Configure site-level fault tolerance

Hyper-V Replica is a new feature in Windows Server 2012 that provides for a virtual machine (VM) a warm standby copy (or *replica virtual machine*) that can exist anywhere in the world. If the primary VM fails, you can manually fail over to the replica VM. Hyper-V Replica can thus provide fault tolerance for a VM even if an entire host site should go offline.

Unlike a failover cluster, Hyper-V Replica doesn't rely on shared storage between the VMs. The replica VM instead begins with its own copy of the primary VM's virtual hard disk. The primary VM then sends updates of its changes (called *replication data*) every 5 to 15 minutes, and this data is repeatedly saved by the replica VM. The replica thus remains up-to-date.

Hyper-V Replica is one of the most important features of Windows Server 2012, and there's no doubt that it will appear on the 70-417 exam. In fact, you'll probably see more than one question about it. Fortunately, it's not an especially difficult feature to understand or implement, so your study efforts in this area will likely reap large dividends on the test.

This section covers the following topic:

- Configure Hyper-V Replication

Configuring Hyper-V physical host servers

It's important to understand the sequence of steps in configuring Hyper-V Replica. The first step is to configure the server-level replication settings for *both* physical Hyper-V hosts, called the primary server and replica server. You can access these settings in Hyper-V Manager by right-clicking a host server in the navigation pane, selecting Hyper-V Settings, and then selecting Replication Configuration in the left column of the Hyper-V Settings dialog box, as shown in Figure 12-13. By default, replication is not enabled, and no options are selected or configured.

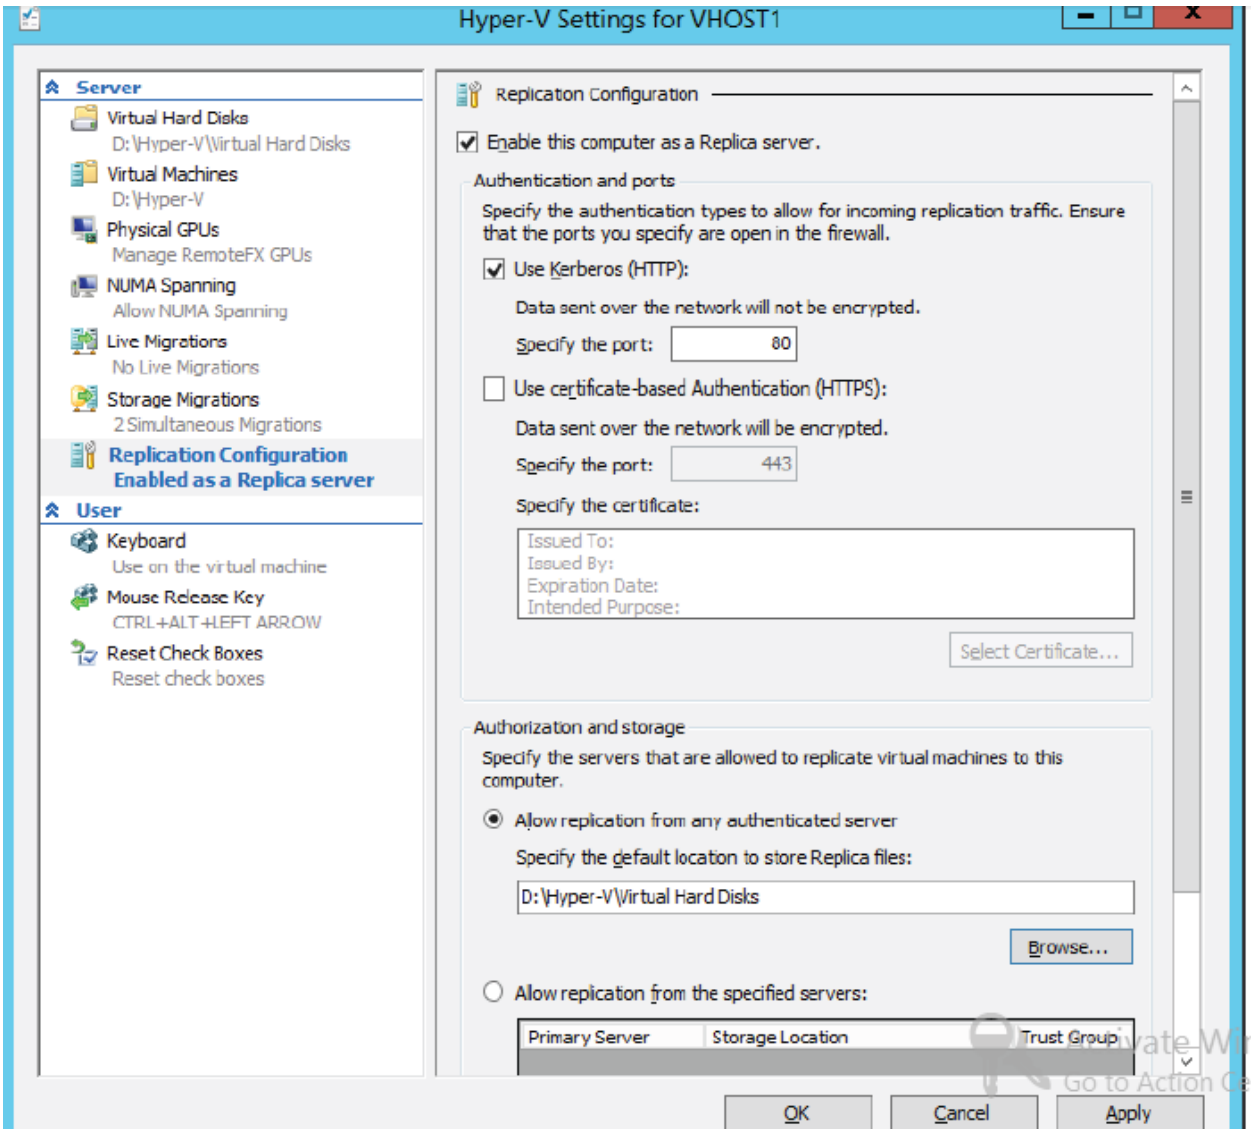

FIGURE 12-13 Host server settings for Hyper-V Replica.

To enable a physical host for Hyper-V Replica, first select the Enable This Computer As A Replica Server check box. Then, configure settings in the Authentication And Ports area and the Authorization And Storage area shown in Figure 12-13. You need to repeat these configuration steps on both primary and replica servers before configuring a VM for replication.

 Authentication And Ports In this area you choose which authentication methods you want to be available later as options when you configure a locally hosted VM for replication. You can enable Kerberos (HTTP), Certificate-Based Authentication (HTTPS), or both.

 You can enable Kerberos (HTTP) only if the local server is domain-joined. The advantage of choosing Kerberos is that it requires no further configuration. The two disadvantages are first that it doesn't encrypt data sent over the network, and second that it can be used for authentication only when the remote host server is located in a trusted domain. Note also that when you choose this authentication protocol, you need to enable the firewall rule named Hyper-V Replica HTTP Listener (TCP-In).

- You can enable Certificate-Based Authentication (HTTPS) regardless of whether the local server is domain-joined. In fact, when the local server is a standalone server, it is the only authentication protocol option. The two advantages of enabling Certificate-Based Authentication (HTTPS) are first that it encrypts replication data, and second that it allows you to replicate with a remote host when there is no trust relationship with that host through Active Directory. The disadvantage of this authentication method is that it is more difficult to configure: It requires you to provide an X.509v3 certificate for which Enhanced Key Usage (EKU) must support both Client Authentication and Server Authentication (through the Computer certificate template, for example) and that specifies (typically) the fully qualified domain name (FQDN) of the local server in the subject name field. The certificate can be self-signed or issued through a public key infrastructure (PKI). When you choose this authentication protocol, you need to enable the firewall rule named Hyper-V Replica HTTPS Listener (TCP-In).

It's important to remember that Windows Server 2012 doesn't automatically enable the firewall rules you need for the authentication protocols you choose. Depending on which protocol(s) you have enabled, you also need to enable the firewall rule "Hyper-V Replica HTTP Listener (TCP-In)", "Hyper-V Replica HTTPS Listener (TCP-In)", or both. You can enable a rule either in Windows Firewall with Advanced Security or by using the Enable-NetFirewallRule -DisplayName command in Windows PowerShell followed by the name of the rule (including quotation marks).

Remember that encrypted replication of a VM requires the host servers to have installed a certificate including both Client Authentication and Server Authentication extensions for EKU.

Authorization And Storage This area allows you to configure security settings on the local server that are used when the local server acts as a replica server. More specifically, your choice here determines the remote primary servers from which the local server will accept replication data. Even if you are configuring your local server as the primary server, the settings here are required so that—if you ever need to fail over to a remote replica—you can later fail back to the local server.
You need to choose one of two security options, both of which also provide a default path you can modify to store replication data:

 Allow Replication From Any Authenticated Server This option is somewhat less secure. When you choose this option, the local server can receive replication data from any authenticated server.

 Allow Replication From The Specified Servers This option requires you to specify the primary server(s) authorized for the local replica server. You can add multiple entries to authorize different primary servers by DNS name. To add an entry authorizing a primary server address, click Add as shown in Figure 12-14. This step opens the Add Authorization Entry dialog box shown in Figure 12-15.

For each entry, a default storage path (the middle field) is already provided, but the other two fields must be filled in manually. In the Specify The Primary Server field, you enter an FQDN that can include a wildcard character (for example, “*.adatum.com”). You also have to provide a tag called a trust group. If you want to allow replication traffic from a set of primary servers, you should assign those primary servers the same trust group name.

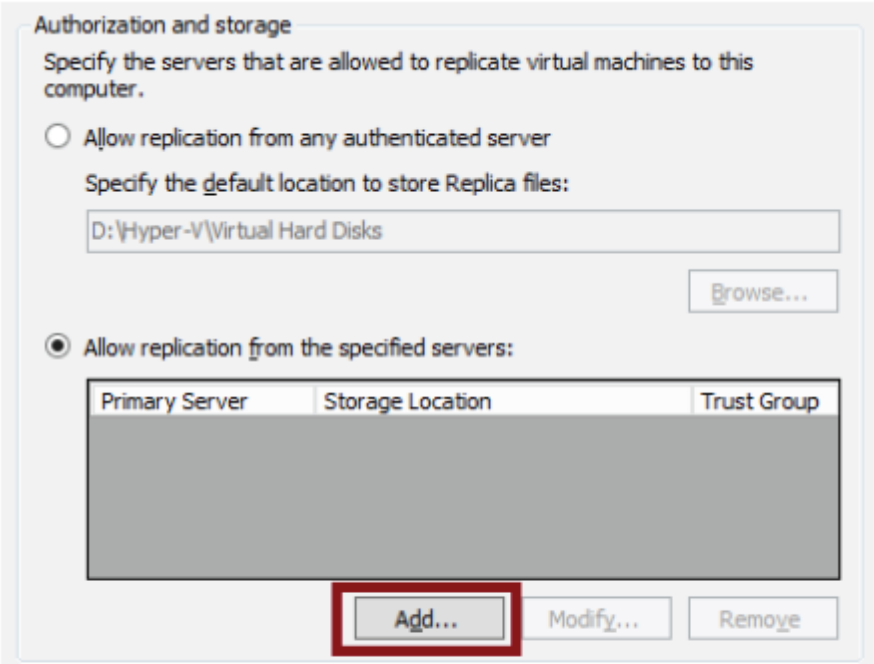

FIGURE 12-14 Authorizing primary servers for the local replica server.

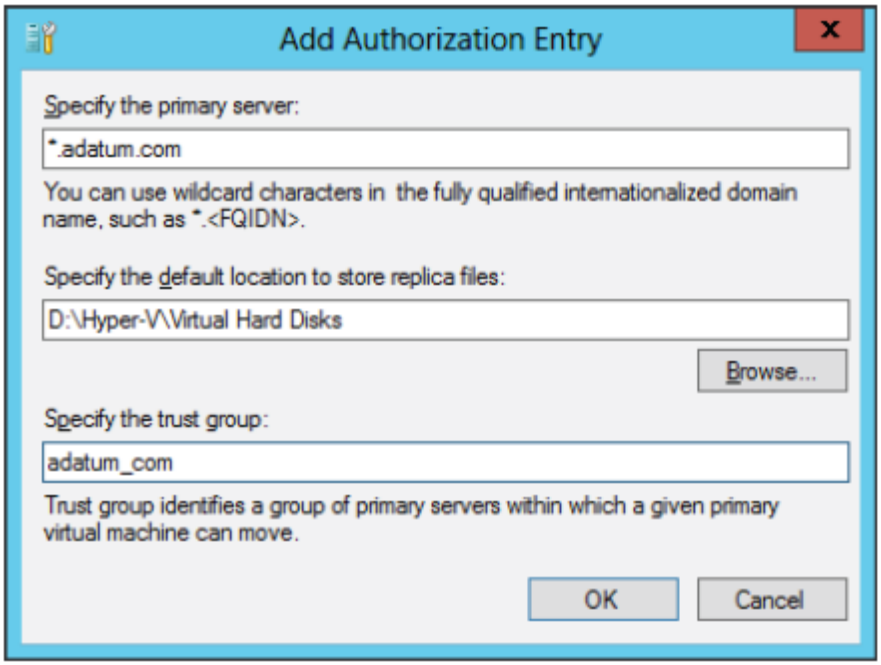

FIGURE 12-15 Adding an authorized primary server address.

How might these settings in the Authorization And Storage area appear on the 70-417 exam? One could imagine a question based on an unsuccessful failover. In such a question, authorization settings might not be configured at all on the replica server. Or, the FQDN provided in the Specify The Primary Server field in Figure 12-15 might be configured incorrectly, and the correct answer fixes that problem. Another possible question could involve a new organizational requirement that security be tightened on a replica server. Incorrect answer choices might refer to IPSec or other security-tightening methods, but the correct answer will refer to adding an authorization entry on the replica server.

Configuring VMs

After you configure both physical host servers, the next step in configuring Hyper-V Replica is to configure the chosen VM for replication on the primary server. Begin by right-clicking the VM and selecting Enable Replication, as shown in Figure 12-16.

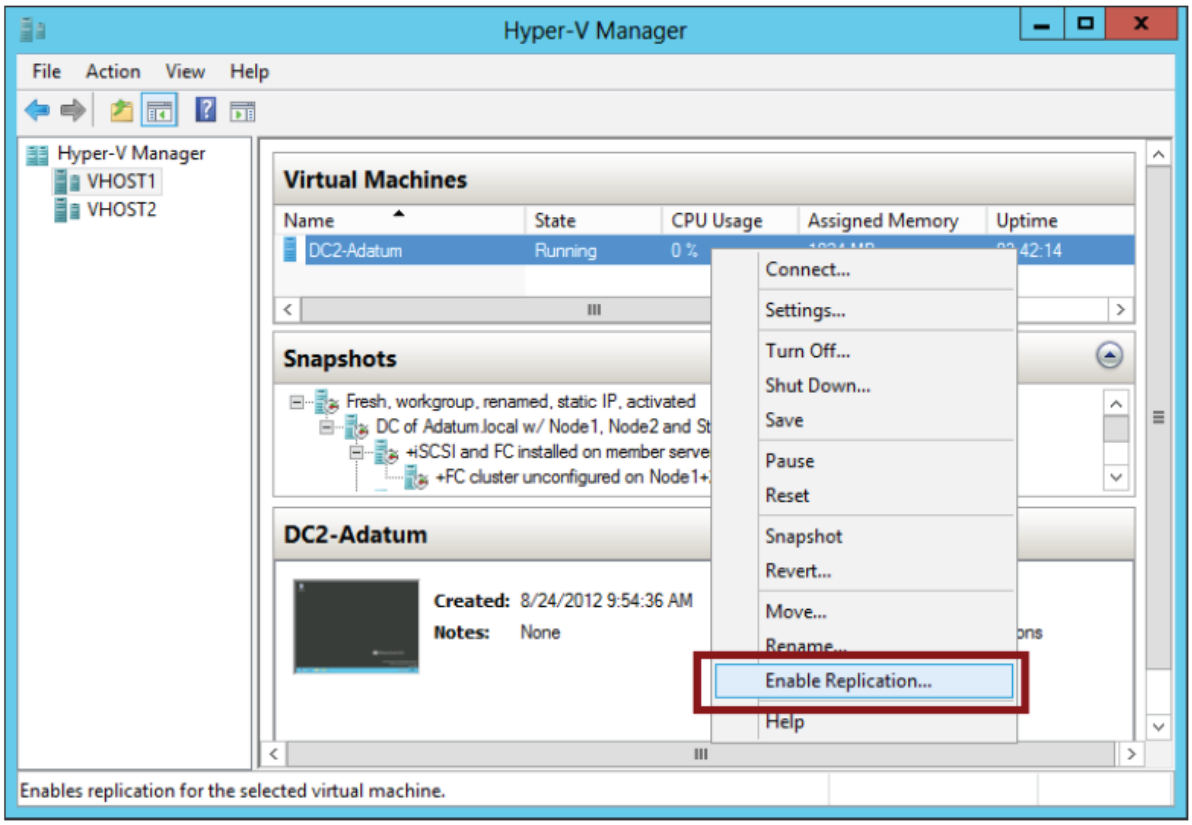

FIGURE 12-16 Creating a replica of a virtual machine.

This step opens the Enable Replication Wizard. The wizard includes the following five configuration pages:
1. **Specify Replica Server page** Use this page to specify the remote replica server by name.

2. **Specify Connection Parameters page** This page, shown in Figure 12-17, asks you to specify which of the authentication types enabled at the server level in Hyper-V Settings you want to use to support this replicated VM. If you have enabled only one of these two authentication methods at the server level, that same method is the only option here. Naturally, the replica server must support the same authentication method.

This page also provides an option that lends itself fairly well to an exam question: the Compress The Data That Is Transmitted Over The Network check box. This compression option reduces bandwidth requirements for replication at the expense of increased processor usage. If this option does appear on the exam, this trade-off is likely to be the key to getting the right answer.

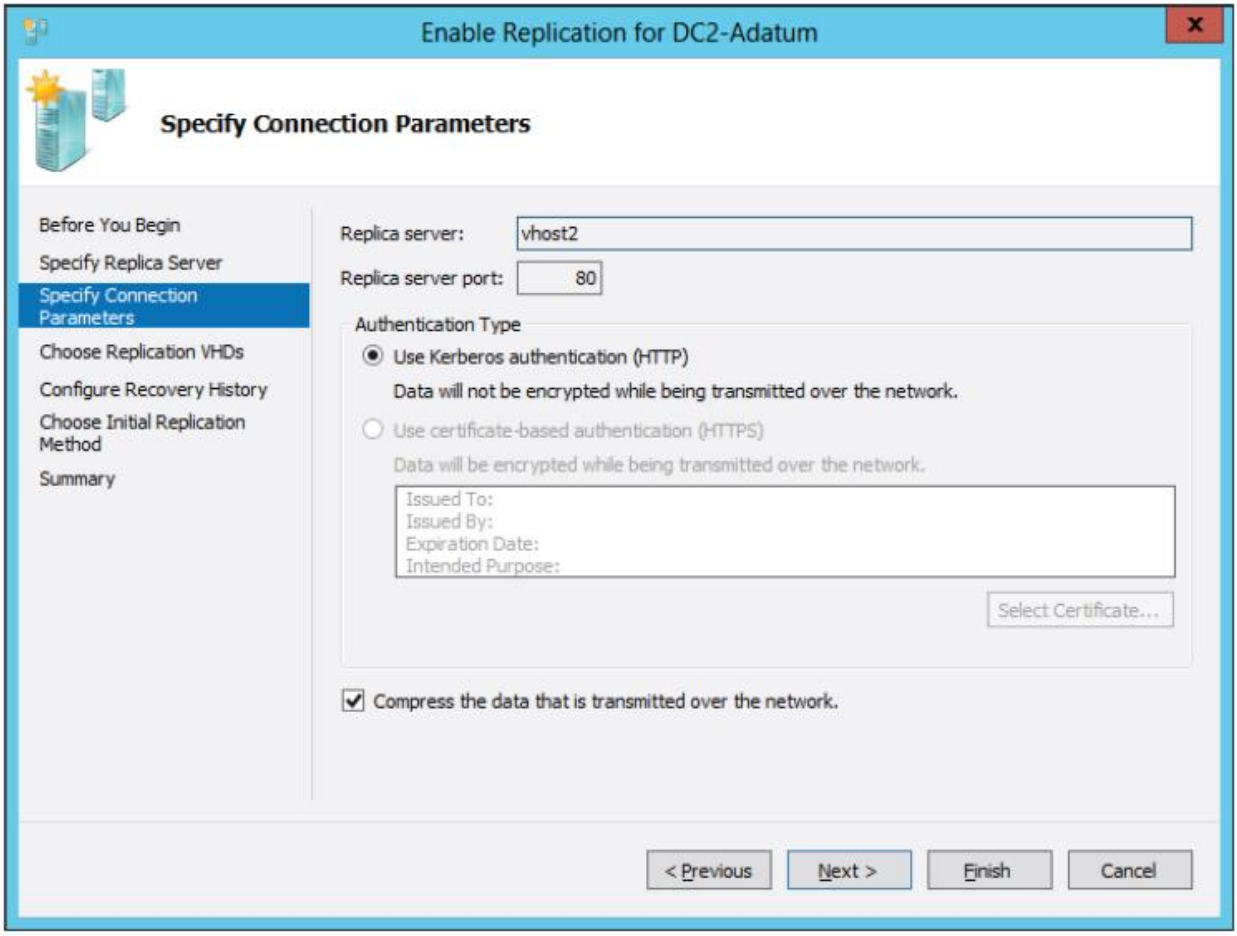

FIGURE 12-17 Selecting authentication and compression settings for a replicated VM.

EXAM TIP

If both authentication types are available for the VM and you want to change the authentication type later, you have to remove replication and complete the Enable Replication wizard again. Before you do, though, make sure that certificate-based authentication is also enabled in the Hyper-V Settings on the remote host server.

3. **Choose Replication VHDs page** By default, all virtual hard disks (VHDs) attached to the VM are enabled for replication. You can use this page to deselect any VMs that you don't want to be replicated.

4. **Configure Recovery History page** This page, shown in Figure 12-18, includes the settings to configure recovery points. These are among the most likely of all Hyper-V Replica settings to appear on the 70-417 exam. By default, the Only The Latest Recovery Point option is selected, and no other options are enabled or configured.

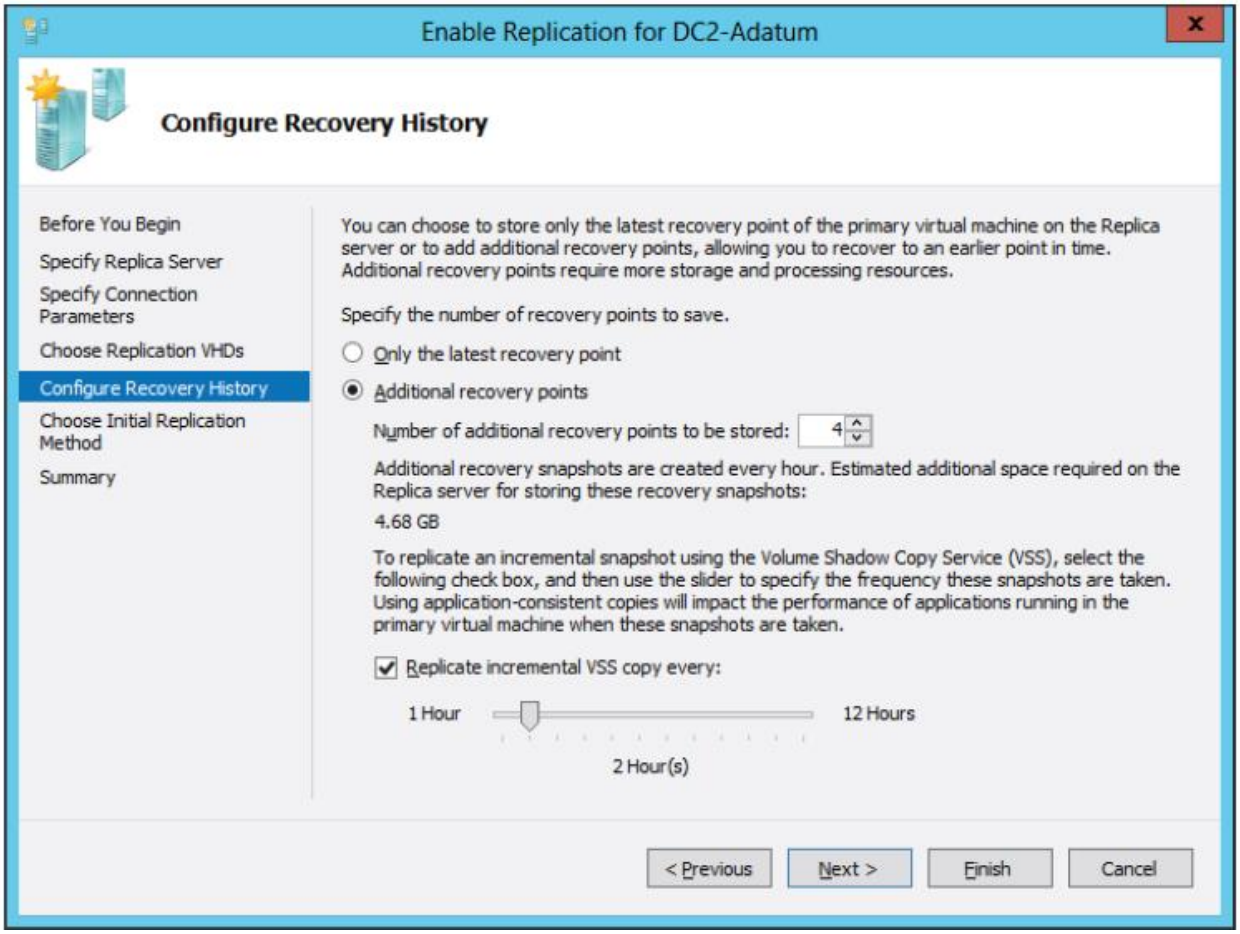

FIGURE 12-18 Configuring additional recovery points.

Recovery points are VM snapshots saved on a replica server. Replication traffic sends a new snapshot from the primary to the replica server every 5 to 15 minutes, but only the latest is saved on the replica by default. Selecting the Additional Recovery Points option configures the replica server to keep one extra snapshot per hour. If you later perform a failover operation at the replica server, you then have the option of recovering either the most recent version of the VM, which is always available, or one of these earlier, hourly snapshots.

A menu of available recovery points on a replica server is shown in Figure 12-19. If the Configure Recovery History page were left at the default setting (Only The Latest Recovery Point), only the first option named Latest Recovery Point would appear in this menu.

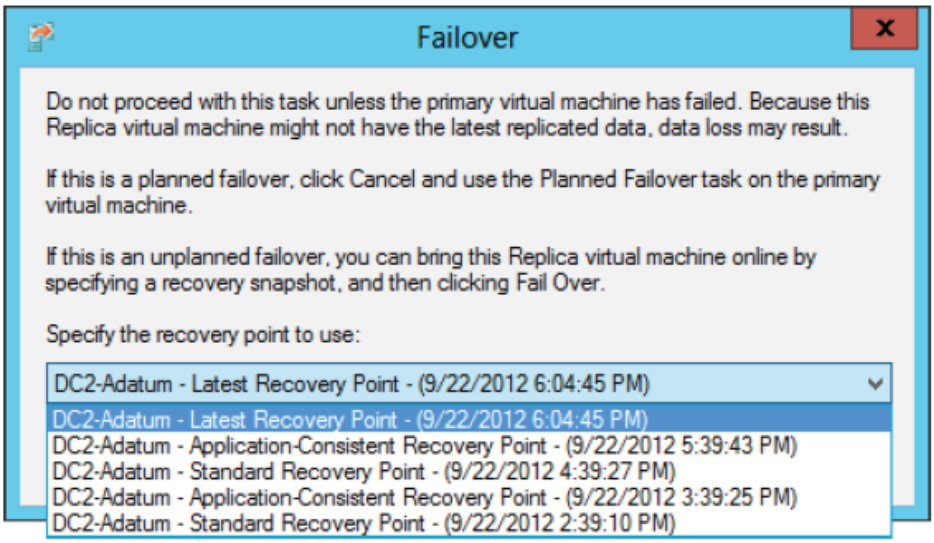

FIGURE 12-19 The latest recovery point and previous hourly snapshots of a VM that can be restored in a failover on the replica server.

When you enable the Additional Recovery Points option on the Configure Recovery History page, the replica server by default will keep an hourly snapshot for each of the past four hours in addition to the latest recovery point. However, you can change this setting if you want to store more (or fewer) of these recovery points on the replica server. The main drawback to keeping many recovery points is the use of storage resources required to do so.

The last configuration settings on the Configure Recovery History page relate to *incremental Volume Shadow Copy Service (VSS) copies*, also known as *application-consistent recovery points*. These are high-quality snapshots taken during moments in which the VM momentarily "quiesces" (gracefully pauses) activity in VSS-aware applications such as Microsoft Exchange and SQL Server. The advantage of these snapshot types is that they help ensure that the failover will be free of errors in these applications. The disadvantage is that they are more processor-intensive and cause important applications to pause briefly. (However, it should be noted that the pause is normally too brief for users to detect.)

You enable incremental VSS copies by selecting the Replicate Incremental VSS Copy Every check box, and then selecting the frequency of the application-consistent recovery point. (You can see these options in Figure 12-18.) If you leave the default frequency of 1 hour, then every recovery point will be an application-consistent recovery point. If you select a frequency of 2 hours, then the standard recovery point will be replaced by an application-consistent recovery point every 2 hours, and so on. Figure 12-20 shows the snapshots stored on a replica server for which incremental VSS copies are scheduled every two hours.

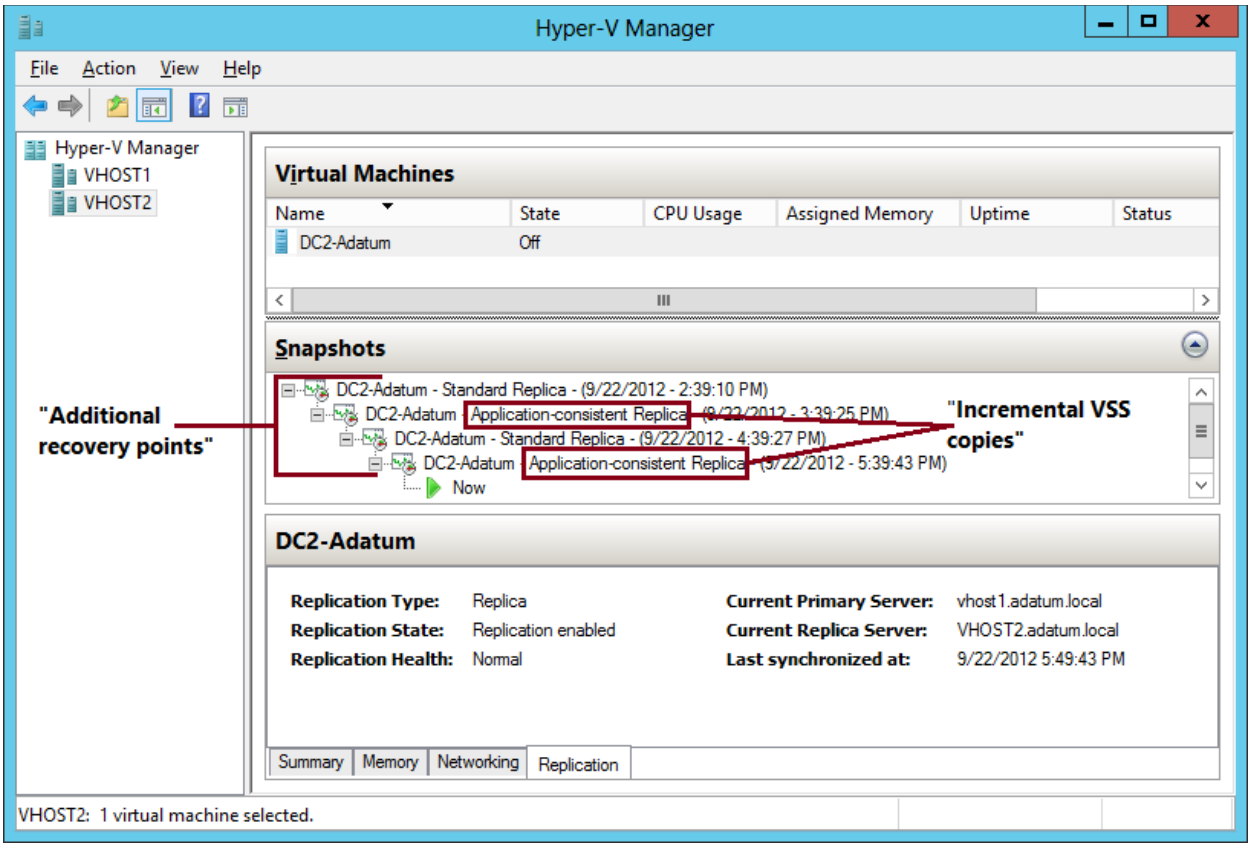

Expect to see a question about application-consistent snapshots on the 70-417 exam.

1. **Choose Initial Replication Method page** This page, shown in Figure 12-21, allows you to specify how the initial copy of the VHDs attached to the primary VM will be sent to the replica server. By default, the VHDs are sent over the network. Sending very large files over a network such as the Internet isn't always a realistic option, however. As an alternative, you can choose the second option, to export the VHDs to external media (and then physically transport them to the replica server). The final option is to use an existing VM on the replica server as the initial copy. You can choose this option if you have restored an exact copy of the VM and its VHDs on the replica server.
2.

This page also allows you to configure the initial network transfer to take place at a specified future time. You can use this option to minimize user disruption.

NOTE

Typically, the initial transfer of the VHD is far more bandwidth-intensive than the updates sent through replication are. After the initial copies of the VHDs are sent, only the changes (deltas) to these VHDs are sent during replication, which occurs every 5 to 15 minutes.

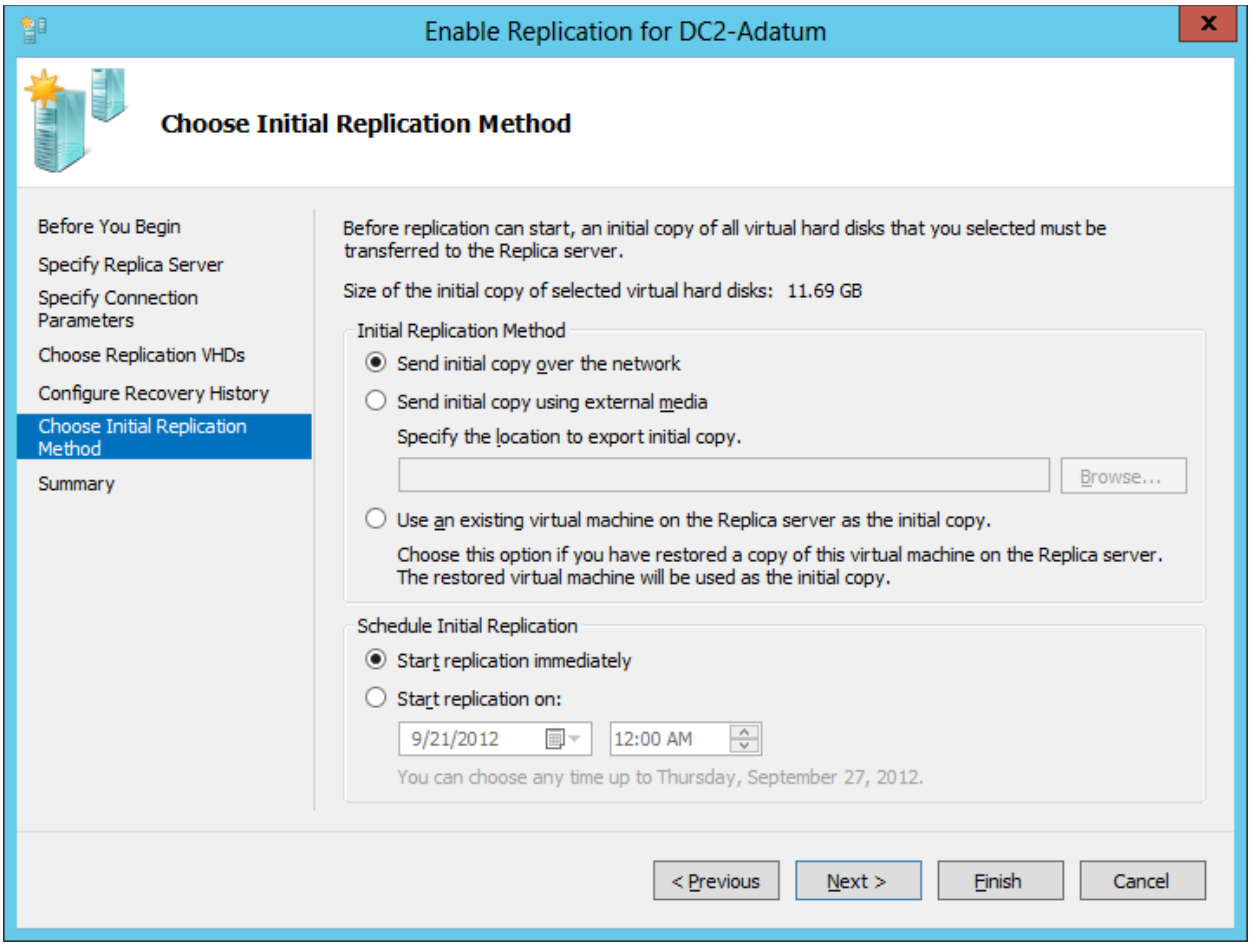

Failover TCP/IP settings

After you enable replication on a VM, you might need to specify the TCP/IP settings that will apply to the replica VM after failover. By default, the replica VM will inherit the same IPv4 and IPv6 configuration as the primary VM. In many cases, however, the replica VM will need a different IP configuration to communicate in its environment.

To assign a different IP configuration to the replica VM, in Hyper-V Manager on the replica server, right-click the replica VM and select Settings from the shortcut menu. In the Settings dialog box, expand Network Adapter in the left column and then select Failover TCP/IP, as shown in Figure 12-22. In the right pane, assign the new IP configuration as appropriate.

Then, on the primary server, assign the original IP configuration in the same settings area. Otherwise, the replica settings will persist if you fail back to the original location. (Remember this last point for the exam.)

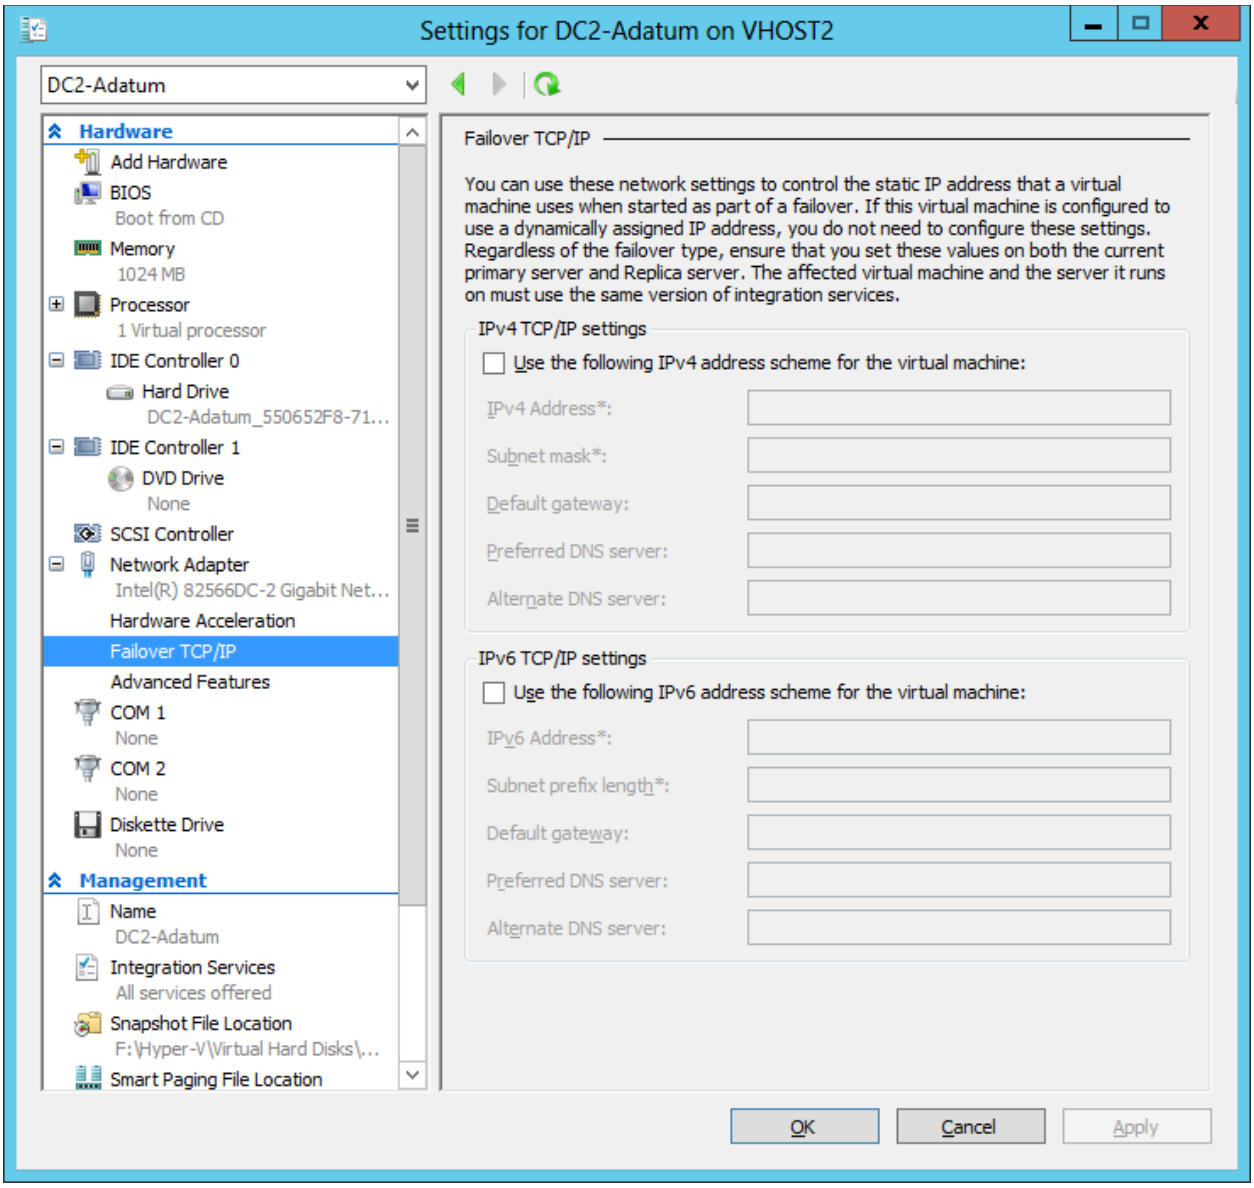

FIGURE 12-22 Assigning a different IP configuration to a replica VM.

Resynchronizing the primary and replica VMs

After you complete the Enable Replication wizard, you can modify the replication settings for a VM in the Settings dialog box for that VM. Replication settings appear in the Management category in the menu on the left, as shown in Figure 12-23.

One configuration setting appears here that does not appear in the Enable Replication wizard: resynchronization. Resynchronization is a highly resource-intensive operation that is performed occasionally between a primary and replica VM. By default, resynchronization can occur at any time. You have the option, however, to restrict resynchronizations to selected off-peak hours. Alternatively, you can opt to perform resynchronization manually.

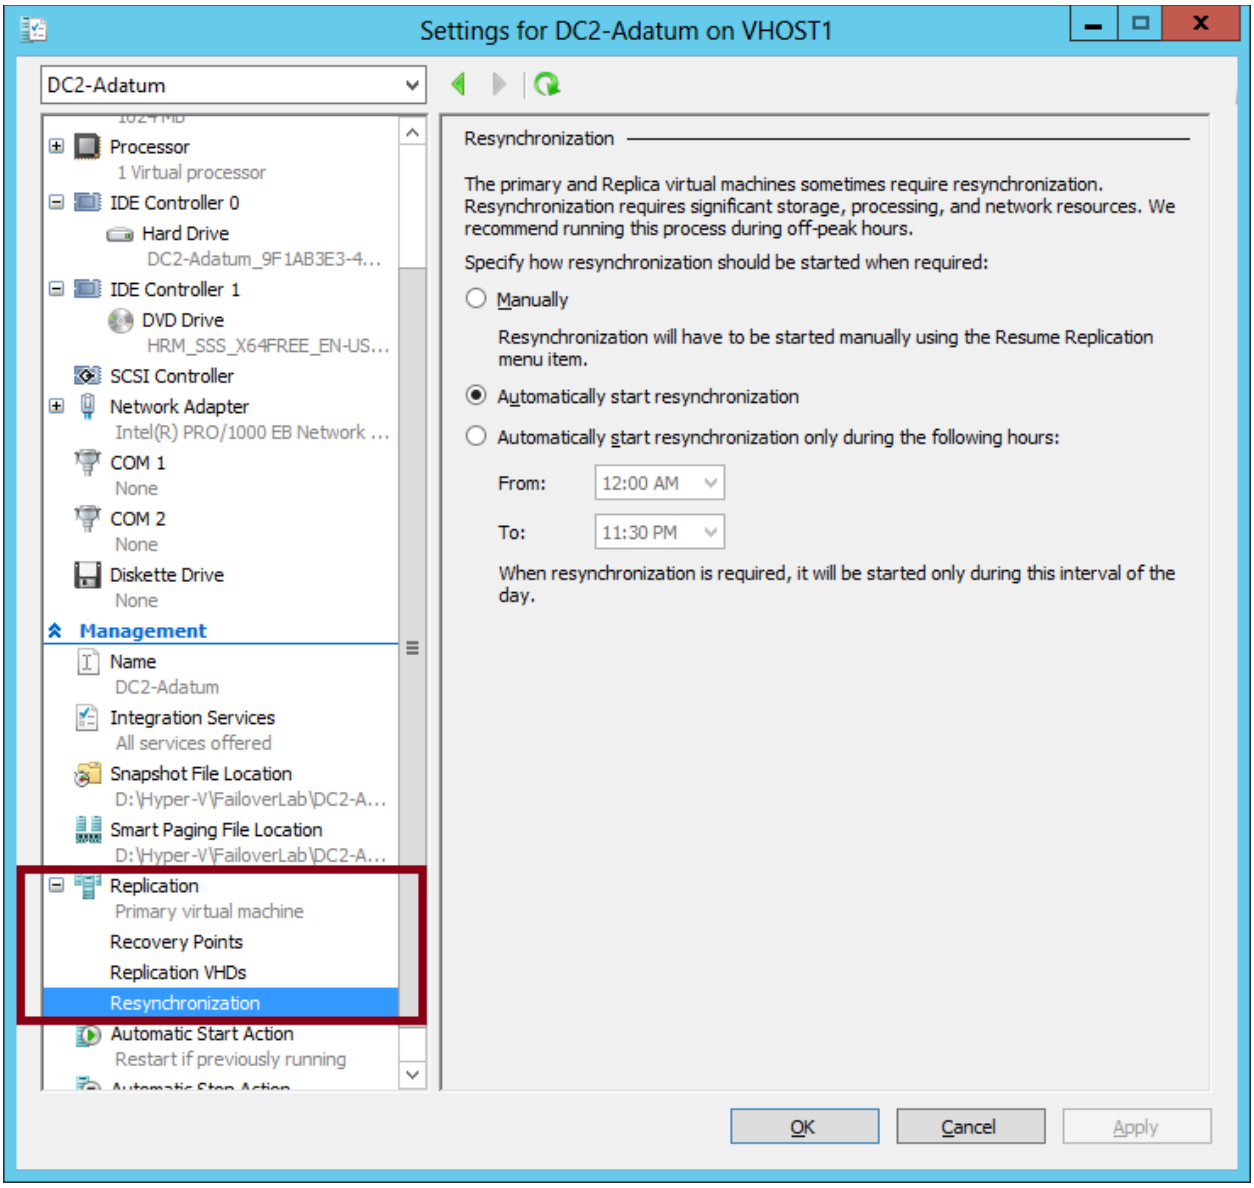

FIGURE 12-23 Replication settings for a VM.

Performing Hyper-V Replica failover

You can perform three types of failovers with Hyper-V Replica after it is configured: planned failovers, unplanned failovers, and test failovers. It's likely you'll see an exam question in which you need to understand the difference among them and when they are used.

 Planned failover A planned failover is the only failover you initiate from the primary server. You use this method whenever you can manually shut down the primary VM, and the primary and replica servers can still communicate.

A planned failover is the preferred failover type because no data is lost. In fact, you cannot even use this option to fail over to the latest recovery point or to any earlier recovery point. With a planned failover, only an exact copy of the current primary VM and its VHDs can be failed over to the replica server.

A planned failover is a good option in the following situations:

 You want to perform host maintenance on the primary server and temporarily want to run the VM from the replica.

• Your primary site is anticipating a possible power outage, and you want to move the VM to the replica site.

 You are expecting a weather emergency such as a flood, and you want to ensure business continuity.

 Your compliance requirements mandate that you regularly run your workloads for certain periods of time from the replica site.

To perform a planned failover, you begin by *shutting down the primary VM*. You then right-click the VM in Hyper-V Manager, click Replication, and then click Planned Failover, as shown in Figure 12-24. The latest updates are then sent to the replica server, the VM is failed over, and the replica VM is automatically started on the remote server. At the end of this operation, the replication relationship is reversed, so what was the replica server becomes the primary server, and vice versa.

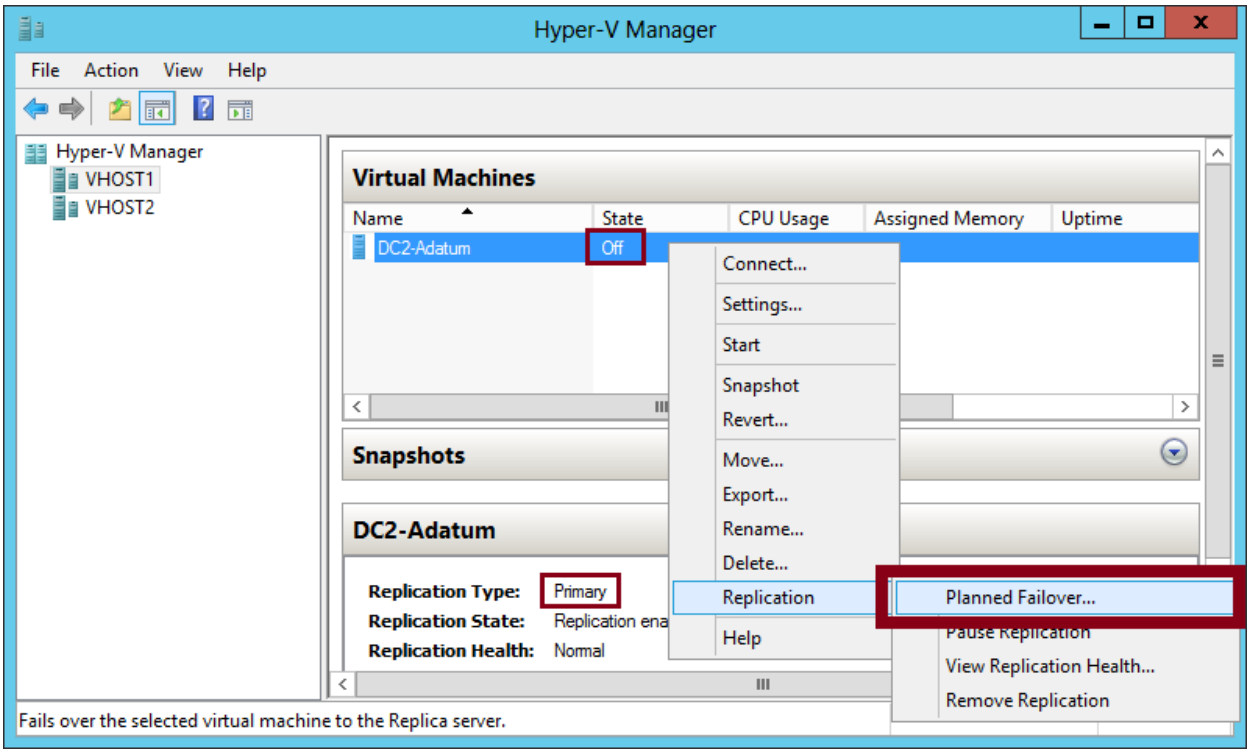

• **(Unplanned) failover** This type of failover is called an unplanned failover in the Windows Server 2012 documentation, but in the actual interface, it's called just "failover." On the 70-417 exam, you might see it referred to either way.

An unplanned failover is performed at the replica server. You perform this failover type when the primary VM fails suddenly and cannot be brought back online.

An unplanned failover is a good option in the following situations:

 Your primary site experiences an unexpected power outage or a natural disaster.

 Your primary site or VM has had a virus attack, and you want to restore your business quickly with minimal data loss by restoring your replica VM to the most recent recovery point before the attack.

To perform an unplanned failover, in Hyper-V Manager on the replica server, right-click the replica VM, click Replication, and then click Failover, as shown in Figure 12-25.

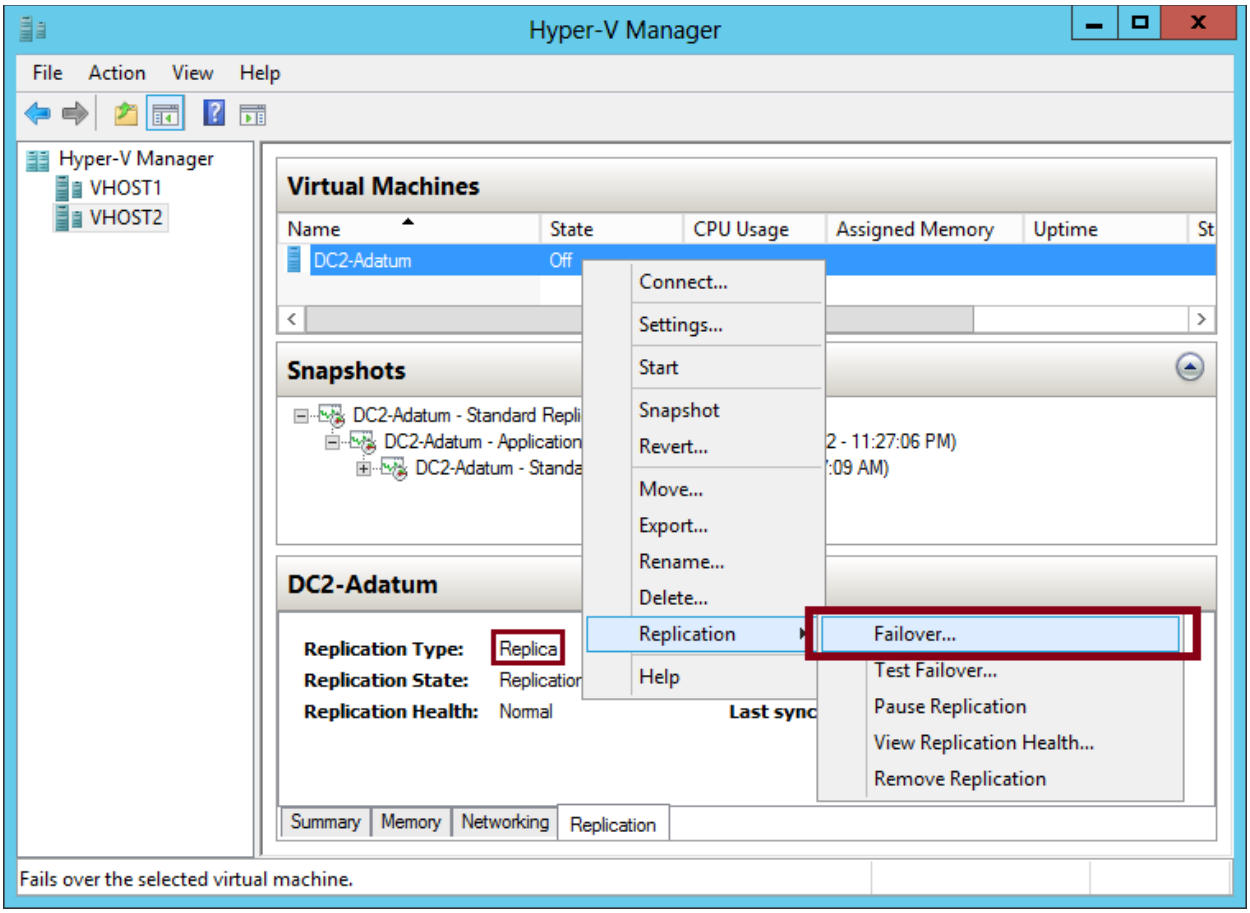

When you perform an unplanned failover, you have to choose a recovery point, as shown earlier in Figure 12-19. The VM is then started on the replica server.

After the replica VM is started, the replica relationship with the primary VM is broken, and replication stops. If at some later point you can bring the original primary VM online, you can resume replication by reversing the replication relationship. After you perform this operation, the local replica server becomes the new primary, and the remote primary becomes the new replica. To reverse replication in this way, right-click the VM on the replica server, click Replication, and then click Reverse Replication, as shown in Figure 12-26. This step starts the Reverse Replication Wizard, which allows you to reenter the settings for the replica.

Another option you can see on the Replication submenu in Figure 12-25 is Cancel Failover. You can safely choose this option after you perform an unplanned failover as long as no changes have been made to the replica. After you cancel a failover, you have to manually resume replication on the primary VM by right-clicking it and selecting Resume Replication. Cancelling a failover is a good idea if you quickly discover after performing an unplanned failover that the primary VM can be brought online.

Remember the Reverse Replication and Cancel Replication options for the exam.

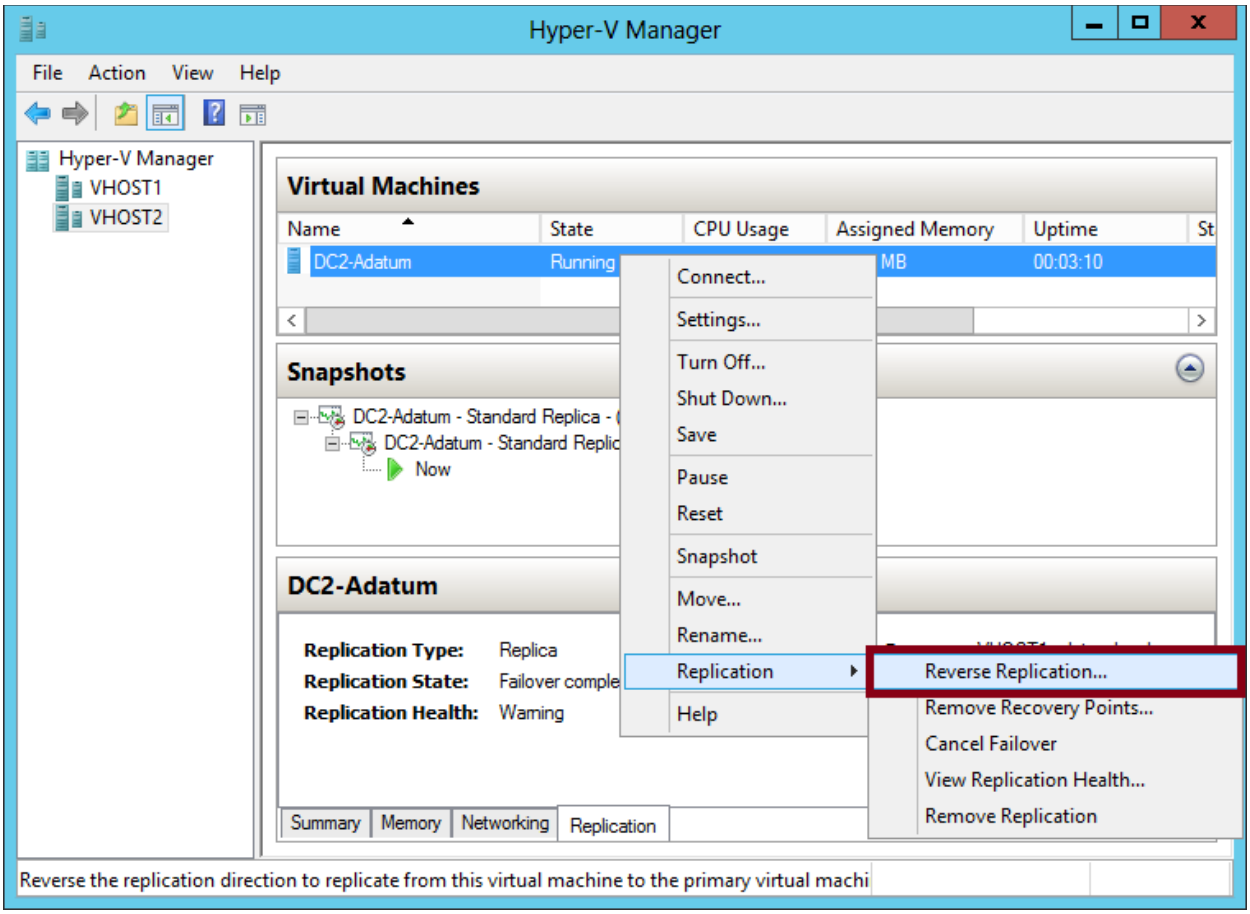

FIGURE 12-26 Reversing replication.

Test failover A test failover is the only failover operation you can perform while the primary VM is still running. The purpose of this failover type is to simulate an unplanned failover so that you can ensure that it will function as planned in case of an emergency.

To perform a test failover, in Hyper-V Manager on the replica server, right-click the replica VM, click Replication, and then click Test Failover. You then have to select a recovery point, just as you do with an unplanned failover. Next, a local, disposable copy of the replica VM is created on the replica server. The new copy of the VM appears in Hyper-V Manager in a stopped state with the tag "- Test." For example, a test failover of a VM named "MyVM1" would result in a new VM called "MyVM1 – Test". You can then start the new VM manually to see if it works as expected.

By default, the virtual network adapters of the test VM are disconnected from all virtual switches. If desired, you can preattach the adapter(s) of the test VM to a virtual switch of your choice. To do so, open the settings of the base replica VM, expand Network Adapter, and then click Test Failover, as shown in Figure 12-27. Make sure you choose a virtual switch that will not create any conflicts in a production network . After you examine the functioning of the test VM, you can safely delete it in Hyper-V Manager.

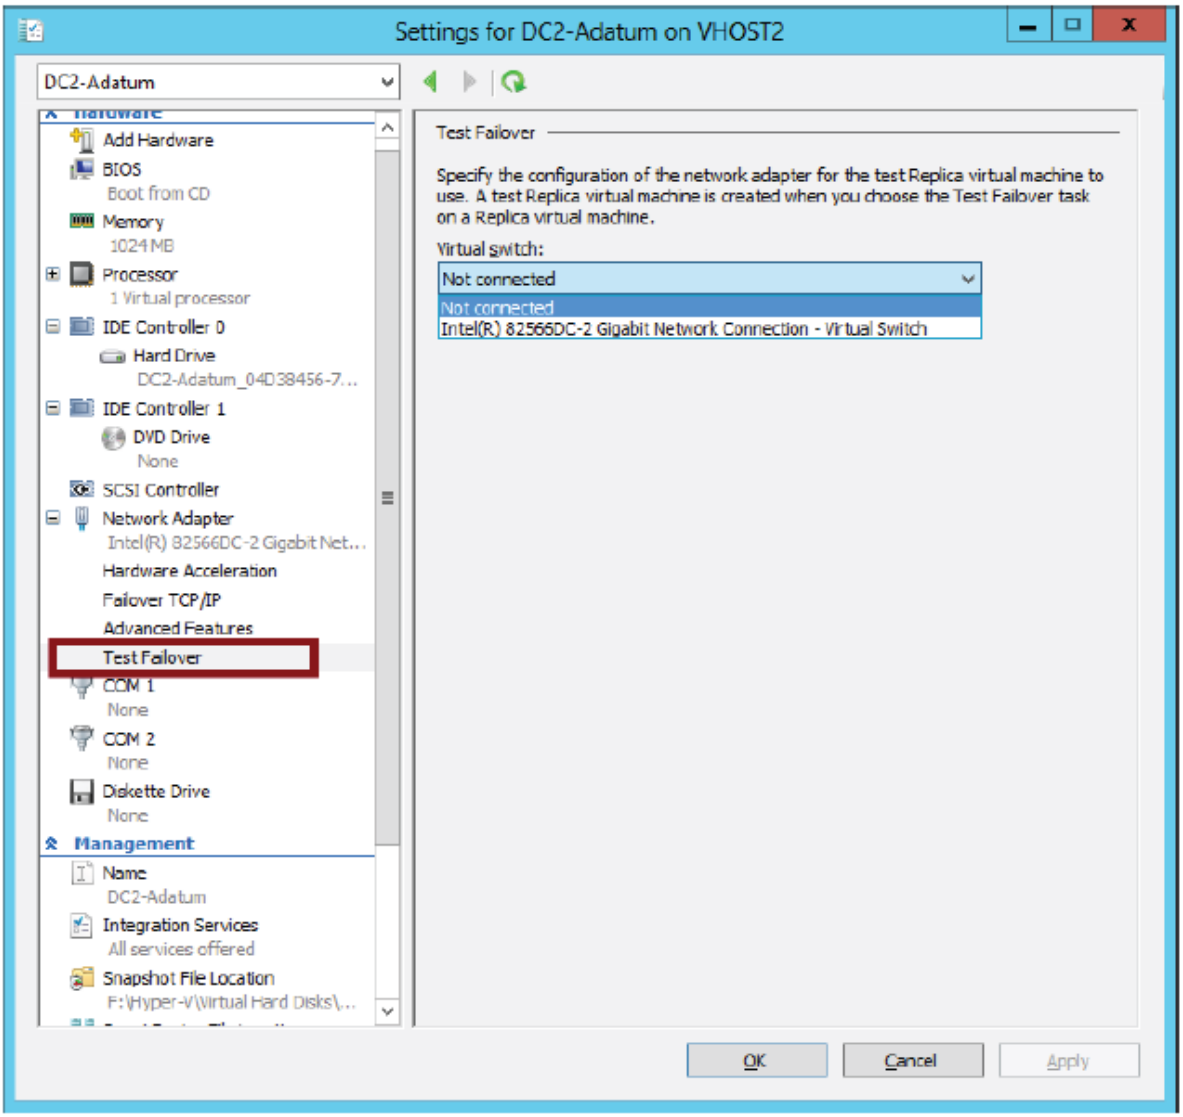

FIGURE 12-27 Preattaching the network adapter of a failover test VM to a virtual switch.

Using Hyper-V Replica in a failover cluster

The configuration steps previously described apply to VMs that are not hosted in a failover cluster. However, you might want to provide an offsite replica VM for a clustered VM. In this scenario, you would provide two levels of fault tolerance. The failover cluster is used to provide local fault tolerance, for example, if a physical node fails within a functioning data center. The offsite replica VM, on the other hand, could be used to recover only from site-level failures, for example, in case of a power outage, weather emergency, or natural disaster.

The steps to configure a replica VM for a clustered VM differ slightly from the normal configuration, but they aren't complicated. The first difference is that you begin by opening Failover Cluster Manager, not Hyper-V Manager. In Failover Cluster Manager, you then have to add a failover cluster role named *Hyper-V Replica Broker* to the cluster. (Remember, the word "role" is now used to describe a hosted service in a failover cluster.)

To add the Hyper-V Replica Broker role, right-click the Roles node in Failover Cluster Manager and select Configure Role. This step opens the High Availability Wizard. In the High Availability Wizard, select Hyper-V Replica Broker, as shown in Figure 12-28.

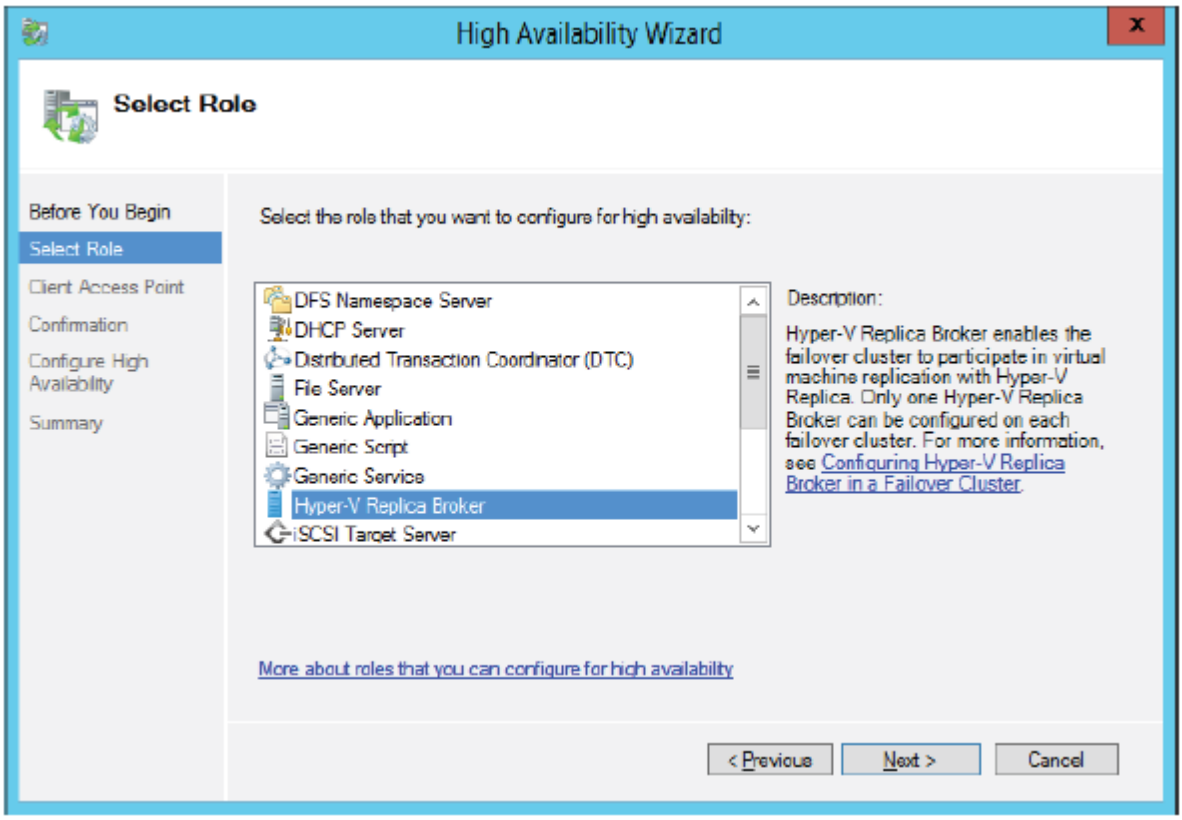

FIGURE 12-28 Adding the Hyper-V Replica Broker role to a failover cluster.

When you choose this role, the High Availability Wizard will then ask you to provide a NetBIOS name and IP address to be used as the connection point to the cluster (called a client access point, or CAP). This step is shown in Figure 12-29.

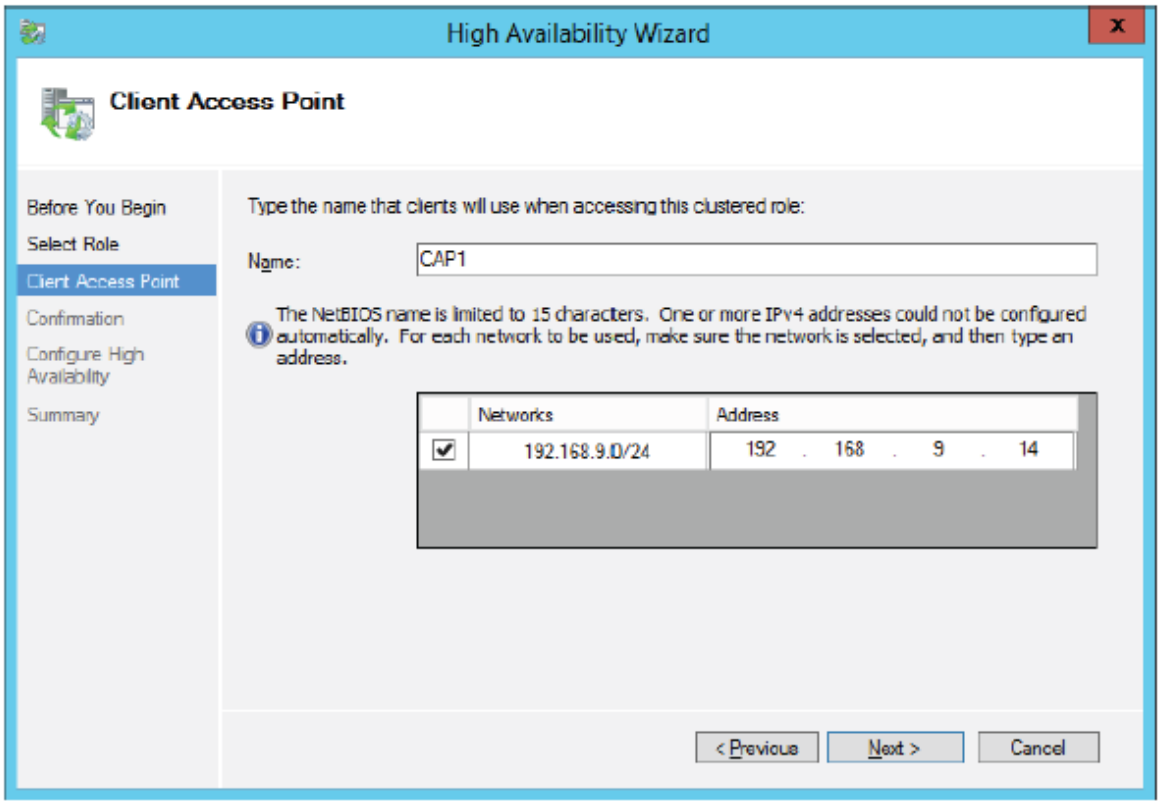

FIGURE 12-29 Providing a name and address for the client access point.

Next, you configure the equivalent of the server replication settings shown earlier in Figure 12-13. To do so, right-click the Hyper-V Replica Broker node in Failover Cluster Manager, and select Replication Settings from the shortcut menu, as shown in Figure 12-30. The difference between the settings here and the settings in Figure 12-13 is that in this case, the settings apply to the entire cluster as a whole.

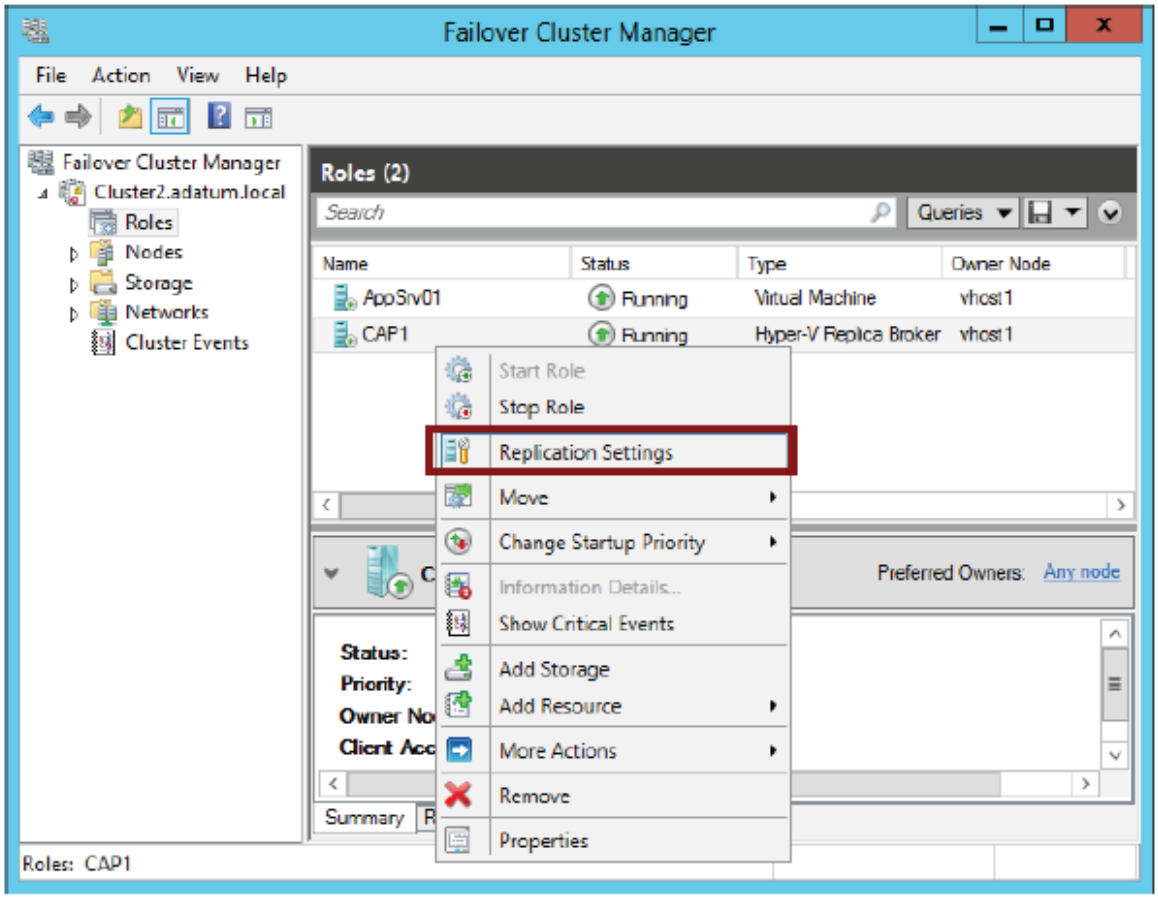

FIGURE 12-30 Configuring replication settings for the cluster.

On the remote Replica server, you configure replication as you normally would, by configuring Hyper-V Settings in Hyper-V Manager as described in the earlier section named “Configuring Hyper-V physical host servers.” However, if you want the remote Replica also to be a multi-node failover cluster, then you would need to configure that remote failover cluster through Failover Cluster Manager (by adding and configuring the Hyper-V Replica Broker role).

After you configure the host server settings, you can configure replication on the VM in Failover Cluster Manager just as you would in Hyper-V Manager. Right-click the clustered VM, click Replication, and then click Enable Replication, as shown in Figure 12-31.

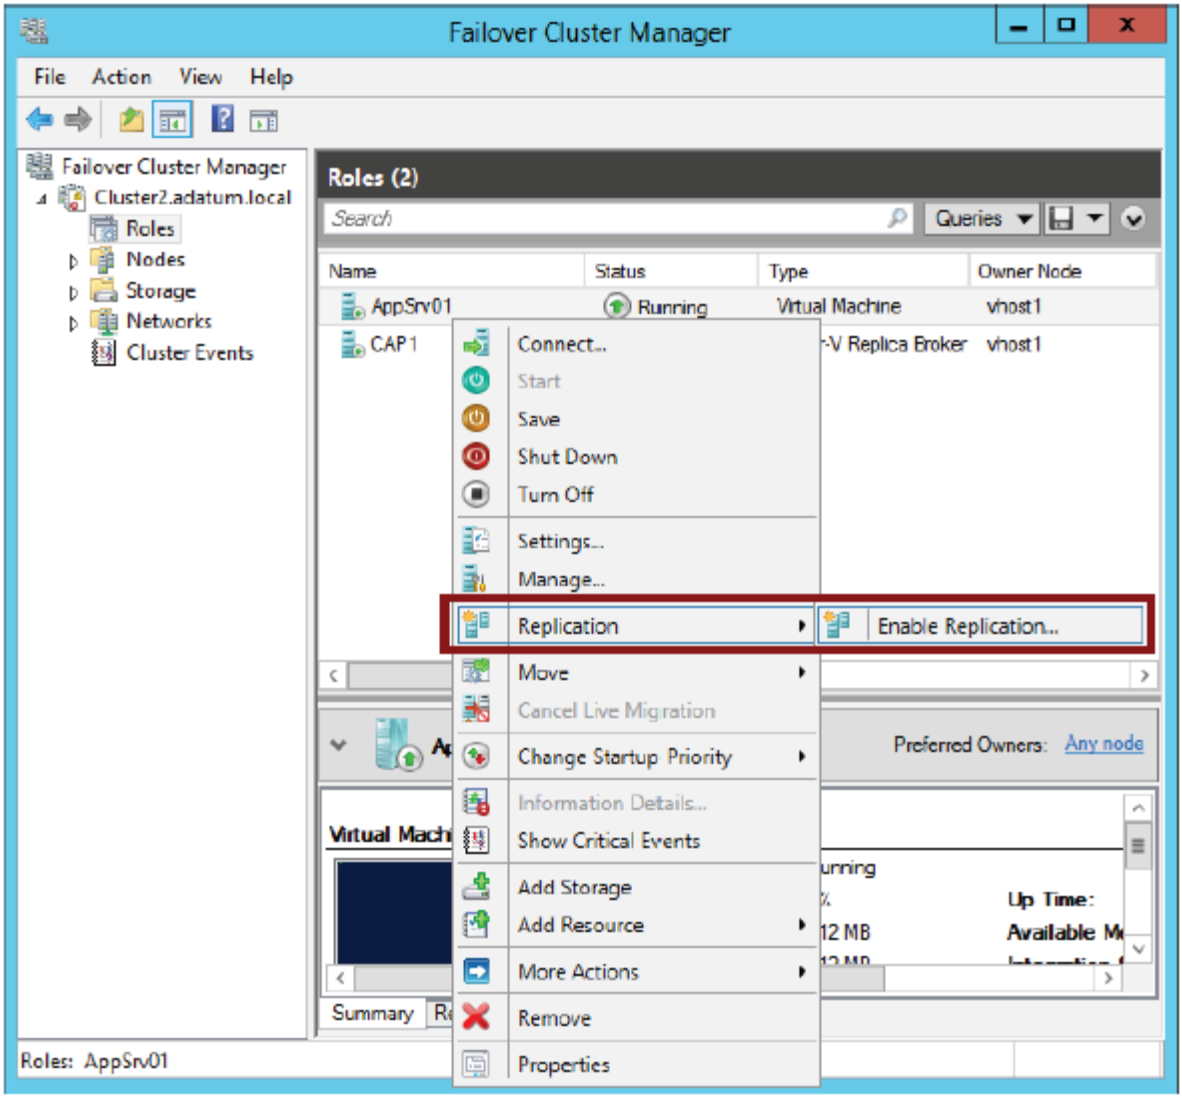

FIGURE 12-31 Enabling replication on a clustered VM.

This step opens the same Enable Replication wizard that you see when you configure replication on a nonclustered VM. The remaining configuration steps are therefore identical.

For the 70-417 exam, there's a good chance you'll be asked about basic concepts related to configuring replication on clustered VMs. Remember first of all that you use Failover Cluster Manager to configure replication for a clustered VM at the primary site but still use Hyper-V Manager at the Replica site. Remember that in Failover Cluster Manager at the primary site, you need to add the Hyper-V Replica Broker role to the failover cluster, and that this role is used to configure Hyper-V Replica "server" settings for the cluster. Finally, you also need to remember that when you configure Hyper-V Replica in a failover cluster, the CAP name and address are used as the server name and address.