

# Answers

## Chapter 1: Case Scenario Answers

---

### Preparing for Server Virtualization

1. You need to recommend a 64-bit server running x64 processors that include both hardware-assisted virtualization and DEP. You also have to make sure that the system's BIOS allows you to enable both of these features.
2. According to Table 1-8, you need 11 logical processors or processor cores. Note that this table does not list virtual processor requirements, but rather actual logical processors. If your system includes four cores per processor, you need three physical processors. Ideally, you will select a host system that includes more than three processor sockets so that it can scale the system up in time if you need to.
3. You should use DAS configured as two drives in a RAID 1 or mirrored configuration for the parent partition. You can use either DAS or remote storage to store the virtual machines or child partitions. However, to improve performance, you should store the virtual machines on a different set of disks than the parent partition. If you use DAS, you should create a separate RAID configuration using independent disks to store the VMs. This storage container needs to have at least 300 GB of disk space and probably more to support VM growth over time. If you use remote storage, you can rely on either NAS or SAN devices. These should also include a RAID configuration. The ideal RAID configuration for hosting virtual machines is RAID 10, which provides the fastest read and write rates.
4. Table 1-8 outlines a requirement for 13 GB of RAM. This means that the minimum configuration for this host server should be 16 GB of RAM. You might want to add more to support future server growth.
5. If you want to create a fault-tolerant configuration for this server, you must create a failover cluster. To do so, you need to make sure the hardware you select has been validated through the Failover Clustering Partner Program. In addition, you must ensure that your child partitions reside on remote storage so that each cluster node can access the VMs in the event of a hardware failure. Your second cluster node should include at least the same configuration as the original node to support failover.

## Chapter 2: Case Scenario Answers

---

### Networking Virtual Machines

1. The production machines should be linked to either a dedicated or external virtual network adapter. This will let them communicate with all of the users, other production servers, the external world, and the host server.
2. The test network should use either an internal or a private network connection. The private connection provides complete isolation, but the internal network lets you upload updates from the host server to the virtual machines and will prove more useful in this case.
3. Both Web server virtual machines should be linked to two virtual network adapters. For the production Web server, the external link will allow it to communicate with the public network and an internal link will let it communicate with the test Web server for Web site updates. The test Web server uses the same adapters but in reverse. The external adapter will let users communicate with the test Web server to modify content and the internal adapter will let it upload approved content to the production Web server.

## Chapter 3: Case Scenario Answers

---

### Deploying SCVMM on Physical or Virtual Platforms

1. Because your environment only consists of 100 VMs and you do not anticipate growth in machine numbers, you can use the recommended settings for an SCVMM deployment for 150 machines or fewer. In fact, you can deploy all of the SCVMM components on a single machine as long as it includes sufficient resources.
2. Deploying SCVMM on a physical computer offers several advantages:
  - When you deploy on a physical system you can ensure that the system is dedicated to the SCVMM role.
  - You can add additional resources such as CPUs or RAM whenever they are required; however, you must shut down the system to do so.
  - A major advantage of deploying SCVMM on a physical computer is that this computer does not have any dependencies on host computer systems. When an application is deployed in a VM, you must start the Hyper-V host before you can start the VM. Because SCVMM is a host and VM management tool, you would not be able to use it until its host server is started and the VM running SCVMM was started.
  - Running an application on a physical computer lets it own all resources. It does not have to compete with other VMs for host resources.

On the other hand, running SCVMM on a physical computer will have drawbacks:

- To deploy SCVMM on a physical computer, you need to launch an acquisition process to obtain the new system. This process can often take weeks to complete. Meanwhile, you would not be able to run SCVMM.
- If additional resources are required, you must once again launch the acquisition process.
- Physical computers require power, space, and cooling resources from the datacenter. These add to the costs of the computer in the long run.

**3.** If you deploy SCVMM on a virtual machine, you can benefit from several advantages:

- Virtual machines have very little overhead cost. If your host servers are properly sized, you can deploy the new VM right away.
- Because your host servers are fault tolerant through the Failover Clustering service, you will be able to configure an SCVMM VM as highly available without additional costs. Doing so on a physical system will add significantly to the cost of the system.
- You can begin the deployment in a test environment, and if everything goes well, you will be able to rapidly move the test VM into production—or even better, rapidly reproduce the VM deployment in production.
- It is always easy to add resources to a VM as long as they are available on the host computer: Just shut it down, add new resources, and start it up again.
- The reason Contoso moved to virtual machines in the first place is the same as most other organizations: Their physical resources were seriously underutilized. Deploying SCVMM on a physical computer does not follow the new datacenter direction Contoso has set by moving to a virtual infrastructure.

On the other hand, running SCVMM on a virtual computer will have drawbacks:

- SCVMM will have to compete with resources with other VMs on the same host computer. This may have an impact on performance.
- Most important, you will not have access to SCVMM in the event of a complete host server shutdown. This means that you will not be able to manage hosts or VMs through SCVMM until at least one host is started and the SCVMM VM is started. You may also have to start a domain controller if your DCs are running in VMs as well. However, this can be done by using both Hyper-V Manager and Failover Clustering Management from a workstation and then switching to SCVMM once it is running.

**4.** SCVMM has a high disk overhead because of the Library and other functions it offers.

This means that the most significant resource for SCVMM in a virtual machine will be disk operations. You should consider creating a set of pass-through disks for the SCVMM data partitions and tying them to the virtual SCSI controller on the VM. You must use one IDE disk to boot the VM—this can be in a traditional VHD—but if you create data partitions on pass-through disks, you gain the best of both worlds. Your SCVMM deployment will gain the advantage of a VM and the performance of a physical computer at the same time.

# Chapter 4: Case Scenario Answers

- 1. You have several options for moving pass-through disks. One of the easiest is to obtain a disk replication tool, but unfortunately, because of the economic downturn, management refuses to let you acquire such a tool. Therefore, you must use another option. One of the best ways to do this is to convert the physical disks into virtual hard drives using the New Virtual Hard Drive Wizard.
- 2. To transfer the complete VM from London to New York, you must use the following process:
  - a. Begin by making sure you have a complete backup of the virtual machine and its pass-through disks.
  - b. Move on to convert the disks from pass-through disks to virtual hard drives. You will use dynamically expanding VHDs in order to limit the amount of data you need to copy over the network from London to New York. Dynamically expanding disks will only contain the actual data in the disk and will therefore be smaller than fixed size disks and will take less time to copy. Use the New Hard Disk Wizard to create the new dynamically expanding VHDs and make sure you select Copy The Contents Of The Specified Hard Drive when you get to the Configure Disk page of the wizard (see Figure A-1). Repeat the operation for each pass-through disk.
  - c. When the drives are converted to VHDs, move to edit the virtual machine's configuration through the Settings dialog box. Disconnect each of the three pass-through disks from the VM and connect the new VHDs to the VM. Make sure you connect the disks in the right order. Save your changes.
  - d. Boot the VM with the new VHDs to make sure it operates properly.

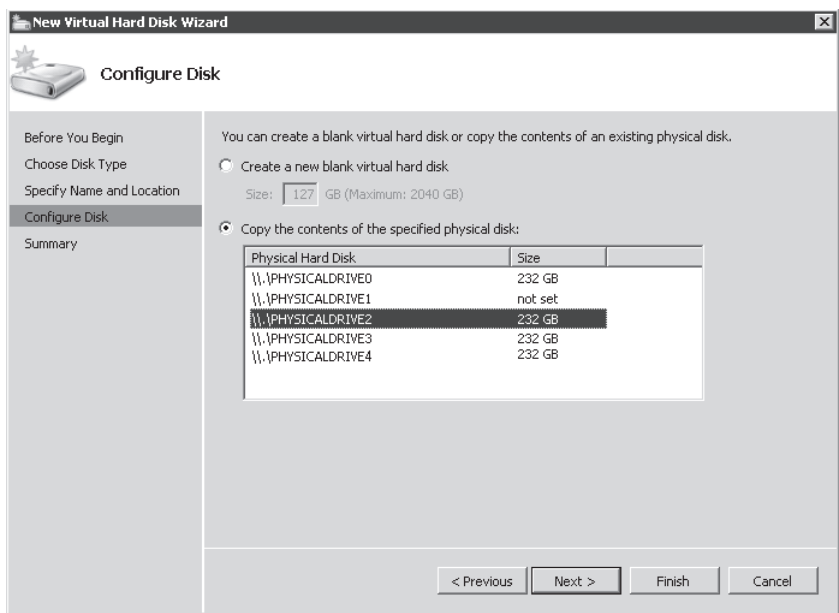


FIGURE A-1 Creating a VHD from a physical hard disk

- e. When the boot test is complete, export the VM from Hyper-V. Use the Export command that is available when you right-click the VM name in Hyper-V Manager. Store the exported VM in a new folder.
  - f. When the export is complete, copy the contents of the entire exported folder from London to New York. You might want to do this during off hours, such as on a weekend, because it will take considerable bandwidth.
  - g. In New York, use Hyper-V Manager to import the copied VM. Use the Import command in the Actions pane to perform the operation.
  - h. When the import operation is complete, open the VM's Settings dialog box to make sure that it is properly connected to network adapters and that other configuration parameters are set properly.
  - i. Boot the VM to ensure that it is working properly. If it is, your operation completed successfully.
3. Now that the pass-through disks have been converted to virtual hard drives, you will have lost the peak performance physical drives guarantee for VMs. However, you can achieve similar performance if you make sure that the new virtual hard disks meet two conditions:
- The virtual hard disks must be converted to fixed size disks. Fixed size disks provide the best VHD performance because Hyper-V does not need to consume processing overhead to manage the disk as the VM runs. You use the Edit Disk Wizard to convert the disks from dynamically expanding VHDs to fixed size VHDs. Make sure your storage container includes enough free space to host both the dynamically expanding disks and the fixed size disks during the conversion.
  - To provide the best performance for your fixed size VHDs, you should store them on three individual LUNs in your storage container and dedicate each LUN to each VHD. Ideally, your storage container will include enough spindles to dedicate separate spindles to each LUN. This will provide near-pass-through disk performance for your new VHDs.

## Chapter 5: Case Scenario Answers

---

1. To support operating system installation through WDS, your VMs must include a Legacy Network Adapter. The WDS bare-metal remote installation process relies on the PXE Boot process or booting from a network adapter, but to boot from a network adapter when no operating system is installed, you must use the Legacy Network Adapter.
2. Using the Legacy Network Adapter to boot from PXE and install the operating system through WDS means that each virtual machine will include both a Legacy Network Adapter and a Network Adapter. To make sure the VMs have a clean configuration and will run at peak performance, you have to remove the Legacy Network Adapter from their configuration each time a new VM is generated. This is needless work and can cause to errors and potential performance issues if the Legacy Network Adapter is not removed.

3. Because you are using SCVMM, you are better off using the Reference VM-Clone-Template process to generate new VMs. Begin with the creation of a reference VM, which is something that needs to be done with WDS anyway. Then clone the reference VM to create a source for the template. Generate a template from the clone and then use the template to generate all new VMs. This is much simpler and faster than using WDS and will lead to fewer potential errors than using a deployment method that is best kept for physical machines.
4. Finally, you should recommend to your manager that you do not need WDS in support of VM deployment. One reason is that although WDS can be used to deploy the operating system in a VM, it does not support deploying or generating the new VM itself. Therefore, you need to use SCVMM to perform this task anyway, making you use two tools to perform one task. In addition, if you use SCVMM on a SAN, you can take advantage of SAN Transfers and generate VMs at high speed, another feature that WDS does not offer but is an integral part of SCVMM.

## Chapter 6: Case Scenario Answers

---

1. Several machines can be migrated manually to improve the end result. This includes the following systems:
  - a. Server05 and Server 06 can easily be created as new machines because they are domain controllers. Install the DNS and AD DS services. This installation migrates the data from AD DS and DNS to the two new VMs. Then you migrate the Operations Master roles to the two new VMs. When this is done, you can remove Server05 and Server 06. Note that you must update DNS entries in DHCP and in any static IP address before you remove the servers.
  - b. The Exchange cluster can be migrated manually as well. Add two new virtual nodes to the cluster, move the Exchange workload onto the new VMs, and remove the physical nodes from the cluster.
  - c. The SQL Server cluster can also be migrated manually. Add two new virtual nodes to the cluster, move the SQL Server workload onto the new VMs, and remove the physical nodes from the cluster.
2. The only machine that must be moved while offline is the legacy Windows NT machine. Because SCVMM does not support the migration of NT, you must use another tool to perform the migration. You can download WinImage from [www.winimage.com](http://www.winimage.com) and even use an evaluation version to perform this migration. Use the Creating Virtual Hard Disk image from physical drive command in the Disk menu to generate the VHD for the NT machine. Then use the new VHD to create a VM on Hyper-V. Remember that this VM will not support Integration Services; therefore, you need to assign legacy network adapters to the VM as well as set the VM to reduce the functionality of the processors.

3. Five machines can be migrated online:
  - a. The two Web servers can be migrated online with SCVMM.
  - b. The two SharePoint servers can also be migrated online with SCVMM.
  - c. The DHCP server can be migrated online as well.
4. When filled in, the table would then look like table 6-8:

**TABLE 6-8** Conversion Results for Lucerne Publishing Server Roles

SERVER NAME	ROLE	CONVERSION TYPE
Server01	Web Server	Online
Server02	Web Server	Online
Server03	SharePoint Portal Server	Online
Server04	SharePoint Portal Server	Online
Server05	Active Directory Domain Services and Global Catalog	Manual
Server06	Active Directory Domain Services	Manual
Server07	SQL Server in Cluster	Manual
Server08	SQL Server in Cluster	Manual
Server09	Legacy App (Win NT)	Offline
Server10	Exchange Server in Cluster	Manual
Server11	Exchange Server in Cluster	Manual
Server12	DHCP	Online

## Chapter 7: Case Scenario Answers

1. You could simply let the Development group create their own VMs through the Self-Service Portal, but because you want to speed up the process, you will create them yourself. Because you already have a template VM in SCVMM, you can begin by generating one machine with the New Virtual Machine Wizard. When you get to the end of the operation, click View Script and save the script as **NewDevVMs.ps1**. Modify the script to include variables. These should cover the VM name, the computer name to assign to the VM and the VHD file name to use with the VM. Replace the appropriate entries in the script with your variables. You can then create a new VM by using the following cmdlet in the SCVMM PowerShell Prompt:

```
.\NewDevVMs.ps1 "MachineName" "ComputerName" "VHDFileName"
```

You then use the Up arrow to restore the cmdlet, change the values in each argument, and apply the script again. Repeat until the 20 VMs have been created.

2. The best way to monitor the new VMs is to use System Center Operations Manager and assign the new machines to be monitored through Performance and Resource Optimization. This will let you easily keep track of the VMs.
3. The last task is a bit trickier than the first. Begin by writing a script to create a checkpoint. Remember that SCVMM checkpoints are the same as Hyper-V snapshots. You place your files in a C:\Toolkit folder. You open Notepad and type the following text:

```
# -----
# Generate CheckPoints Script
# -----
Add-PSSnapin Microsoft.SystemCenter.VirtualMachineManager
Get-VMMServer -Computersname "SCVMM01.Contoso.com"
Get-VM -Name "VM01" | New-VMCheckpoint -RunAsynchronously
Get-VM -Name "VM02" | New-VMCheckpoint -RunAsynchronously
Get-VM -Name "VM03" | New-VMCheckpoint -RunAsynchronously
Get-VM -Name "VM04" | New-VMCheckpoint -RunAsynchronously
Get-VM -Name "VM05" | New-VMCheckpoint -RunAsynchronously
Get-VM -Name "VM06" | New-VMCheckpoint -RunAsynchronously
Get-VM -Name "VM07" | New-VMCheckpoint -RunAsynchronously
Get-VM -Name "VM08" | New-VMCheckpoint -RunAsynchronously
Get-VM -Name "VM09" | New-VMCheckpoint -RunAsynchronously
Get-VM -Name "VM10" | New-VMCheckpoint -RunAsynchronously
Get-VM -Name "VM11" | New-VMCheckpoint -RunAsynchronously
Get-VM -Name "VM12" | New-VMCheckpoint -RunAsynchronously
Get-VM -Name "VM13" | New-VMCheckpoint -RunAsynchronously
Get-VM -Name "VM14" | New-VMCheckpoint -RunAsynchronously
Get-VM -Name "VM15" | New-VMCheckpoint -RunAsynchronously
Get-VM -Name "VM16" | New-VMCheckpoint -RunAsynchronously
Get-VM -Name "VM17" | New-VMCheckpoint -RunAsynchronously
Get-VM -Name "VM18" | New-VMCheckpoint -RunAsynchronously
Get-VM -Name "VM19" | New-VMCheckpoint -RunAsynchronously
Get-VM -Name "VM20" | New-VMCheckpoint -RunAsynchronously
# -----
```

The first line loads the SCVMM PowerShell snap-in into the generic Windows PowerShell prompt. The second puts the focus on the SCVMM01 server. The remainder generate checkpoints for each VM. You use the `-RunAsynchronously` parameter because you want to get the prompt back as soon as possible. This parameter starts the job but does not wait until it is complete to return the prompt. Save the file as **NewCheckpoint.ps1**.

Next, you create a command file named **NewCheckpoint.cmd** and add the following command to it:

```
PowerShell -Command "& 'C:\Toolkit\NewCheckpoint.ps1'"
```

Save the file. Now you use Server Manager on SCVMM01 to go to the Configuration, Task Scheduler node and click Create Basic Task. Name the task **New CheckPoints**, set it



to run daily, set the start time, choose Start A Program as the action, type in **C:\Toolkit\NewCheckPoint.cmd**. **Right-click the newly created task** and open the Properties dialog box when the task is created. Set Run Whether User Is Logged On Or Not, move to the Triggers tab and click New, set the task to Daily and add the second time you want to run the task. Save the settings. When you do so, you are prompted for the account to use to run the task. You make sure you assign an account with local logon privileges and add its password. You also run the task once to make sure it works properly. You can view the new checkpoints in the SCVMM Administrator Console to verify that the checkpoints have been created.

You're done. You can now go back to converting your physical machines.

## Chapter 8: Case Scenario Answers

---

1. Your resource pool includes several different components in addition to the host servers. Specifically, it includes:
  - **Host servers** Currently three host servers are in production and two are in test and development environments.
  - **Domain controllers** Two virtual machines have been generated to run as domain controllers for the resource pool. These DCs run a utility forest that has no link to the production environment.
  - **Administrator workstations** Two workstations are used to administer the resource pool. Both are virtual machines running Windows Vista.
  - **System Center Virtual Machine Management Server** One server is running SCVMM in a VM. This server is also the Lucerne resource pool database server.
  - **SCVMM Library Server** One additional VM is running the Library Server role using pass-through disks to store resources in the Lucerne SAN.
  - **System Center Essentials** Another VM is running System Center Essentials to support controlled configuration management, updates to both hosts and VMs, and system monitoring.
2. The resource pool components are running on two networks. All host servers have a dedicated physical network adapter for management purposes only and a second adapter has been dedicated to resource pool traffic from the VMs that are part of the utility domain. All resource pool traffic is segregated at all times.
3. All virtual machines are running from redirected locations in the SAN. Each of the storage folders that contains VMs is secured using proper NTFS access control lists and all other resources are contained within the SCVMM Library, which makes them secure by default. Overall, your environment is very secure and is accessed only by trusted individuals.

## Chapter 9: Case Scenario Answers

---

1. To repair the host server, you must replace the faulty disk controller and then restore the latest full backup of the server. You can use the Windows Server 2008 installation media to launch the Complete PC Restore tool as soon as the disk controller has been repaired. This will restore the operating system and the VMs as they were when the last backup was created. Remember that when you restore the host and VMs, all VMs are replaced on the host server because WSB backups and restores are application-based.
2. Although you were able to restore the operating systems and the VMs up to the Friday data, several changes that were made on the weekend inside the VMs have been lost. To restore these changes, you need to locate the replicated VHDs and replace them on the host server in the proper location for each VM. You choose to overwrite the existing backup VHDs because the VHDs from Friday do not include the information from Sunday. Then you can go into each VM and perform a restore of the entire backup to update the VM to its latest version.
3. The best way to avoid this type of issue in the future is to create a host server failover cluster and to join HostServer03 to this cluster. If HostServer03 has issues again, the VMs it contains will automatically be failed over to other host nodes. There will be a slight interruption in service, but at least the VMs will be safe and will not need to be restored. If you place the host operating system on one volume and the guest VMs on another, as you must in a failover cluster, you only need to restore the host operating system volume and not the VM volume. The VMs will automatically be added to the host as soon as it is repaired and rejoins the cluster.

## Chapter 10: Case Scenario Answers

---

You begin the preparation of your report. As you begin, you dig further into the matter. Your responses to the questions are as follows:

1. You have set aside pass-through logical units (LUNs) in your storage area network for Exchange. You will therefore use only pass-through disks attached to virtual SCSI interfaces. You use SCSI interfaces because they let you attach larger disks. You use pass-through disks because they will grant you the best performance levels for Exchange VMs.
2. You have examined the various clustering options in Exchange and have decided that you will create the Exchange VMs on clustered host servers, but you will not make the Exchange VMs highly available. By creating stand-alone VMs on cluster hosts, you make sure that the Quick Migration feature is not available for the Exchange VMs and therefore cannot be used by mistake. However, you want to make your Exchange VMs highly available. To do so, you configure two VMs with pass-through disks acting as direct-attached storage. You then

configure Exchange using the Cluster Continuous Replication model, which is a multi-site clustering model using DAS and the Exchange transaction log replication engine to support high availability. While your two VMs are still in the same site, this multi-site clustering solution makes it very easy to create a highly available e-mail server configuration. You also make sure both nodes are installed on different host servers to avoid a single point of failure.

3. If maintenance operations are required on the host server running the active node of the Exchange VM cluster, you must use the Failover Clustering Management console to fail over the Exchange e-mail service. Then you can perform maintenance activities on the host server. Fail the VM node back after the host has been updated and rebooted. Repeat the operation for the second node.

