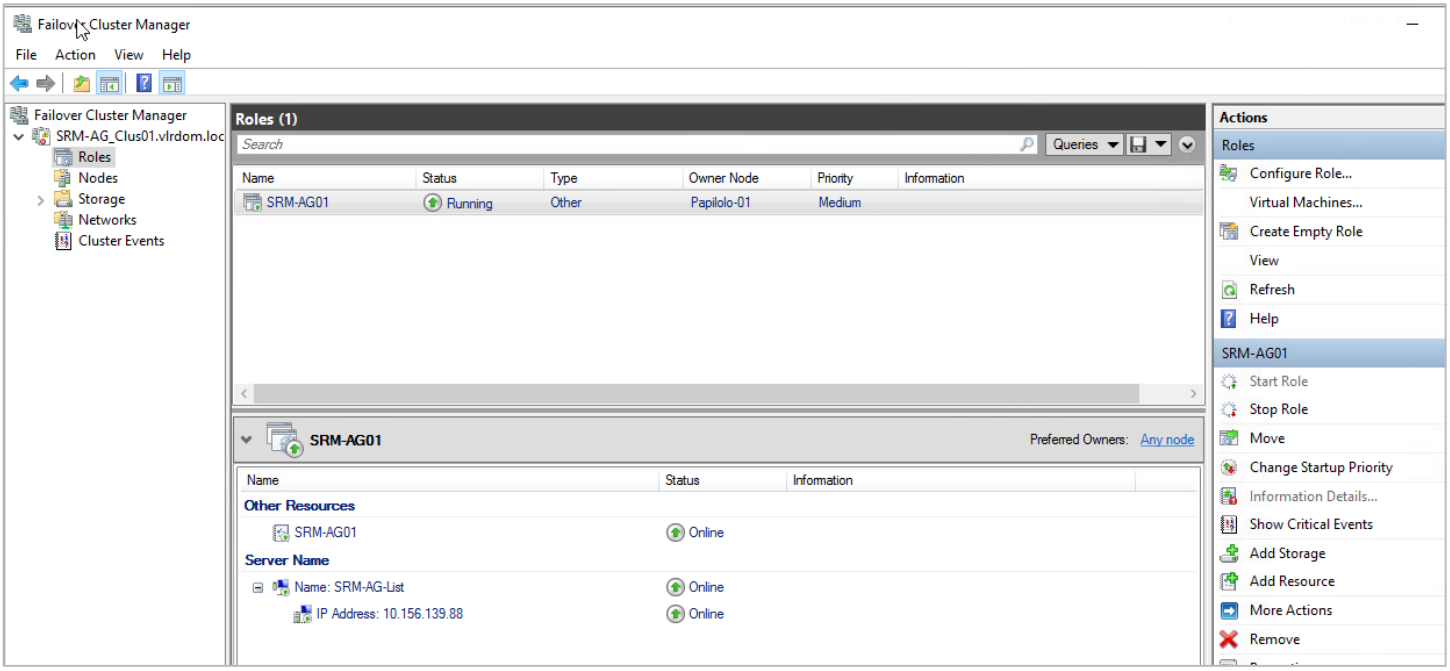

All the participating nodes should be up and operational.

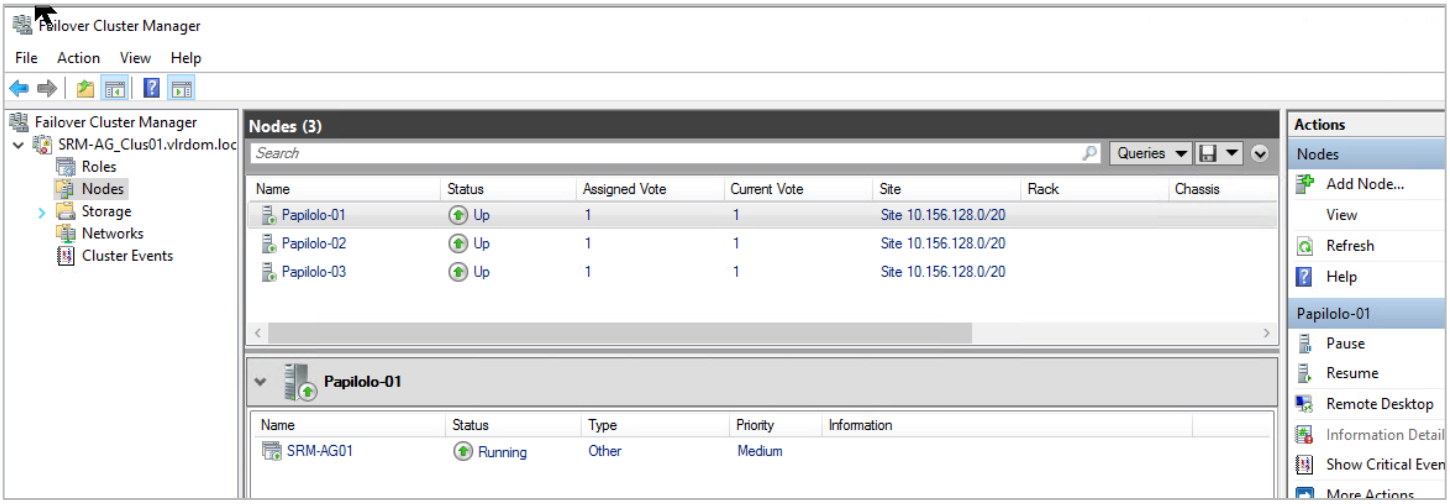

The Microsoft SQL Server listener resource should also be up and available.

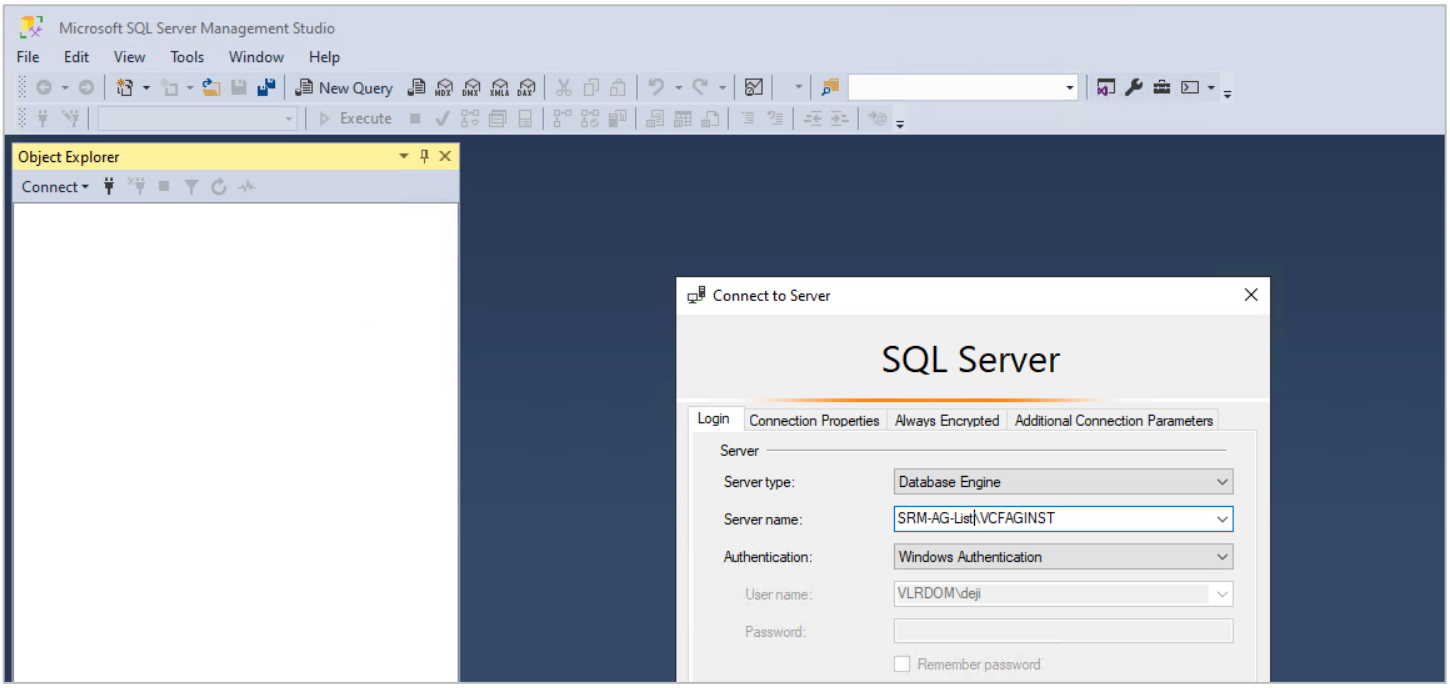

3. Because this is a test recovery exercise, confirm that your production SQL Server cluster is still up and functional at the protected site.

Here's the production SQL Server instance and its recovered copy, side-by-side. They should both be operational and processing queries.

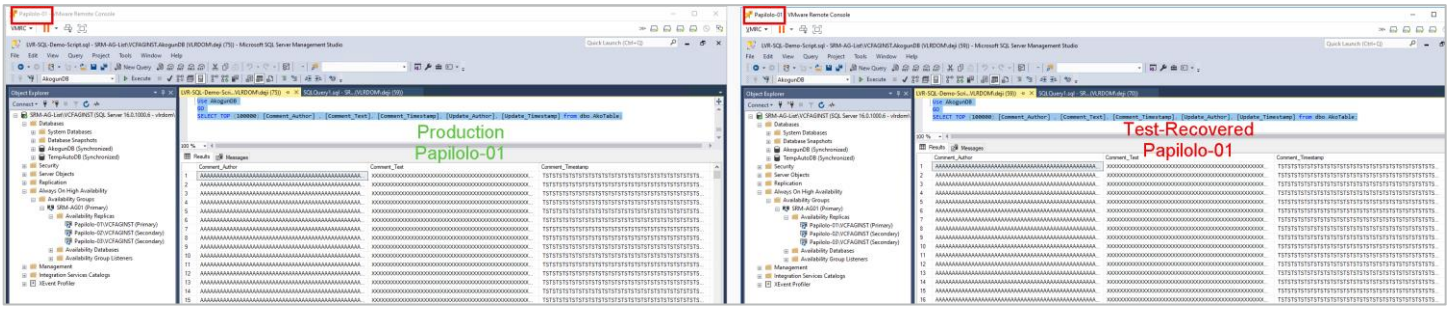

In a test recovery exercise, the recovered workloads should not be able to communicate with the production environment because they are recovered into the VMware Live Site Recovery test network we specified in previous steps. However, all workloads recovered into this test network can communicate with each other. This allows the admins and operators to more robustly test and verify the integrity of the recovery process and ascertain the availability and accessibility of their services.

Admins can fail over a client VM to the test network in VMware Live Site Recovery, or they can connect a regular client VM to the fenced-off port group in vCenter and connect to the recovered workloads to perform any validation or integrity tests required for their disaster recovery plans.

You have successfully performed test recovery of the recovery plan. If there were any failures, misconfigurations, or unexpected behaviors, you could correct them by editing the plan and retesting the changes without disrupting services in production.

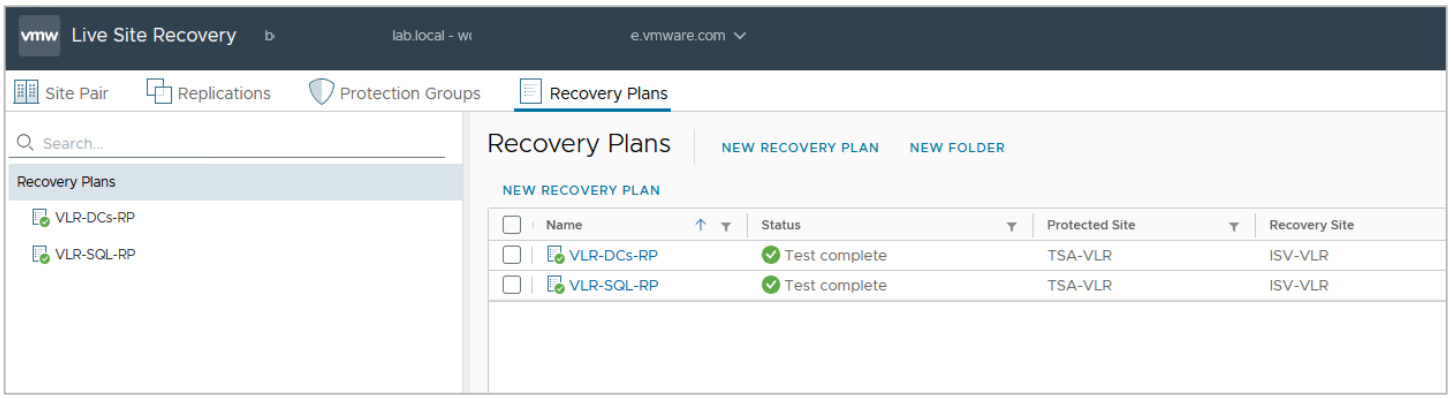

Clean up after the test recovery

Now that you're done with the test recovery, you need to clean up the test environment.

1. Select each tested recovery plan and click **Cleanup**.

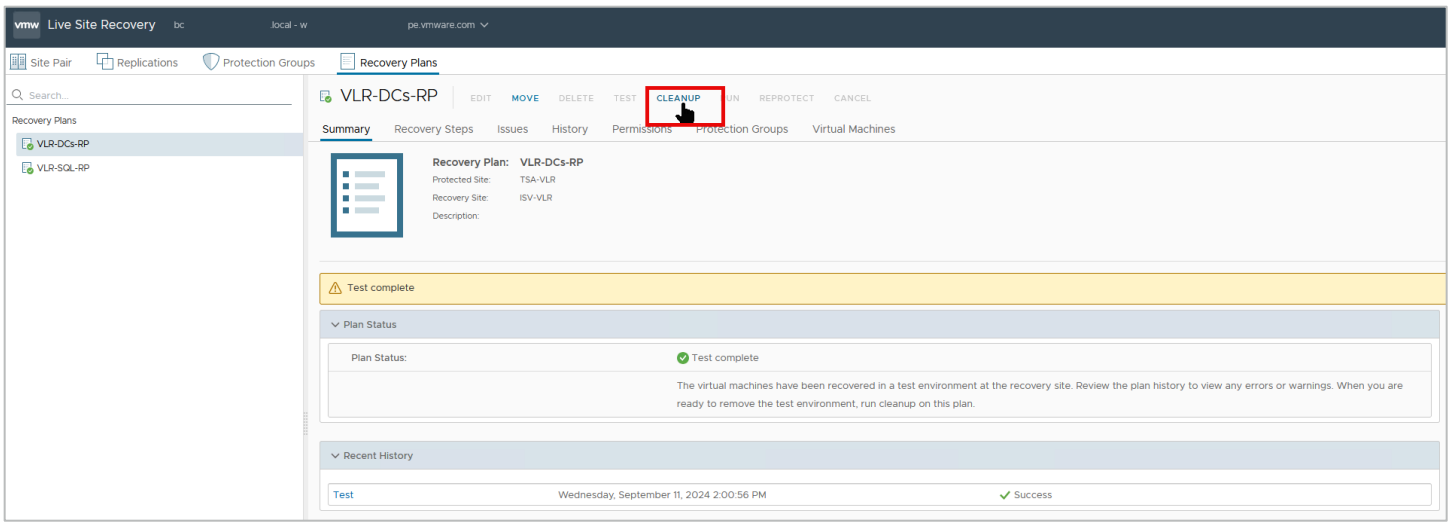

2. Click **Next** to confirm.

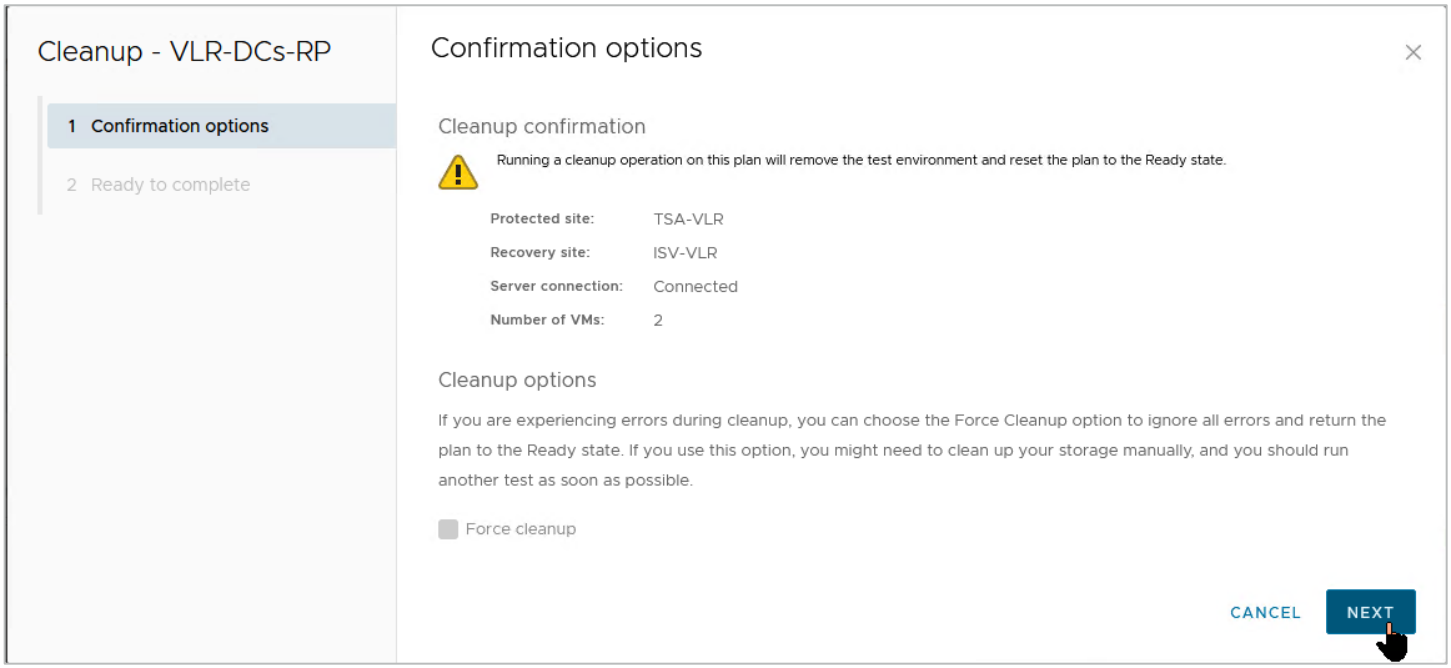

3. Click **Finish** to commit the changes.

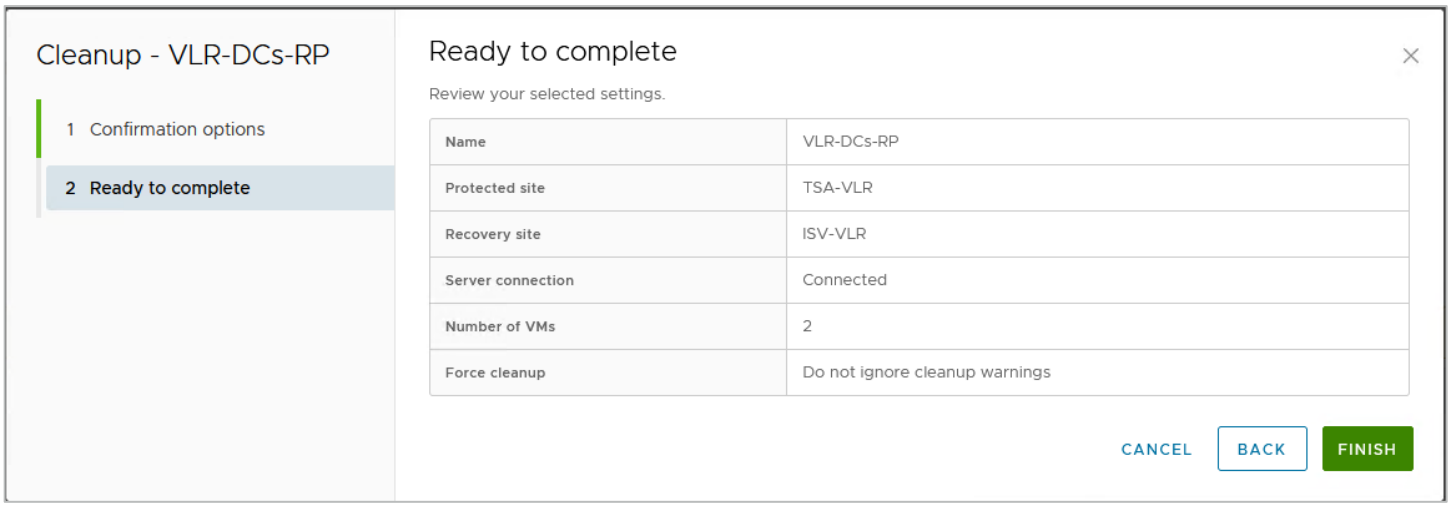

The recovery plans have returned to **Ready** for another test or invoke in an actual disaster event.

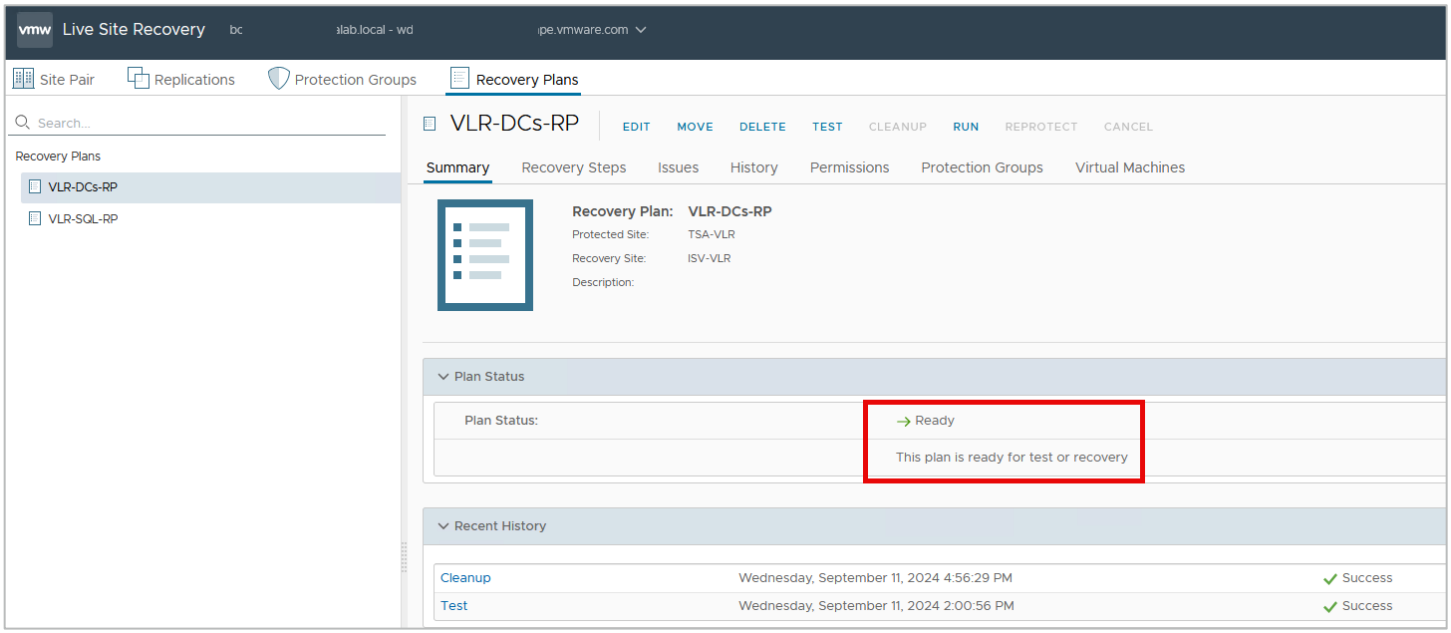

Also, the recovered VMs are powered off and restored to their previous placeholder states.

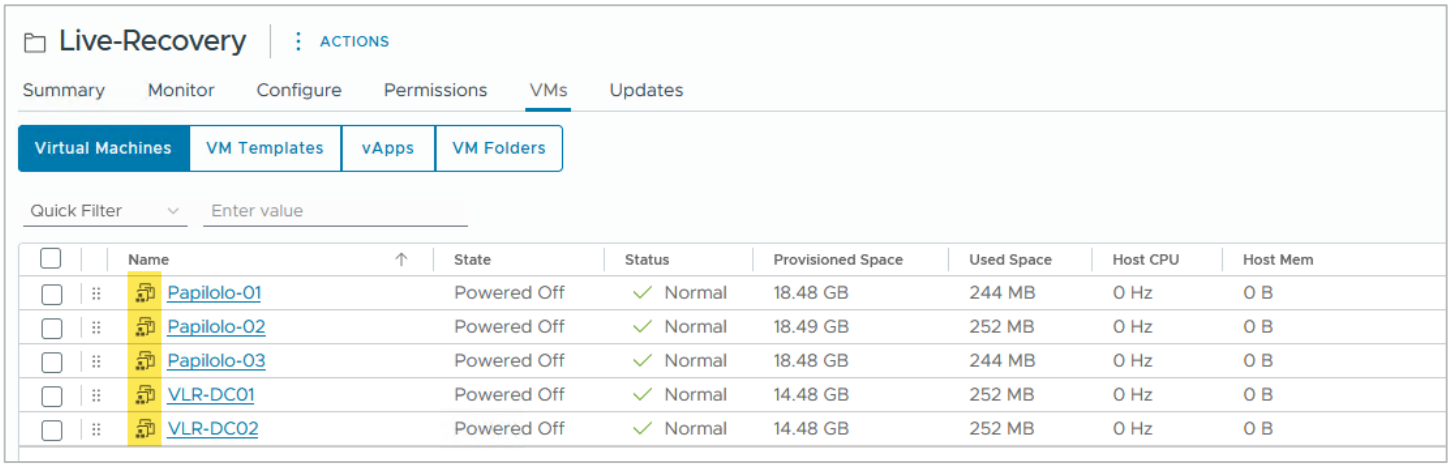

Perform a real disaster recovery

Performing mocked-up or simulated disaster recovery exercises is one of the best features of VMware Live Site Recovery. Knowing that you're adequately prepared to recover an infrastructure in real disaster events gives you peace of mind. It also helps your organization satisfy compliance, regulatory, and other legal requirements. A simulated failure and recovery isn't usually the desired outcome for investment in a robust BCDR solution like VMware Live Site Recovery, though. What the solution can do for you in a real disaster event is always the end goal. We'll now demonstrate VMware Live Site Recovery's capabilities in a disaster event.

A disaster event is a catastrophic event that impacts IT services in a production environment. It implies that all servers and services in that specific environment are unavailable and must be reinstantiated or reinstated in another environment for business continuity.

Except for a few considerations and cosmetic differences, the process of performing real disaster recovery exercises is not much different from the test disaster recovery process you previously conducted. We'll highlight those differences in this section.

1. Select the recovery plan and click **Run** to initiate a disaster recovery exercise.

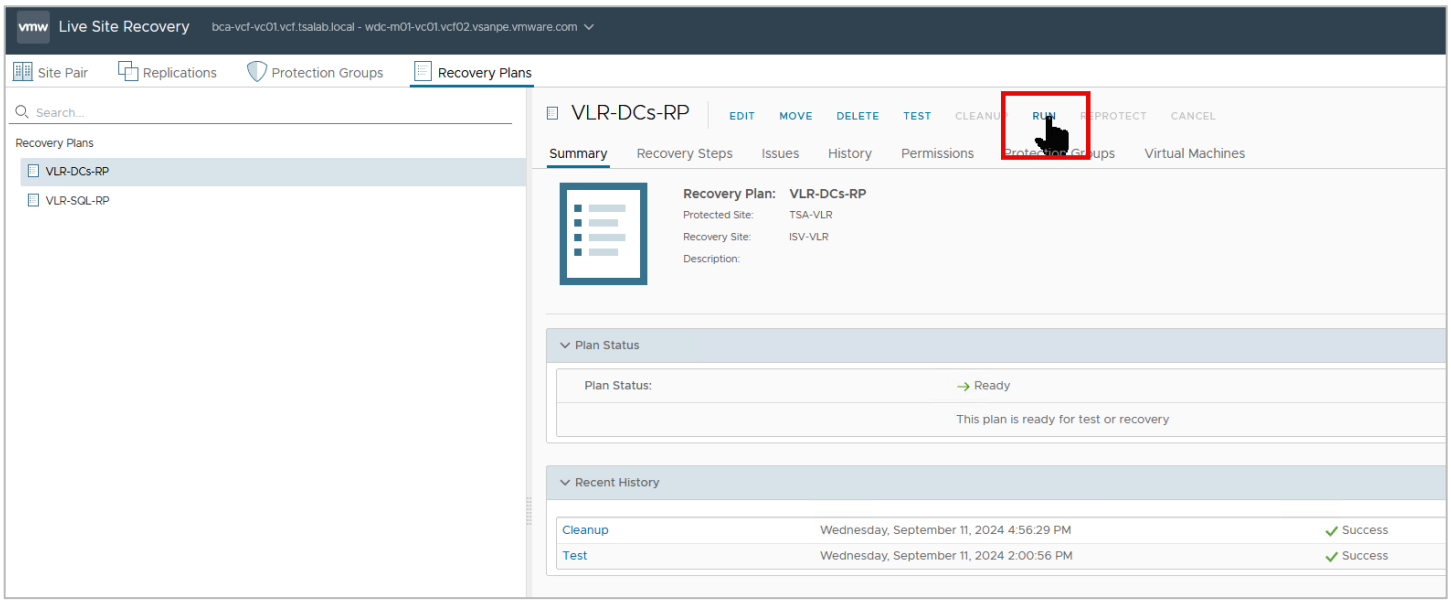

VMware Live Site Recovery provides two types of disaster recovery operations:

- **Planned recovery:** This is good for proactively relocating business-critical workloads from one datacenter to another for any business reasons. For example, if a natural disaster event is predicted for the area where the workloads are currently located, you can invoke the recovery plans to move them to another site in a controlled fashion. In this mode, the recovery operation will (among other things) perform an up-to-date synchronization between the two sites to ensure that changes in flight are committed to the replicated copies of the workloads at the recovery site. The process will also attempt to power off the workloads at the protected site to avoid service collision. If these attempts fail, the recovery will be stopped.
- **Disaster recovery:** This is for situations where the workloads at the protected sites are no longer available. When this option is invoked, VMware Live Site Recovery attempts a last-minute replication and a controlled power-off of the VMs at the protected site. The recovery continues even if VMware Live Site Recovery cannot successfully perform these steps. When you indicate a disaster recovery, the system assumes an actual disaster event makes the protected site unreachable, and the services or servers there are unavailable.

2. Select the checkbox to acknowledge you understand the action is disruptive. If you miss this checkbox, you're prompted to acknowledge the disaster recovery to make sure you don't accidentally initiate one.

3. Select **Disaster recovery**.
4. Click **Next**.

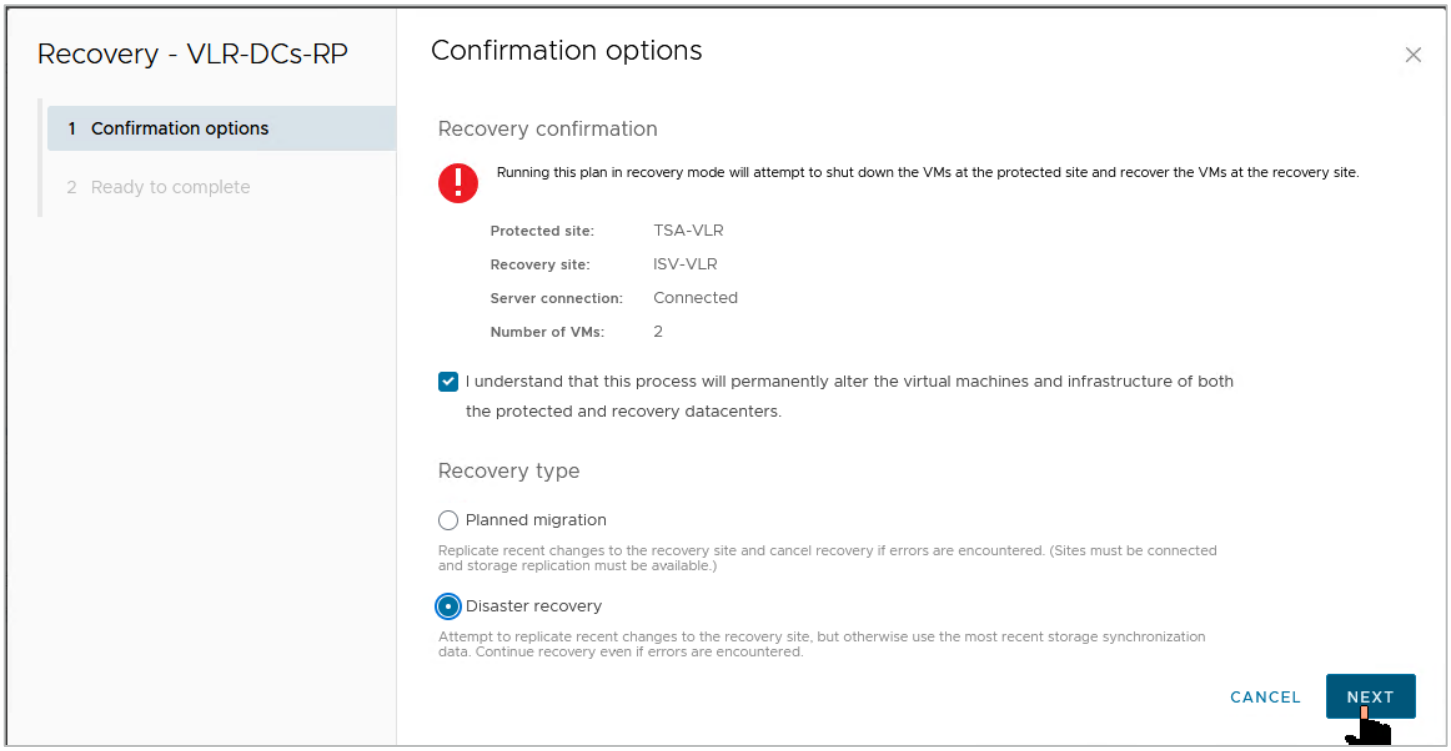

5. Click **Finish** to begin the disaster recovery.

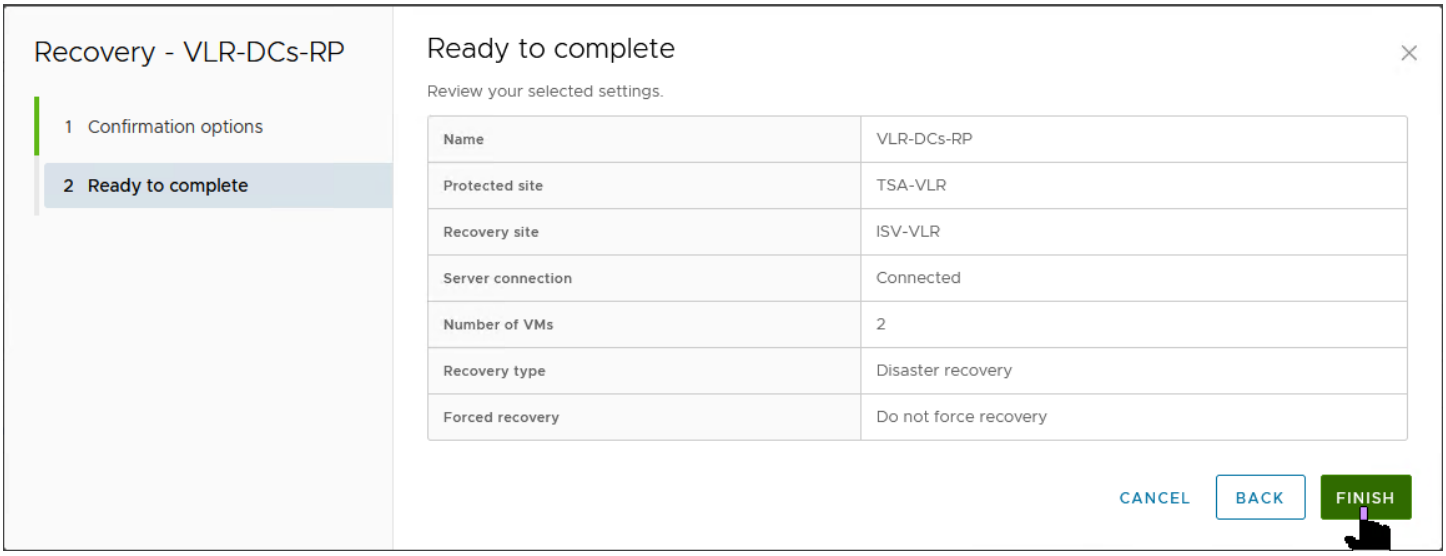

Here, you'll see VMware Live Site Recovery powering off the protected VMs at the protected site before it starts to recover them at the recovery site. The power-off and synchronization attempts succeeded because your

protected site wasn't offline. If it were, these tasks wouldn't succeed, and the recovery process would continue.

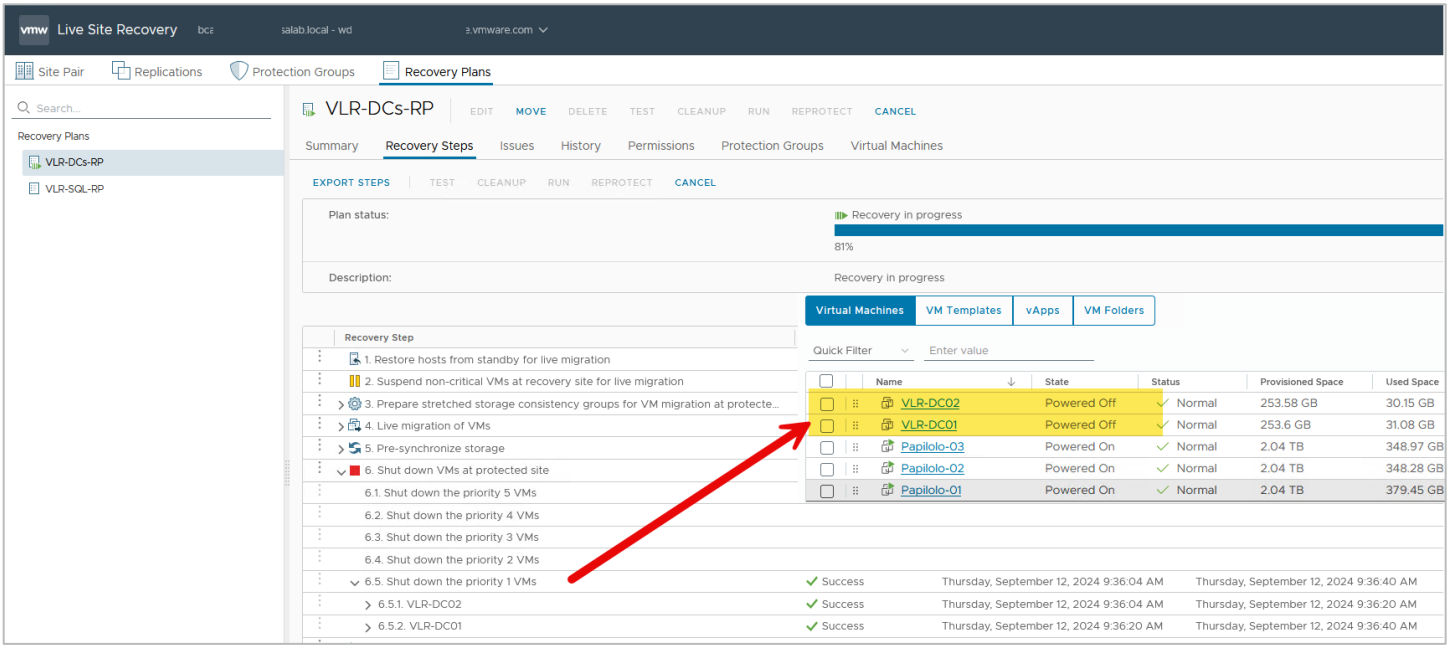

The recovery was successfully completed.

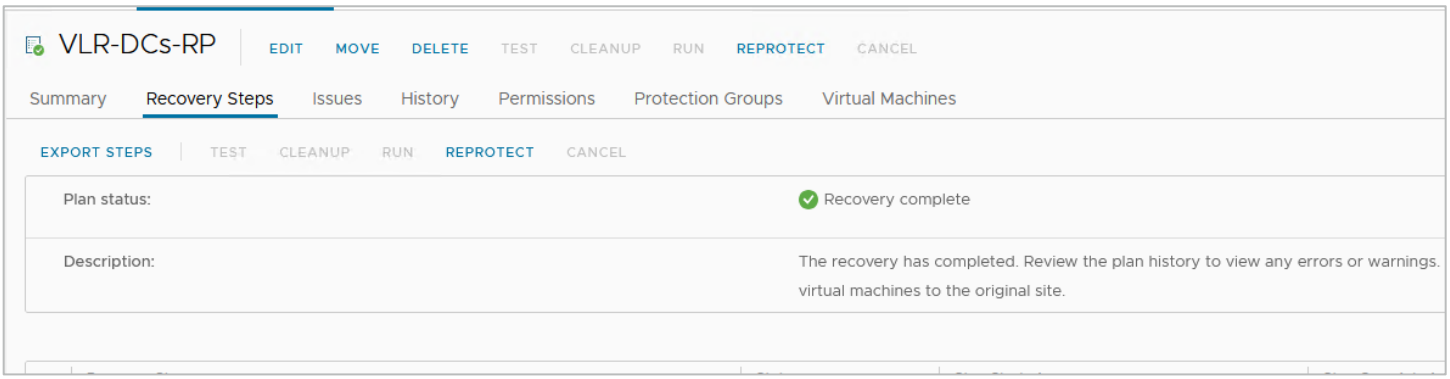

The Domain Controllers in your recovery plan are now running and providing services in the recovery site. Business continuity is restored with just a few mouse clicks.

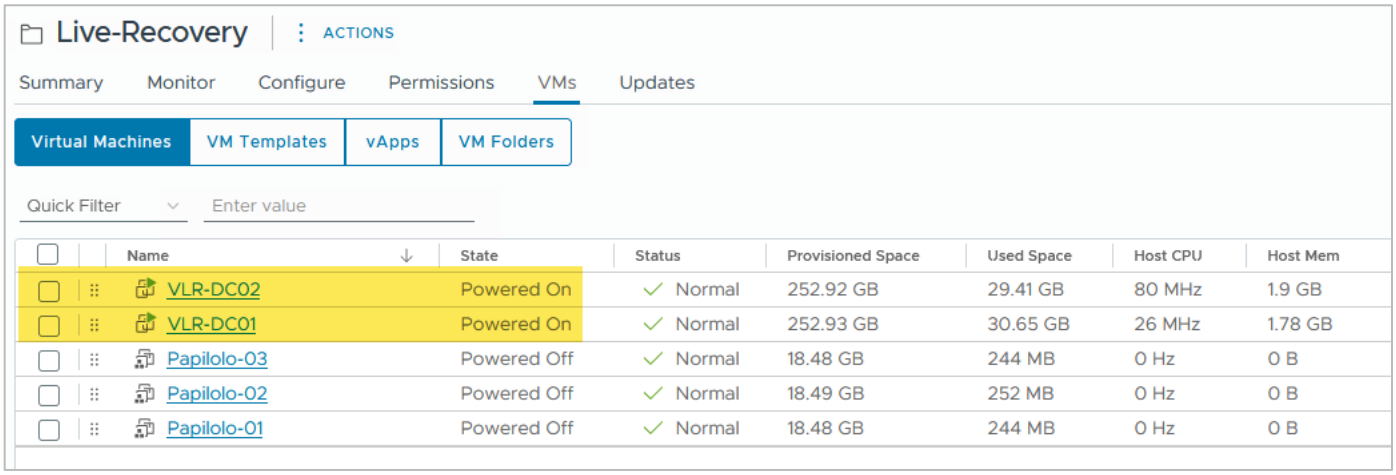

6. Go ahead and invoke the rest of your recovery plans.

The VMs are now powered on at the recovery site and powered off at the protected site.

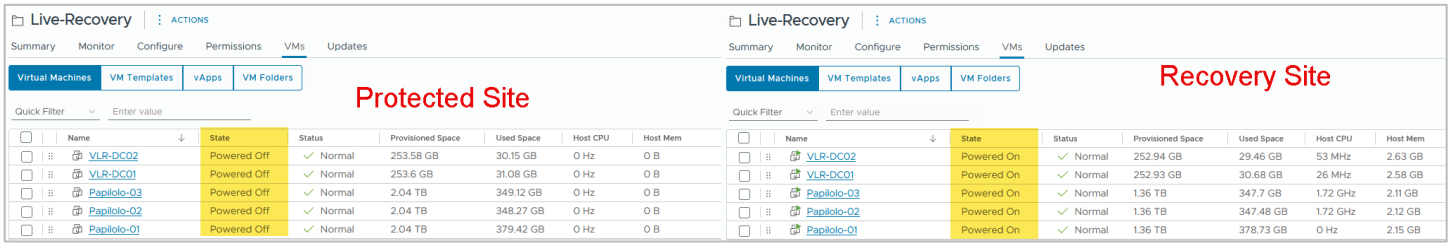

Reprotect business-critical applications after a disaster event

What is **Reprotect needed** as shown in the **Recovery Plans** tab?

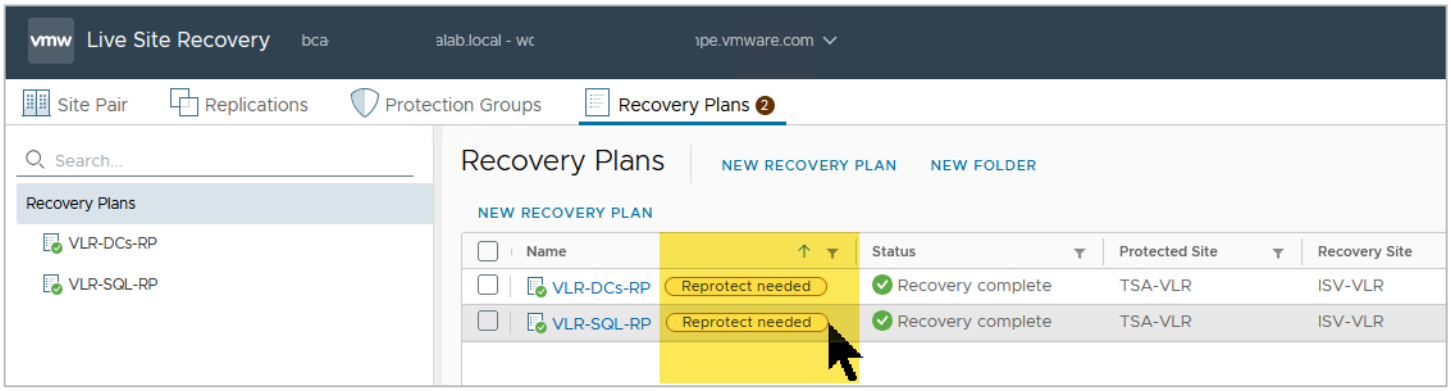

VMware Live Site Recovery makes it easy to quickly configure protection for the VMs after a disaster recovery operation. In the immediate aftermath of a real disaster event, the recovered VMs don't have any protection (because the original site is unavailable). After the disaster is over and you're ready to resume operations at that site, you simply reprotect the VMs.

7. Select the recovery plan containing the VMs you want to protect.

8. Click on **Reprotect**.

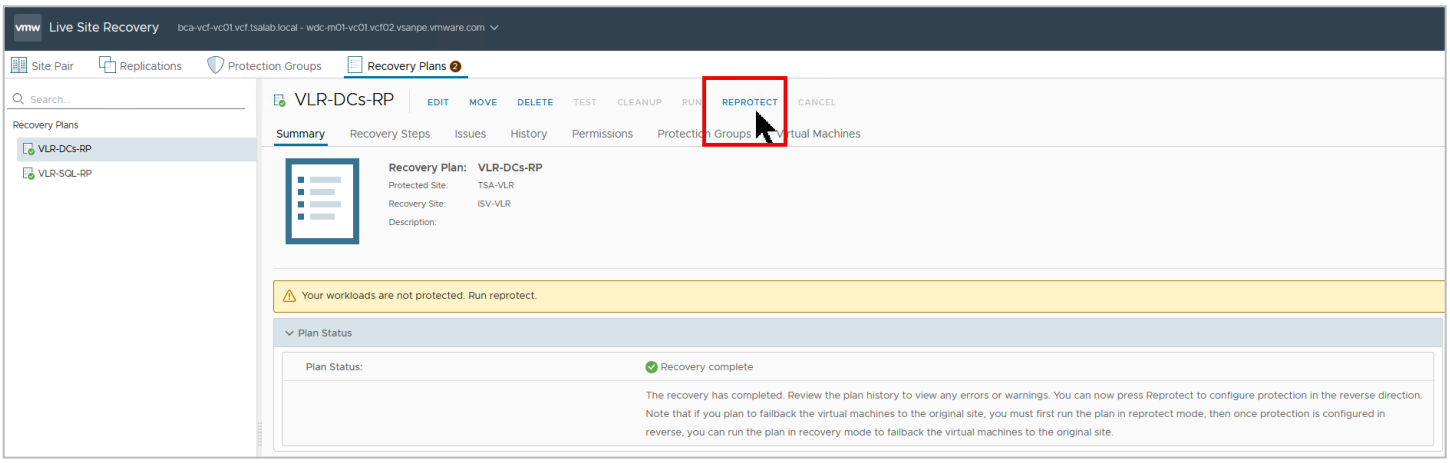

You'll notice that the source → target direction has now been automatically reversed. The original recovery site is now the protected site (and vice versa) because the VMs are now running at the original recovery site.

9. Select **I understand that this operation cannot be undone** to signal that you understand the effects and implications of the action you're about to perform.

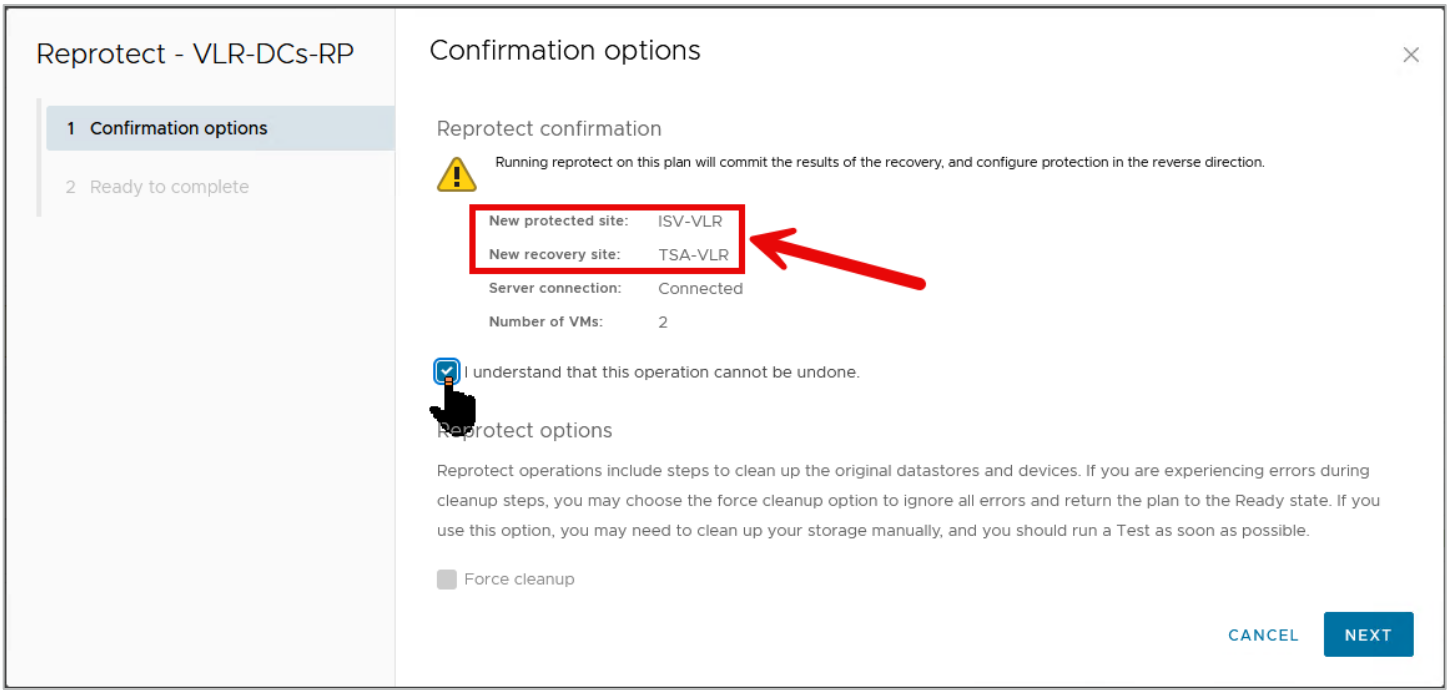

10. Click **Next** and then **Finish** on the next screen.

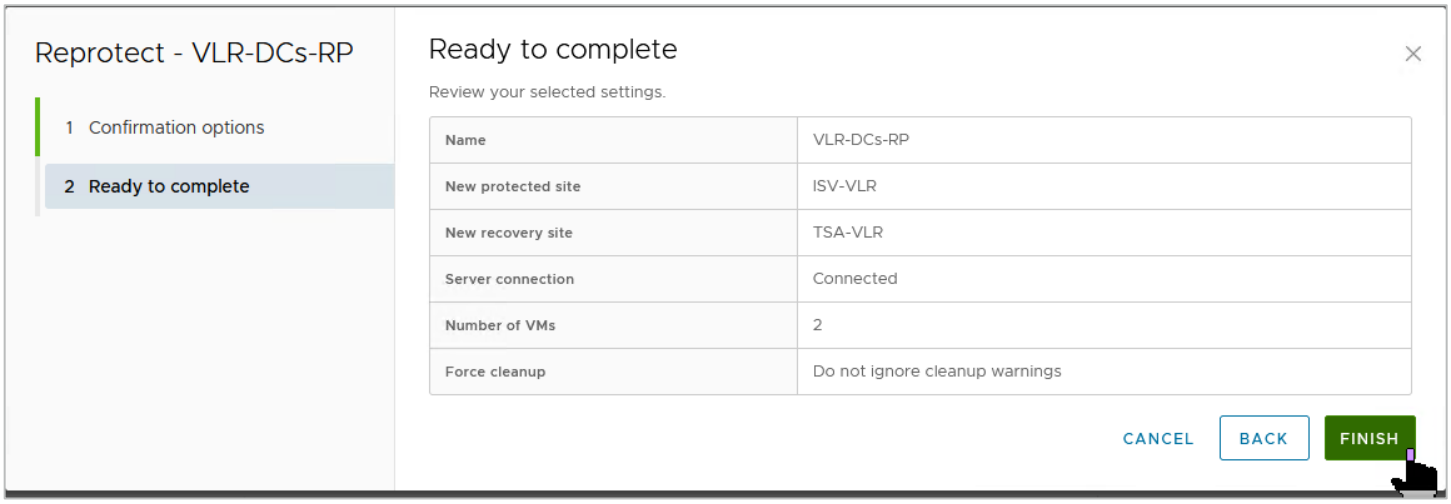

When the VMs are reprotected, they're converted to placeholders at the new recovery site (formerly the protected site).

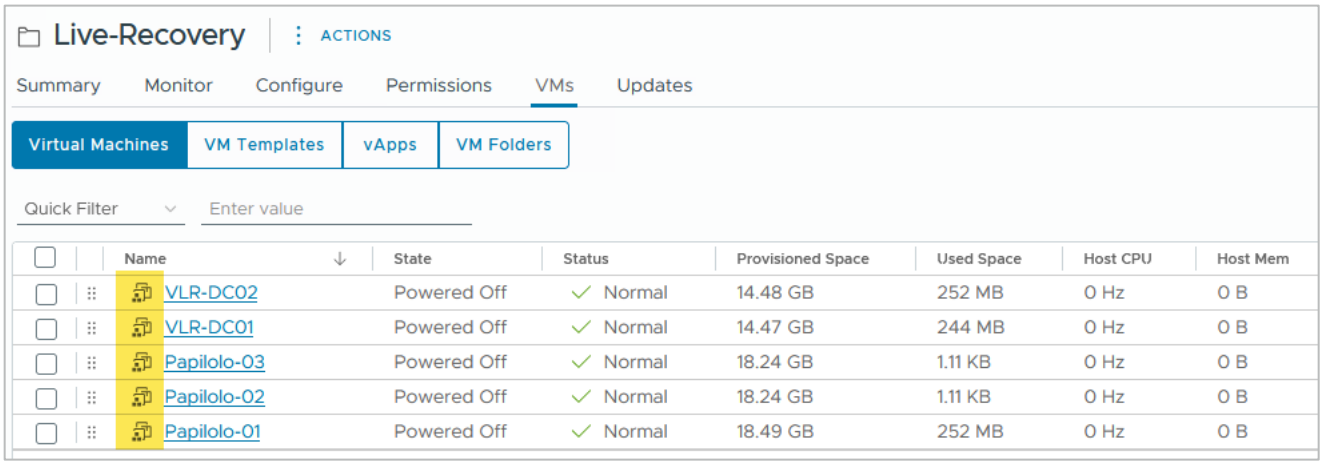

Modify the in-guest script after a disaster recovery operation

You previously configured a **Run Command on Recovered VM** task in the **Post Power On Steps** section for the Domain Controller and SQL Server recovery plans.

We called a script to reboot the Domain Controller recovery plan. No modification is necessary for this step when we reprotect these VMs. However, the SQL Server recovery plan deserves some attention because the script must make site/subnet-specific configuration changes to both the Windows cluster and SQL Server Always On. Therefore, you need to modify the original script with the correct information:

- Now that the VM is running in the recovery site, you can log in and edit the script itself.
- Or, you can edit the recovery [lan and specify a different script to be used in the post power-on steps as you did previously. You'll do this process below.

11. Connect and log into VMware Live Site Recovery on the new recovery site.
12. Select the recovery plan, click the **Virtual Machines** tab, and then select the VM with the recovery steps to modify.
13. Click **Configure Recovery.**

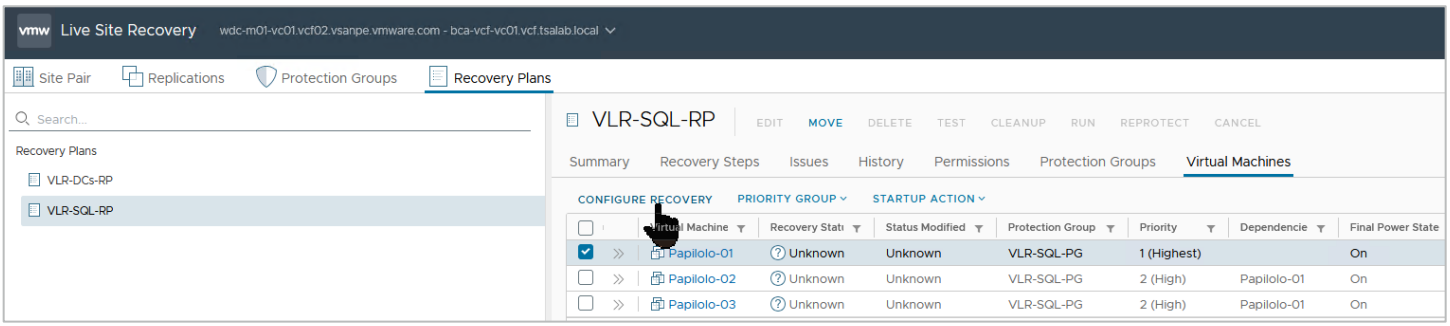

14. Select the step to modify (here, it's **Post Power On Steps**), and then click **Edit**.

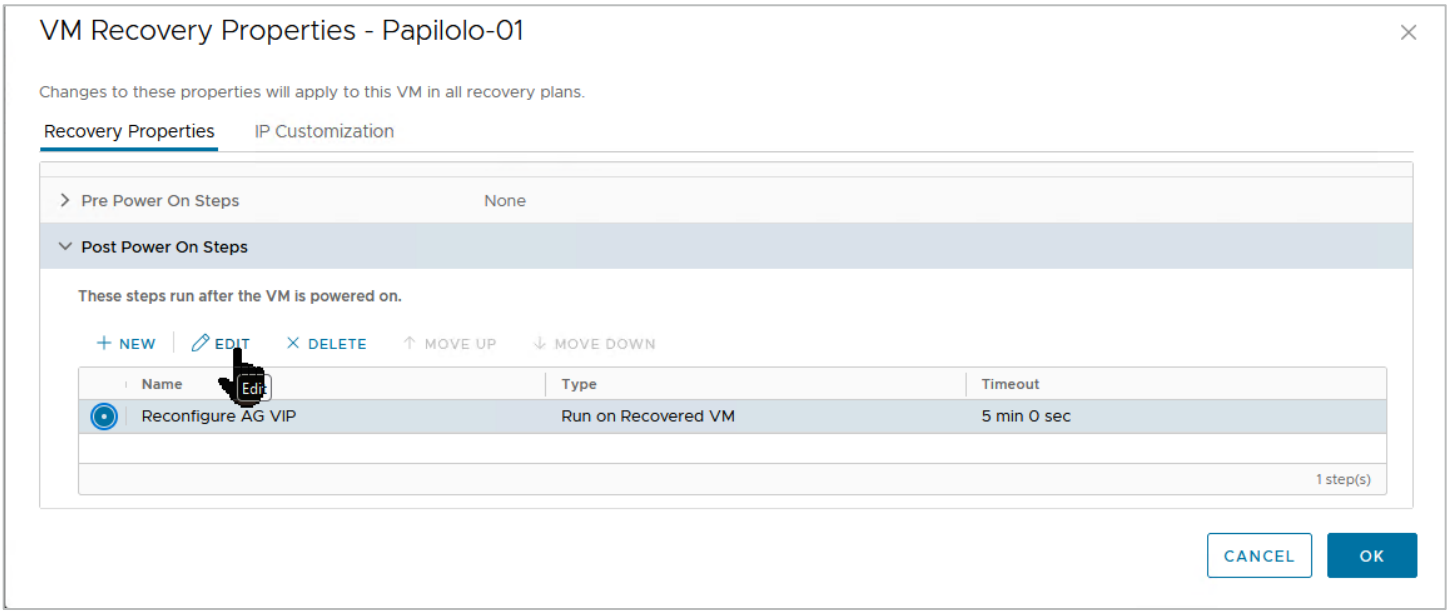

15. Type in the command to run. In our case, we're calling another scrip: **`Change-Cluster-AG-VIP-Reverse.ps1`.** This is also located in the VM.
16. Click **Save**.

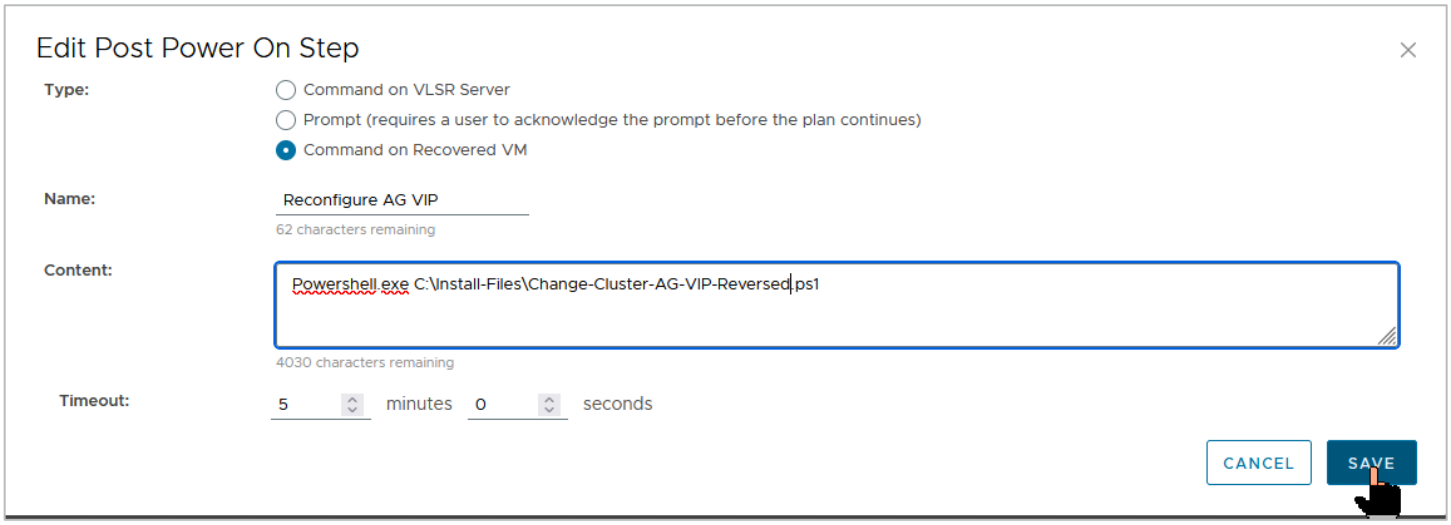

17. Click **OK** to commit the changes.

Here's the script we used for this exercise. It also appears in the appendix.

Change-Cluster-AG-VIP-Reversed.ps1

```
# Change-Cluster-AG-VIP.ps1 (For reconfiguring recovered MS SQL Server cluster properties)
Import-Module FailoverClusters

# Let's Force-Start our Cluster first

# Immediately post-recovery, the whole Cluster is down

Start-ClusterNode -FQ

# Let's define our new IP address and subnet mask for the Cluster IP Address
$newClusIP = "10.156.138.87"  # Replace with your new IP address
$newClusMask = "255.255.240.0"    # Replace with your subnet mask

# Get the IP Address of the Cluster resource
$setNewClusIP = Get-ClusterResource -Name "SRM-AG01_Clus_IP"

# Set the new IP address and subnet mask for the Cluster resource
$setNewClusIP | Set-ClusterParameter -Name Address -Value $newClusIP
$setNewClusIP | Set-ClusterParameter -Name SubnetMask -Value $newClusMask

########## Next, we modify the AG VIP
# Let's define our new IP address and subnet mask for the AG VIP Address
$newAGIP = "10.156.138.88"  # Replace with your new IP address
$newAGMask = "255.255.240.0"    # Replace with your subnet mask

# Get the IP Address of the AG resource
$setNewAGIP = Get-ClusterResource -Name "SRM-AG01-IP"

# Set the new IP address and subnet mask for the AG resource
$setNewAGIP | Set-ClusterParameter -Name Address -Value $newAGIP
$setNewAGIP | Set-ClusterParameter -Name SubnetMask -Value $newAGMask

# Bring the resources offline
Stop-ClusterResource "SRM-AG01_Clus_IP"
Stop-ClusterResource "SRM-AG01_SRM-AG-List"
Stop-ClusterResource "Cluster Name"
Stop-ClusterResource "SRM-AG01"
Stop-ClusterResource "SRM-AG01-IP"

# We now start up everything
Start-ClusterResource "SRM-AG01"
Start-ClusterResource "SRM-AG01-IP"
Start-ClusterResource "SRM-AG01_SRM-AG-List"
Start-ClusterResource "SRM-AG01_Clus_IP"
Start-ClusterResource "Cluster Name"
```

Conclusion

We’ve reached the end of our demonstration on how to prepare and configure a set of virtualized business-critical application VMs in a vSphere-based infrastructure to be protected against a disaster event and recovered with VMware Live Site Recovery to restore business continuity.

We showcased a multi-tiered application stack that required special considerations. We covered using in-guest scripting to complement VMware Live Site Recovery’s automated workflow and capabilities.

We demonstrated how to use VMware Live Site Recovery to conduct:

- Test recovery operations for compliance purposes and verify our recovery plans’ reliability on demand.
- A real disaster recovery operation and reconfigure the recovered workloads to be protected again after we achieved stability.

We also provided the in-guest scripts used in these exercises in the appendix.

We hope you've found this comprehensive documentation useful for your own purposes. Thank you.

Appendix A: Sample scripts

We (Broadcom) provide the following sample scripts for illustration purposes only. We don't assure, warranty, or guarantee their suitability for your purposes and usage. We don't provide support for these scripts. We disclaim any responsibility for any adverse effect that might result from your use of these sample scripts.

Run-Post-Script.ps1

This reboots the first Domain Controller that VMware Live Site Recovery recovers.

```
Write-Output "Rebooting VM to complete recovery..." $(Get-Date) > c:\install-files\recovery.txt
shutdown -r -t 60
```

Change-Cluster-AG-VIP.ps1

This reconfigures the recovered SQL Server cluster properties.

```
Import-Module FailoverClusters

# Let's Force-Start our Cluster first

# Immediately post-recovery, the whole Cluster is down

Start-ClusterNode -FQ

# Let's define our new IP address and subnet mask for the Cluster IP Address
$newClusIP = "10.156.139.87"  # Replace with your new IP address
$newClusMask = "255.255.240.0"    # Replace with your subnet mask

# Get the IP Address of the Cluster resource
$setNewClusIP = Get-ClusterResource -Name "SRM-AG01_Clus_IP"

# Set the new IP address and subnet mask for the Cluster resource
$setNewClusIP | Set-ClusterParameter -Name Address -Value $newClusIP
$setNewClusIP | Set-ClusterParameter -Name SubnetMask -Value $newClusMask

########## Next, we modify the AG VIP
# Let's define our new IP address and subnet mask for the AG VIP Address
$newAGIP = "10.156.139.88"  # Replace with your new IP address
$newAGMask = "255.255.240.0"    # Replace with your subnet mask

# Get the IP Address of the AG resource
$setNewAGIP = Get-ClusterResource -Name "SRM-AG01-IP"

# Set the new IP address and subnet mask for the AG resource
$setNewAGIP | Set-ClusterParameter -Name Address -Value $newAGIP
$setNewAGIP | Set-ClusterParameter -Name SubnetMask -Value $newAGMask
```

```
# Bring the resources offline
Stop-ClusterResource "SRM-AG01_Clus_IP"
Stop-ClusterResource "SRM-AG01_SRM-AG-List"
Stop-ClusterResource "Cluster Name"
Stop-ClusterResource "SRM-AG01"
Stop-ClusterResource "SRM-AG01-IP"

# We now start up everything
Start-ClusterResource "SRM-AG01"
Start-ClusterResource "SRM-AG01-IP"
Start-ClusterResource "SRM-AG01_SRM-AG-List"
Start-ClusterResource "SRM-AG01_Clus_IP"
Start-ClusterResource "Cluster Name"
```

Change-Cluster-AG-VIP-Reversed.ps1

You'd use this when the recovered SQL Server VM is reprotected.

```
Import-Module FailoverClusters

# Let's Force-Start our Cluster first

# Immediately post-recovery, the whole Cluster is down

Start-ClusterNode -FQ

# Let's define our new IP address and subnet mask for the Cluster IP Address
$newClusIP = "10.156.138.87"  # Replace with your new IP address
$newClusMask = "255.255.240.0"    # Replace with your subnet mask

# Get the IP Address of the Cluster resource
$setNewClusIP = Get-ClusterResource -Name "SRM-AG01_Clus_IP"

# Set the new IP address and subnet mask for the Cluster resource
$setNewClusIP | Set-ClusterParameter -Name Address -Value $newClusIP
$setNewClusIP | Set-ClusterParameter -Name SubnetMask -Value $newClusMask

########## Next, we modify the AG VIP
# Let's define our new IP address and subnet mask for the AG VIP Address
$newAGIP = "10.156.138.88"  # Replace with your new IP address
$newAGMask = "255.255.240.0"    # Replace with your subnet mask

# Get the IP Address of the AG resource
$setNewAGIP = Get-ClusterResource -Name "SRM-AG01-IP"

# Set the new IP address and subnet mask for the AG resource
$setNewAGIP | Set-ClusterParameter -Name Address -Value $newAGIP
$setNewAGIP | Set-ClusterParameter -Name SubnetMask -Value $newAGMask

# Bring the resources offline
Stop-ClusterResource "SRM-AG01_Clus_IP"
Stop-ClusterResource "SRM-AG01_SRM-AG-List"
Stop-ClusterResource "Cluster Name"
```

```
Stop-ClusterResource "SRM-AG01"
Stop-ClusterResource "SRM-AG01-IP"

# We now start up everything
Start-ClusterResource "SRM-AG01"
Start-ClusterResource "SRM-AG01-IP"
Start-ClusterResource "SRM-AG01_SRM-AG-List"
Start-ClusterResource "SRM-AG01_Clus_IP"
Start-ClusterResource "Cluster Name"
```

References

- Installation, setup, configuration and/or administration of VMware vSphere infrastructure
- Installation, setup, configuration and/or administration of specific VMware vSphere-based cloud infrastructure
- Installation, setup, configuration and/or administration of VMware Live Site Recovery
- Virtualizing Active Directory Domain Services on VMware vSphere
- Architecting Microsoft SQL Server on VMware vSphere
- Installation, setup, configuration and/or administration of Microsoft Active Directory Domain Services or Domain Controllers
- Installation, setup, configuration and/or administration of Microsoft SQL Server, Windows Failover Cluster or Always On
- VMware vSphere Client

About the author

Deji Akomolafe – Staff Solutions Architect and Microsoft Applications Practice Leads (Broadcom – VCF Division)

Acknowledgments

A very special thanks to the following people for their invaluable contributions to this guide:

- Chen Wei - Product Marketing Management (Broadcom – VCF Division)
- Julie Brodeur – Technical Writer (Broadcom – VCF Division)
- Michael McLaughlin - Product Marketing Engineer (Broadcom – VCF Division)
- Ka Kit Wong - Product Marketing Engineer (Broadcom – VCF Division)