Vendor: Microsoft

Exam Code: 070-646Big5

Exam Name: Pro:Windows Server 2008, Server Administrator

Version: Demo

QUESTION NO: 1

You need to recommend a Windows Server 2008 R2 server configuration that meets the following requirements:

- Supports the installation of Microsoft SQL Server 2008
- Provides redundancy for SQL services if a single server fails

What should you recommend?

A. Install a Server Core installation of Windows Server 2008 R2 Enterprise on two servers. Configure the servers in a failover cluster.
B. Install a full installation of Windows Server 2008 R2 Standard on two servers. Configure Network Load Balancing on the two servers.
C. Install a full installation of Windows Server 2008 R2 Enterprise on two servers. Configure Network Load Balancing on the two servers.
D. Install a full installation of Windows Server 2008 R2 Enterprise on two servers. Configure the servers in a failover cluster.

Answer: D
Explanation:
Fail Over Clustering, which is available on the Enterprise edition (not on standard) will provide fail over as required.

Windows Server 2008 Enterprise Edition
Windows Server 2008 Enterprise Edition is the version of the operating system targeted at large businesses. Plan to deploy this version of Windows 2008 on servers that will run applications such as SQL Server 2008 Enterprise Edition and Exchange Server 2007. These products require the extra processing power and RAM that Enterprise Edition supports. When planning deployments, consider Windows Server 2008 Enterprise Edition in situations that require the following technologies unavailable in Windows Server 2008 Standard Edition:

Failover Clustering I-ail over clustering is a technology that allows another server to continue to service client requests in the event that the original server fails. Clustering is covered in more detail in Chapter 11. "Clustering and High Availability." You deploy failover clustering on mission-critical servers to ensure that important resources are available even if a server hosting those resources fails.

QUESTION NO: 2

Your network consists of a single Active Directory domain. Your main office has an Internet connection.

Your company plans to open a branch office. The branch office will connect to the main office by using a WAN link. The WAN link will have limited bandwidth. The branch office will not have access to the Internet. The branch office will contain 30 Windows Server 2008 R2 servers.

You need to plan the deployment of the servers in the branch office.

The deployment must meet the following requirements:

- Installations must be automated.
- Computers must be automatically activated.
- Network traffic between the offices must be minimized.

What should you include in your plan?

A. In the branch office, implement Key Management Service (KMS), a DHCP server, and Windows Deployment Services (WDS).
B. Use Multiple Activation Key (MAK) Independent Activation on the servers. In the main office, implement a DHCP server and Windows Deployment Services (WDS).
C. In the main office, implement Windows Deployment Services (WDS). In the branch office, implement a DHCP server and implement the Key Management Service (KMS).
D. Use Multiple Activation Key (MAK) Independent Activation on the servers. In the main office, implement a DHCP server. In the branch office, implement Windows Deployment Services (WDS).

Answer: A

Explanation:

The key here is that bandwidth from the branch to the main office is limited and there is no direct link to MS.

WDS and Product Activation

Although product activation does not need to occur during the actual installation process, administrators considering using WDS to automate deployment should also consider using volume activation to automate activation. Volume activation provides a simple centralized method that systems administrators can use for the activation of large numbers of deployed servers. Volume activation allows for two types of keys and three methods of activation. The key types are the Multiple Activation Key (MAK) and the Key Management Services (KMS) key.

Multiple Activation Keys allow activation of a specific number of computers. Each successful activation depletes the activation pool. For example, a MAK key that has 100 activations allows for

the activation of 100 computers. The Multiple Activation Key can use the MAK Proxy Activation and the MAK Independent Activation activation methods. MAK Proxy Activation uses a centralized activation request on behalf of multiple products using a single connection to Microsoft’s activation servers. MAK Independent Activation requires that each computer activates individually against Microsoft's activation servers.
The Branch office has no internet connection, so MAK is not the solution.

KMS requires at least 25 computers connecting before activation can occur, and activation must be renewed by reconnecting to the KMS server every 180 days.
You can use KMS and MAK in conjunction with one another. The number of computers, how often they connect to the network, and whether there is Internet connectivity determines which solution you should deploy. You should deploy MAK if substantial numbers of computers do not connect to the network for more than 180 days. If there is no Internet connectivity and more than 25 computers, you should deploy KMS. If there is no Internet connectivity and less than 25 computers, you will need to use MAK and activate each system over the telephone.

QUESTION NO: 3

Your network contains a Webbased Application that runs on Windows Server 2003. You plan to migrate the Webbased Application to Windows Server 2008 R2. You need to recommend a server configuration to support the Webbased Application.

The server configuration must meet the following requirements:

- Ensure that the Application is available to all users if a single server fails
- Support the installation of .NET Applications
- Minimize software costs

What should you recommend?

A. Install the Server Core installation of Windows Server 2008 R2 Standard on two servers. Configure the servers in a Network Load Balancing cluster.
B. Install the full installation of Windows Server 2008 R2 Web on two servers. Configure the servers in a Network Load Balancing cluster.
C. Install the full installation of Windows Server 2008 R2 Enterprise on two servers. Configure the servers in a failover cluster.
D. Install the full installation of Windows Server 2008 R2 Datacenter on two servers. Configure the servers in a failover cluster.

Answer: B

Explanation:

Web Edition meets the requirements

Windows Web Server 2008 R2

Windows Web Server 2008 R2 is designed to function specifically as a Web application server.

Other roles, such as Windows Deployment Server and Active Directory Domain Services (AD DS), are not supported on Windows Web Server 2008 R2. You deploy this server role either on a screened subnet to support a website viewable to external hosts or as an intranet server. As appropriate given its stripped-down role, Windows Web Server 2008 R2 does not support the high-powered hardware configurations that other editions of Windows Server 2008 R2 do. Windows Web Server 2008 R2 has the following properties:

Supports a maximum of 32 GB of RAM and 4 sockets in symmetric multiprocessing (SMP) configuration

You should plan to deploy Windows Web Server 2008 R2 in the Server Core configuration, which minimizes its attack surface, something that is very important on a server that interacts with hosts external to your network environment. You should plan to deploy the full version of Windows Web Server 2008 R2 only if your organization's web applications rely on features that are not available in the Server Core version of Windows Web Server 2008 R2. Unlike the Server Core version of Windows Web Server 2008, Windows Web Server 2008 R2 supports a greater amount of Internet Information Services (IIS) functionality.

Configuring Windows Network Load Balancing

While DNS Round Robin is a simple way of distributing requests, Windows Server 2008 NLB is a much more robust form of providing high availability to applications. Using NLB, an administrator can configure multiple servers to operate as a single cluster and control the usage ot the cluster in near real-time.

Why Failover Cluster will not work.

Contrast DNS Round Robin and NLB with Failover Clustering, another availability technology in Windows Server 2008. Formerly known as server clustering, Failover Clustering creates a group of computers that all have access lo the same data store or disk resource or network share. The applicationsjunning on aJailoverCluster must be cluster-aware. Failover Clustering has had some changes since Windows Server 2003. Lesson 2 will cover these changes.

QUESTION NO: 4

Your company purchases 15 new 64bit servers as follows:

- Five of the servers have a single processor.
- Five of the servers have a single dual core processor.
- Five of the servers have two quad core processors.

You plan to deploy Windows Server 2008 R2 on the new servers by using Windows Deployment Services (WDS). You need to recommend a WDS install image strategy that meets the following requirements:

- Minimizes the number of install images
- Supports the deployment of Windows Server 2008 R2

What should you recommend?

A. one install image file that contains three install images
B. one install image file that contains a single install image
C. two install image files that each contain a single install image
D. three install image files that each contain a single install image

Answer: B
Explanation:
You only need one image per processor type
Windows Deployment Services Images
Windows Deployment Services uses two different types of images: install images and boot images. Install images are the operating system images that will be deployed to Windows Server 2008 or Windows Vista client computers. A default installation image is located in the \Sources directory of the Windows Vista and Windows Server 2008 installation DVDs. If you are using WDS to deploy Windows Server 2008 to computers with different processor architectures, you will need to add separate installation images for each architecture to the WDS server. Architecture-specific images can be found on the architecture-specific installation media. For example, the Itanium image is located on the Itanium installation media and the x64 default installation image is located on the x64 installation media. Although you can create custom images, you only need to have one image per processor architecture. For example, deploying Windows Server 2008 Enterprise Edition x64 to a computer with 1 x64 processor and to a computer with 8 x64 processors in SMP configuration only requires access to the default x64 installation image. Practice exercise 2 at the end of this lesson covers the specifics ol adding a default installation image to a WDS server.

QUESTION NO: 5

Your network contains a single Active Directory site.

You plan to deploy 1,000 new computers that will run Windows 7 Enterprise. The new computers have Preboot Execution Environment (PXE) network adapters.

You need to plan the deployment of the new computers to meet the following requirements:

·Support 50 simultaneous installations of Windows 7

·Minimize the impact of network operations during the deployment of the new computers

·Minimize the amount of time required to install Windows 7 on the new computers

What should you include in your plan?

A. Deploy the Windows Deployment Services (WDS) server role. Configure the IP Helper tables on all routers.
B. Deploy the Windows Deployment Services (WDS) server role. Configure each WDS server by using native mode.
C. Deploy the Windows Deployment Services (WDS) server role and the Transport Server feature. Configure the Transport Server to use a custom network profile.
D. Deploy the Windows Deployment Services (WDS) server role and the Transport Server feature. Configure the Transport Server to use a static multicast address range.

Answer: D
Explanation:
http://technet.microsoft.com/en-us/library/cc726564%28WS.10%29.aspx
http://technet.microsoft.com/en-us/library/cc725964%28WS.10%29.aspx

WDS Multicast Server
Updated: November 21, 2007
Applies To: Windows Server 2008
The multicast server deploys an image to a large number of client computers concurrently without overburdening the network. When you create a multicast transmission for an image, the data is sent over the network only once, which can drastically reduce the network bandwidth that is used.

Using Transport Server
Updated: May 8, 2008
Applies To: Windows Server 2008
This topic only applies to Windows Server 2008. If you have Windows Server 2008 R2, see Configuring Transport Server.

You have two options when installing the Windows Deployment Services role in Windows Server 2008. You can install both the Deployment Server and Transport Server role services (which is the default) or you can install only the Transport Server role service. The second configuration is for advanced scenarios, such as environments without Active Directory Domain Services (AD DS), Domain Name System (DNS), or Dynamic Host Configuration Protocol (DHCP). You can configure Transport Server to enable you to boot from the network using Pre-Boot Execution Environment (PXE) and Trivial File Transfer Protocol (TFTP), a multicast server, or both. Note that Transport Server does not contain or support the Windows Deployment Services image store.

Configure how to obtain IP addresses. If multiple servers are using multicast functionality on a network (Transport Server, Deployment Server, or another solution), it is important that each server is configured so that the multicast IP addresses do not collide. Otherwise, you may encounter excessive traffic when you enable multicasting. Note that each Windows Deployment Services server will have the same default range. To work around this issue, specify static ranges that do not overlap to ensure that each server is using a unique IP address, or configure each of the servers to obtain multicast addresses from a Multicast Address Dynamic Client Allocation Protocol (MADCAP) server.

The server architectures are illustrated in the following diagram. The blue parts are installed with Transport Server and the Deployment Server. The grey parts are installed with the Deployment Server only. The yellow parts are not installed with either, but can be written using guidelines in the Windows SDK.

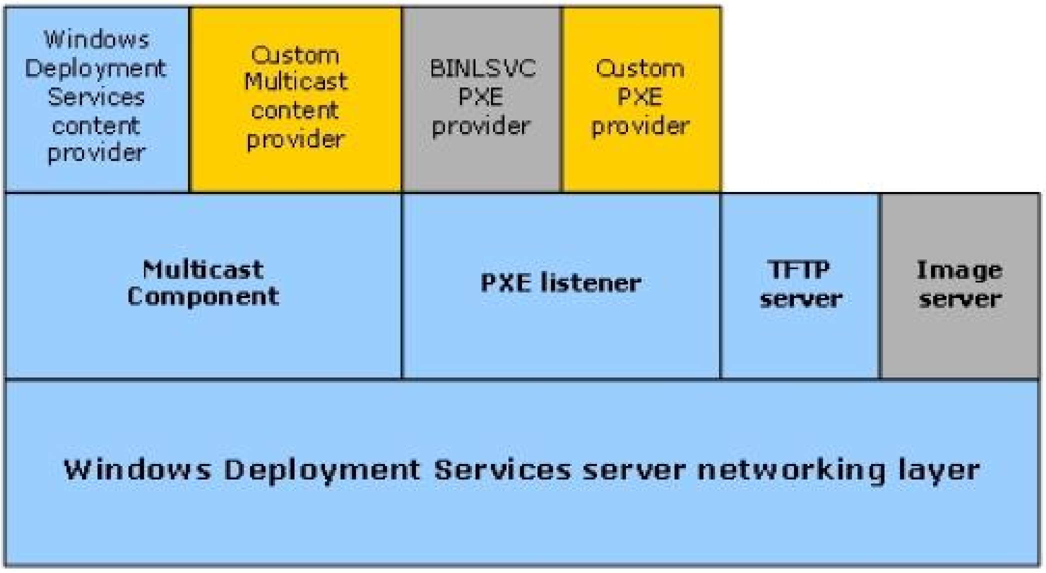
Windows Deployment Services content provider
Custom Multicast content provider
BINLSVC PXE provider
Custom PXE provider
Multicast Component
PXE listener
TFTP server
Image server
Windows Deployment Services server networking layer

QUESTION NO: 6

Your network consists of a single Active Directory site that includes two network segments. The network segments connect by using a router that is RFC 1542 compliant.

You plan to use Windows Deployment Services (WDS) to deploy Windows Server 2008 R2 servers. All new servers support PreBoot Execution Environment (PXE).

You need to design a deployment strategy to meet the following requirements:

- Support Windows Server?2008 R2
- Deploy the servers by using WDS in both network segments
- Minimize the number of servers used to support WDS

What should you include in your design?

A. Deploy one server. Install WDS and DHCP on the server. Configure the IP Helper tables on the router between the network segments.
B. Deploy two servers. Install WDS and DHCP on both servers. Place one server on each of the network segments. Configure both servers to support DHCP option 60.
C. Deploy two servers. Install WDS and DHCP on both servers. Place one server on each of the network segments. Configure both servers to support DHCP option 252.
D. Deploy two servers. Install WDS and DHCP on one server. Install DHCP on the other server. Place one server on each of the network segments. Configure both servers to support DHCP option 60.

Answer: A

Explanation:

http://support.microsoft.com/kb/926172

IP Helper table updates

The PXE network boot method uses DHCP packets for communication. The DHCP packets serve a dual purpose. They are intended to help the client in obtaining an IP address lease from a DHCP server and to locate a valid network boot server. If the booting client, the DHCP server, and the network boot server are all located on the same network segment, usually no additional configuration is necessary. The DHCP broadcasts from the client reach both the DHCP server and the network boot server.

However, if either the DHCP server or the network boot server are on a different network segment than the client, or if they are on the same network segment but the network is controlled by a switch or a router, you may have to update the routing tables for the networking equipment in order to make sure that DHCP traffic is directed correctly.
Such a process is known as performing IP Helper table updates. When you perform this process, you must configure the networking equipment so that all DHCP broadcasts from the client computer are directed to both a valid DHCP server and to a valid network boot server.

Note: It is inefficient to rebroadcast the DHCP packets onto other network segments. It is best to only forward the DHCP packets to the recipients that are listed in the IP Helper table.

After the client computer has obtained an IP address, it contacts the network boot server directly in order to obtain the name and the path of the network boot file to download. Again, this process is handled by using DHCP packets.

Note: We recommend that you update the IP Helper tables in order to resolve scenarios in which the client computers and the network boot server are not located on the same network segment.

QUESTION NO: 7

Your company has 250 branch offices. Your network contains an Active Directory domain. The domain controllers run Windows Server 2008 R2. You plan to deploy Readonly Domain Controllers (RODCs) in the branch offices.

You need to plan the deployment of the RODCs to meet the following requirements:

- Build each RODC at the designated branch office.
- Ensure that the RODC installation source files do not contain cached secrets.
- Minimize the bandwidth used during the initial synchronization of Active Directory Domain Services (AD?DS).

What should you include in your plan?

A. Use Windows Server Backup to perform a full backup of an existing domain controller. Use the backup to build the new RODCs.
B. Use Windows Server Backup to perform a custom backup of the critical volumes of an existing domain controller. Use the backup to build the new RODCs.
C. Create a DFS namespace that contains the Active Directory database from one of the existing domain controllers. Build the RODCs by using an answer file.
D. Create an RODC installation media. Build the RODCs from the RODC installation media.

Answer: D
Explanation:
http://technet.microsoft.com/en-us/library/cc770654%28WS.10%29.aspx

Installing AD DS from Media
Applies To: Windows Server 2008, Windows Server 2008 R2
You can use the Ntdsutil.exe tool to create installation media for additional domain controllers that you are creating in a domain. By using the Install from Media (IFM) option, you can minimize the replication of directory data over the network. This helps you install additional domain controllers in remote sites more efficiently.

Ntdsutil.exe can create four types of installation media, as described in the following table.
You must use read-only domain controller (RODC) installation media to install an RODC. For RODC installation media, the ntdsutil command removes any cached secrets, such as passwords. You can create RODC installation media either on an RODC or on a writeable domain controller. You must use writeable domain controller installation media to install a writeable domain controller. You can create writeable domain controller installation media only on a writeable domain controller.

If the source domain controller where you create the installation media and the destination server where you plan to install ActiveDirectory Domain Services (ADDS) both run Windows Server2008 with Service Pack2 or later or Windows Server2008R2, and if you are using Distributed File System (DFS) Replication for SYSVOL, you can run the ntdsutil ifm command with an option to include the SYSVOL shared folder in the installation media. If the installation media includes SYSVOL, you must use Robocopy.exe to copy the installation media from the source domain controller to the destination server. For more information, see Installing an Additional Domain Controller by Using IFM.

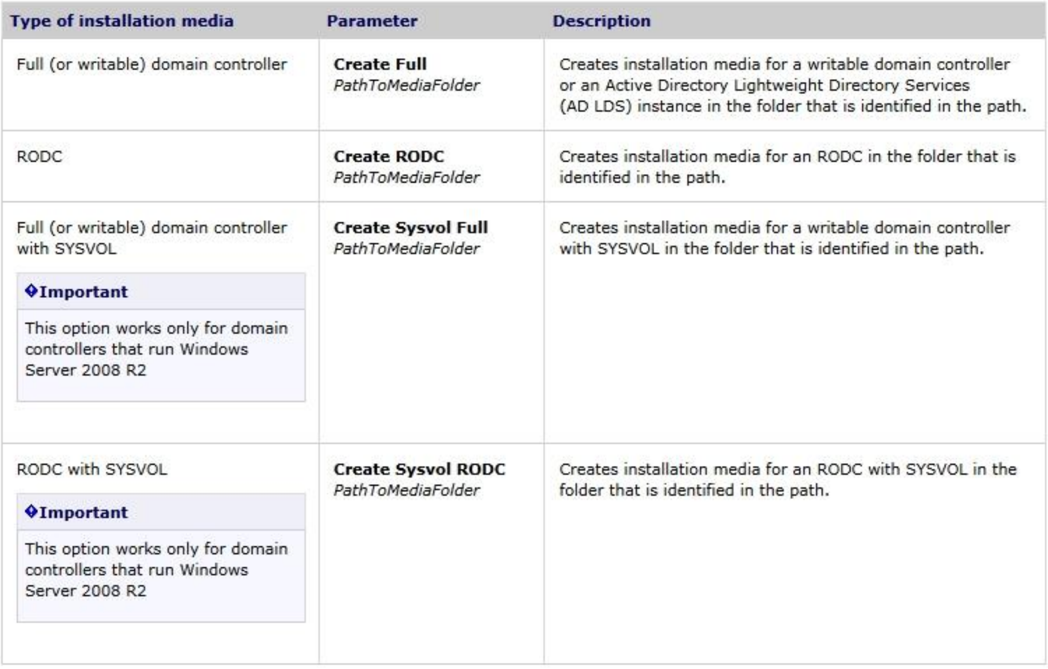

Type of installation media	Parameter	Description
Full (or writable) domain controller	**Create Full** *PathToMediaFolder*	Creates installation media for a writable domain controller or an Active Directory Lightweight Directory Services (AD LDS) instance in the folder that is identified in the path.
RODC	**Create RODC** *PathToMediaFolder*	Creates installation media for an RODC in the folder that is identified in the path.
Full (or writable) domain controller with SYSVOL **◆Important** This option works only for domain controllers that run Windows Server 2008 R2	**Create Sysvol Full** *PathToMediaFolder*	Creates installation media for a writable domain controller with SYSVOL in the folder that is identified in the path.
RODC with SYSVOL **◆Important** This option works only for domain controllers that run Windows Server 2008 R2	**Create Sysvol RODC** *PathToMediaFolder*	Creates installation media for an RODC with SYSVOL in the folder that is identified in the path.

QUESTION NO: 8

Your network consists of a single Active Directory domain. The network is located on the 172.16.0.0/23 subnet.

The company hires temporary employees. You provide user accounts and computers to the temporary employees. The temporary employees receive computers that are outside the Active

Directory domain. The temporary employees use their computers to connect to the network by using wired connections and wireless connections.

The company's security policy specifies that the computers connected to the network must have the latest updates for the operating system.

You need to plan the network's security so that it complies with the company's security policy.

What should you include in your plan?

A. Implement a Network Access Protection (NAP) strategy for the 172.16.0.0/23 subnet.
B. Create an extranet domain within the same forest. Migrate the temporary employees' user accounts to the extranet domain. Install the necessary domain resources on the 172.16.0.0/23 subnet.
C. Move the temporary employees' user accounts to a new organizational unit (OU). Create a new Group Policy object (GPO) that uses an intranet Microsoft Update server. Link the new GPO to the new OU.
D. Create a new subnet in a perimeter network. Relocate the wireless access point to the perimeter network. Require authentication through a VPN server before allowing access to the internal resources.

Answer: A

Explanation:

http://technet.microsoft.com/en-us/library/dd125338%28WS.10%29.aspx

Network Access Protection Design Guide

Updated: October 6, 2008

Applies To: Windows 7, Windows Server 2008, Windows Server 2008 R2, Windows Vista

Network Access Protection (NAP) is one of the most anticipated features of the WindowsServer®2008 operating system. NAP is a new platform that allows network administrators to define specific levels of network access based on a client’s identity, the groups to which the client belongs, and the degree to which the client complies with corporate governance policy. If a client is not compliant, NAP provides a mechanism for automatically bringing the client into compliance (a process known as remediation) and then dynamically increasing its level of network access. NAP is supported by Windows Server2008R2, Windows Server2008, Windows7, WindowsVista®, and Windows® XP with Service Pack 3 (SP3). NAP includes an application programming interface that developers and vendors can use to integrate their products and leverage this health state validation, access enforcement, and ongoing compliance evaluation. For more information about the NAP API, see Network Access Protection (http://go.microsoft.com/fwlink/?LinkId=128423).

The following are key NAP concepts:

NAP Agent.
A service included with Windows Server2008, WindowsVista, and Windows XP with SP3 that collects and manages health information for NAP client computers.

NAP client computer.
A computer that has the NAP Agent service installed and running, and is providing its health status to NAP server computers.
NAP-capable computer.
A computer that has the NAP Agent service installed and running and is capable of providing its health status to NAP server computers. NAP-capable computers include computers running Windows Server2008, WindowsVista, and Windows XP with SP3.
Non-NAP-capable computer. A computer that cannot provide its health status to NAP server components. A computer that has NAP agent installed but not running is also considered non-NAP-capable.
Compliant computer.
A computer that meets the NAP health requirements that you have defined for your network. Only NAP client computers can be compliant.
Noncompliant computer.
A computer that does not meet the NAP health requirements that you have defined for your network. Only NAP client computers can be noncompliant.
Health status.
Information about a NAP client computer that NAP uses to allow or restrict access to a network. Health is defined by a client computer's configuration state. Some common measurements of health include the operational status of Windows Firewall, the update status of antivirus signatures, and the installation status of security updates. A NAP client computer provides health status by sending a message called a statement of health (SoH).
NAP health policy server.
A NAP health policy server is a computer running Windows Server2008 with the Network Policy Server (NPS) role service installed and configured for NAP. The NAP health policy server uses NPS policies and settings to evaluate the health of NAP client computers when they request access to the network, or when their health state changes. Based on the results of this evaluation, the NAP health policy server instructs whether NAP client computers will be granted full or restricted access to the network.

QUESTION NO: 9

Your company has a main office and two branch offices. The main office is located in London. The branch offices are located in New York and Paris.

Your network consists of an Active Directory forest that contains three domains named contoso.com, paris.contoso.com, and newyork.contoso.com. All domain controllers run Windows Server 2008 R2 and have the DNS Server server role installed.

The domain controllers for contoso.com are located in the London office. The domain controllers for paris.contoso.com are located in the Paris office. The domain controllers for newyork.contoso.com are located in the New York office.

A domain controller in the contoso.com domain has a standard primary DNS zone for contoso.com. A domain controller in the paris.contoso.com domain has a standard primary DNS zone for paris.contoso.com. A domain controller in the newyork.contoso.com domain has a standard primary DNS zone for newyork.contoso.com.

You need to plan a name resolution strategy for the Paris office that meets the following requirements:

- If a WAN link fails, clients must be able to resolve hostnames for contoso.com.
- If a WAN link fails, clients must be able to resolve hostnames for newyork.contoso.com.
- The DNS servers in Paris must be updated when new authoritative DNS servers are added to newyork.contoso.com.

What should you include in your plan?

A. Configure conditional forwarding for contoso.com. Configure conditional forwarding for newyork.contoso.com.
B. Create a standard secondary zone for contoso.com. Create a standard secondary zone for newyork.contoso.com.
C. Convert the standard zone into an Active Directoryintegrated zone. Add all DNS servers in the forest to the root hints list.
D. Create an Active Directoryintegrated stub zone for contoso.com. Create an Active Directoryintegrated stub zone for newyork.contoso.com.

Answer: B

Explanation:

http://technet.microsoft.com/en-us/library/cc771640.aspx

http://technet.microsoft.com/en-us/library/cc771898.aspx

Understanding Zone Delegation

Applies To: Windows Server 2008, Windows Server 2008 R2

Domain Name System (DNS) provides the option of dividing up the namespace into one or more zones, which can then be stored, distributed, and replicated to other DNS servers. When you are deciding whether to divide your DNS namespace to make additional zones, consider the following reasons to use additional zones:

- You want to delegate management of part of your DNS namespace to another location or department in your organization.
- You want to divide one large zone into smaller zones to distribute traffic loads among multiple servers, improve DNS name resolution performance, or create a more-fault-tolerant DNS environment.
- You want to extend the namespace by adding numerous subdomains at once, for example, to accommodate the opening of a new branch or site.

Secondary zone

When a zone that this DNS server hosts is a secondary zone, this DNS server is a secondary source for information about this zone. The zone at this server must be obtained from another remote DNS server computer that also hosts the zone. This DNS server must have network access to the remote DNS server that supplies this server with updated information about the zone. Because a secondary zone is merely a copy of a primary zone that is hosted on another server, it cannot be stored in AD DS.

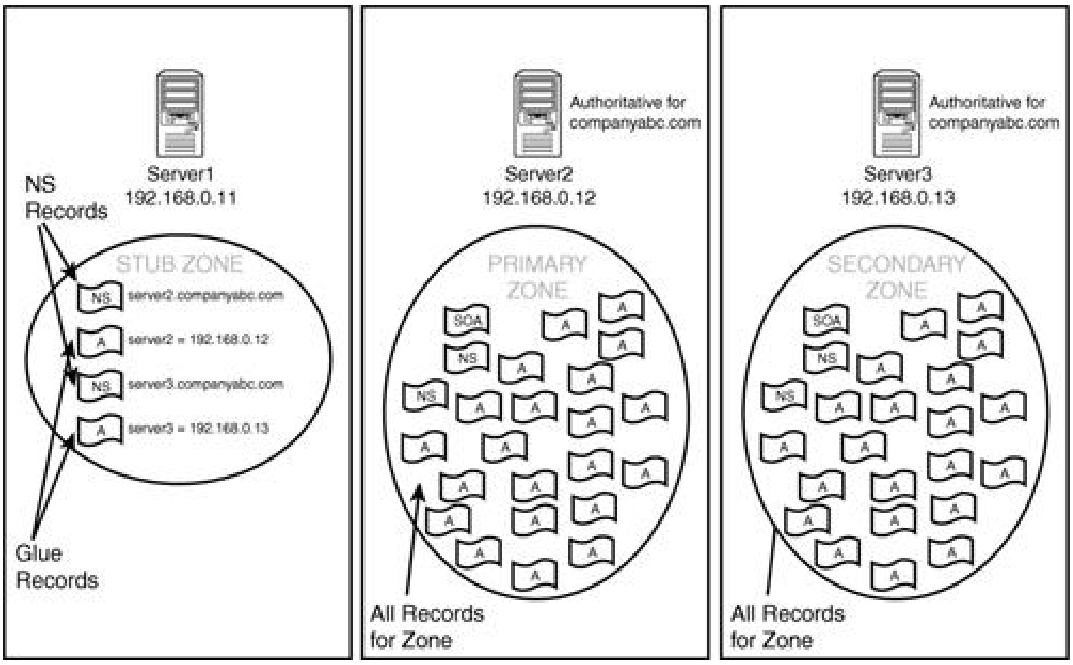

QUESTION NO: 10

Your network consists of a single Active Directory domain. All domain controllers run Windows Server 2008 R2.

You need to implement a Certificate Services solution that meets the following requirements:

- Automates the distribution of certificates for internal users
- Ensures that the network's certificate infrastructure is as secure as possible
- Gives external users access to resources that use certificate based authentication

What should you do?

A. Deploy an online standalone root certification authority (CA). Deploy an offline standalone root CA.
B. Deploy an offline enterprise root certification authority (CA). Deploy an offline enterprise subordinate CA.
C. Deploy an offline standalone root certification authority (CA). Deploy an online enterprise subordinate CA. Deploy an online standalone subordinate CA.
D. Deploy an online standalone root certification authority (CA). Deploy an online enterprise subordinate CA. Deploy an online standalone subordinate CA.

Answer: C

Explanation:

Certification authority hierarchies

The Microsoft public key infrastructure (PKI) supports a hierarchical certification authority (CA) model. A certification hierarchy provides scalability, ease of administration, and consistency with a growing number of commercial and other CA products.

In its simplest form, a certification hierarchy consists of a single CA. However, in general, a hierarchy will contain multiple CAs with clearly defined parent-child relationships. In this model, the child subordinate certification authorities are certified by their parent CA-issued certificates, which bind a certification authority's public key to its identity. The CA at the top of a hierarchy is referred to as the root authority, or root CA. The child CAs of the root CAs are called subordinate certification authorities (CAs).

A root certification authority (CA) is the top of a public key infrastructure (PKI) and generates a self-signed certificate. This means that the root CA is validating itself (self-validating). This root CA could then have subordinate CAs that effectively trust it. The subordinate CAs receive a certificate signed by the root CA, so the subordinate CAs can issue certificates that are validated by the root CA. This establishes a CA hierarchy and trust path.
http://social.technet.microsoft.com/wiki/contents/articles/2900.offline-root-certification-authority-ca.aspx

Certification authority hierarchies

The Microsoft public key infrastructure (PKI) supports a hierarchical certification authority (CA) model. A certification hierarchy provides scalability, ease of administration, and consistency with a growing number of commercial and other CA products.
In its simplest form, a certification hierarchy consists of a single CA. However, in general, a

hierarchy will contain multiple CAs with clearly defined parent-child relationships. In this model, the child subordinate certification authorities are certified by their parent CA-issued certificates, which bind a certification authority's public key to its identity. The CA at the top of a hierarchy is referred to as the root authority, or root CA. The child CAs of the root CAs are called subordinate certification authorities (CAs).

Authentication and Authorization

Stand-alone CAs use local authentication for certificate requests, mainly through the Web enrollment interface.
Stand-alone CAs provide an ideal service provider or commercial PKI provider platform for issuing certificates to users outside of an Active Directory environment where the user identity is separately verified and examined before the request is submitted to the CA.

Offline and Online CAs

Traditionally, the decision of whether to use either an online or offline CAs involves a compromise between availability and usability versus security. The more sensitive that the key material is and the higher the security requirements are, the less accessible the CA should be to users.

Specifying CA Roles

An ideal PKI hierarchy design divides the responsibility of the CAs. A topology that is designed with requirements that have been carefully considered provides the most flexible and scalable enterprise configuration. In general, CAs are organized in hierarchies. Single tier hierarchies might not provide adequate security compartmentalization, extensibility and flexibility. Hierarchies with more than three tiers might not provide additional value regarding security, extensibility and flexibility.

The most important consideration is protecting the highest instance of trust as much as possible. Single-tier hierarchies are based on the need to compartmentalize risk and reduce the attack surface that is available to users who have malicious intent. A larger hierarchy is much more difficult to administer, with little security benefit.

Depending on the organization's necessities, a PKI should consist of two or three logical levels that link several CAs in a hierarchy. Administrators who understand the design requirements for a three-level topology may also be able to build a two-level topology.
A three-tier CA hierarchy consists of the following components:

A root CA that is configured as a stand-alone CA without a network connection
One or more intermediate CAs that are configured as stand-alone CAs without a network connection
One or more issuing CAs that are configured as enterprise CAs that are connected to the network

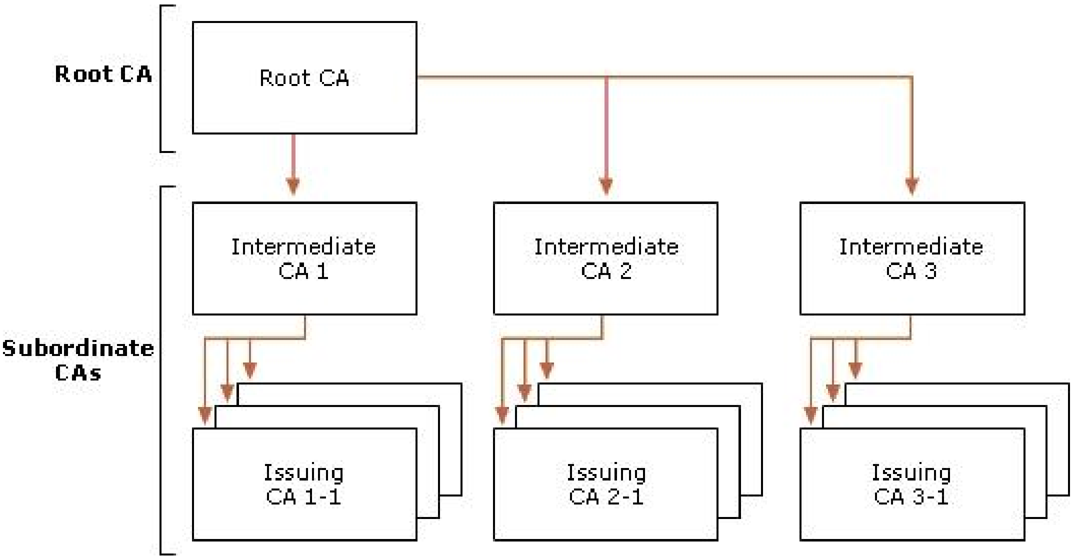

Also worth a look though it refers to windows 2003
http://technet.microsoft.com/en-us/library/cc779714%28WS.10%29.aspx

QUESTION NO: 11

Your network contains an Active Directory forest named contoso.com.

You plan to deploy a new child domain named branch.contoso.com. The child domain will contain two domain controllers. Both domain controllers will have the DNS Server server role installed. All users and computers in the branch office will be members of the branch.contoso.com domain.

You need to plan the DNS infrastructure for the child domain to meet the following requirements:

- Ensure resources in the root domain are accessible by fully qualified domain names.
- Ensure resources in the child domain are accessible by fully qualified domain names.
- Provide name resolution services in the event that a single server fails for a prolonged period of time.
- Automatically recognize when new DNS servers are added to or removed from the contoso.com domain.

What should you include in your plan?

A. On both domain controllers, add a conditional forwarder for contoso.com and create a standard primary zone for branch.contoso.com.
B. On both domain controllers, modify the root hints to include the domain controllers for contoso.com. On one domain controller, create an Active Directory-integrated zone for

branch.contoso.com.
C. On one domain controller create an Active Directory-integrated zone for branch.contoso.com and create an Active Directory-integrated stub zone for contoso.com.
D. On one domain controller, create a standard primary zone for contoso.com. On the other domain controller, create a standard secondary zone for contoso.com.

Answer: C

Explanation:

http://technet.microsoft.com/en-us/library/cc772101.aspx

http://technet.microsoft.com/en-us/library/cc771898.aspx

Understanding DNS Zone Replication in Active Directory Domain Services

Applies To: Windows Server 2008, Windows Server 2008 R2

You can store Domain Name System (DNS) zones in the domain or application directory partitions of Active Directory Domain Services (AD DS). A partition is a data structure in AD DS that distinguishes data for different replication purposes. For more information, see Understanding Active Directory Domain Services Integration.
The following table describes the available zone replication scopes for AD DS-integrated DNS zone data.

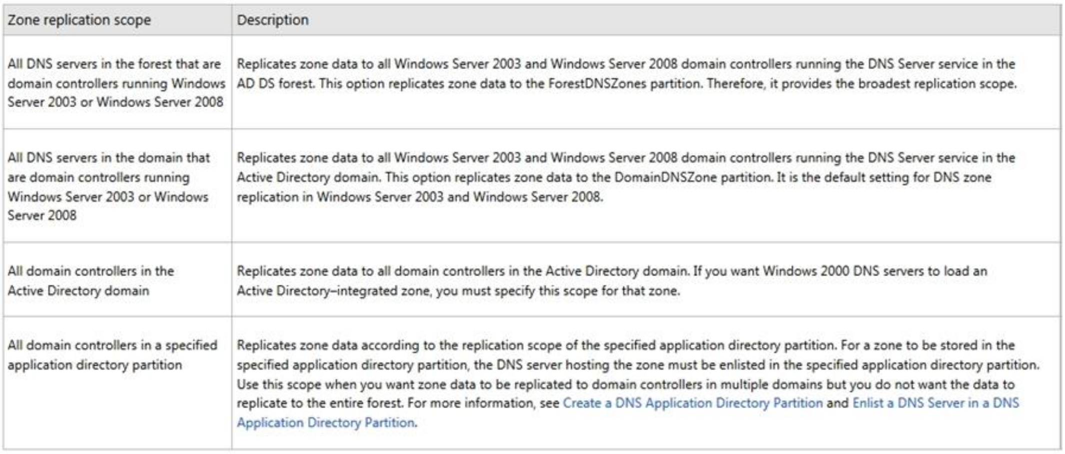

Zone replication scope	Description
All DNS servers in the forest that are domain controllers running Windows Server 2003 or Windows Server 2008	Replicates zone data to all Windows Server 2003 and Windows Server 2008 domain controllers running the DNS Server service in the AD DS forest. This option replicates zone data to the ForestDNSZones partition. Therefore, it provides the broadest replication scope.
All DNS servers in the domain that are domain controllers running Windows Server 2003 or Windows Server 2008	Replicates zone data to all Windows Server 2003 and Windows Server 2008 domain controllers running the DNS Server service in the Active Directory domain. This option replicates zone data to the DomainDNSZone partition. It is the default setting for DNS zone replication in Windows Server 2003 and Windows Server 2008.
All domain controllers in the Active Directory domain	Replicates zone data to all domain controllers in the Active Directory domain. If you want Windows 2000 DNS servers to load an Active Directory–integrated zone, you must specify this scope for that zone.
All domain controllers in a specified application directory partition	Replicates zone data according to the replication scope of the specified application directory partition. For a zone to be stored in the specified application directory partition, the DNS server hosting the zone must be enlisted in the specified application directory partition. Use this scope when you want zone data to be replicated to domain controllers in multiple domains but you do not want the data to replicate to the entire forest. For more information, see Create a DNS Application Directory Partition and Enlist a DNS Server in a DNS Application Directory Partition.

When you decide which replication scope to choose, consider that the broader the replication scope, the greater the network traffic caused by replication. For example, if you decide to have AD DS-integrated DNS zone data replicated to all DNS servers in the forest, this will produce greater network traffic than replicating the DNS zone data to all DNS servers in a single AD DS domain in that forest.

AD DS-integrated DNS zone data that is stored in an application directory partition is not replicated to the global catalog for the forest The domain controller that contains the global catalog can also host application directory partitions, but it will not replicate this data to its global catalog.

AD DS-integrated DNS zone data that is stored in a domain partition is replicated to all domain

controllers in its AD DS domain, and a portion of this data is stored in the global catalog. This setting is used to support Windows 2000.
If an application directory partition's replication scope replicates across AD DS sites, replication will occur with the same intersite replication schedule as is used for domain partition data.

By default, the Net Logon service registers domain controller locator (Locator) DNS resource records for the application directory partitions that are hosted on a domain controller in the same manner as it registers domain controller locator (Locator) DNS resource records for the domain partition that is hosted on a domain controller.

Primary zone
When a zone that this DNS server hosts is a primary zone, the DNS server is the primary source for information about this zone, and it stores the master copy of zone data in a local file or in AD DS. When the zone is stored in a file, by default the primary zone file is named rone_name.dns and it is located in the %windir%\System32\Dns folder on the server.
Secondary zone
When a zone that this DNS server hosts is a secondary zone, this DNS server is a secondary source for information about this zone. The zone at this server must be obtained from another remote DNS server computer that also hosts the zone. This DNS server must have network access to the remote DNS server that supplies this server with updated information about the zone. Because a secondary zone is merely a copy of a primary zone that is hosted on another server, it cannot be stored in AD DS.
Stub zone
When a zone that this DNS server hosts is a stub zone, this DNS server is a source only for information about the authoritative name servers for this zone. The zone at this server must be obtained from another DNS server that hosts the zone. This DNS server must have network access to the remote DNS server to copy the authoritative name server information about the zone.
You can use stub zones to:
• Keep delegated zone information current. By updating a stub zone for one of its child zones regularly, the DNS server that hosts both the parent zone and the stub zone will maintain a current list of authoritative DNS servers for the child zone.
• Improve name resolution. Stub zones enable a DNS server to perform recursion using the stub zone's list of name servers, without having to query the Internet or an internal root server for the DNS namespace.
• Simplify DNS administration. By using stub zones throughout your DNS infrastructure, you can distribute a list of the authoritative DNS servers for a zone without using secondary zones. However, stub zones do not serve the same purpose as secondary zones, and they are not an alternative for enhancing redundancy and load sharing.

There are two lists of DNS servers involved in the loading and maintenance of a stub zone:

- The list of master servers from which the DNS server loads and updates a stub zone. A master server may be a primary or secondary DNS server for the zone. In both cases, it will have a complete list of the DNS servers for the zone.
- The list of the authoritative DNS servers for a zone. This list is contained in the stub zone using name server (NS) resource records.

When a DNS server loads a stub zone, such as widgets.tailspintoys.com, it quenes the master servers, which can be in different locations, for the necessary resource records of the authoritative servers for the zone widgets.tailspintoys.com. The list of master servers may contain a single server or multiple servers, and it can be changed anytime.

QUESTION NO: 12

Your network is configured as shown in the following diagram.

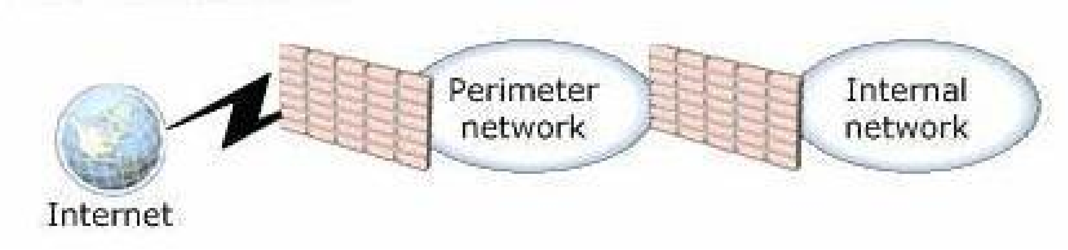

You deploy an enterprise certification authority (CA) on the internal network. You also deploy a Microsoft Online Responder on the internal network. You need to recommend a secure method for Internet users to verify the validity of individual certificates.

The solution must minimize network bandwidth.

What should you recommend?

A. Deploy a subordinate CA on the perimeter network.
B. Install a standalone CA and the Network Device Enrollment Service (NDES) on a server on the perimeter network.
C. Install a Network Policy Server (NPS) on a server on the perimeter network. Redirect authentication requests to a server on the internal network.
D. Install Microsoft Internet Information Services (IIS) on a server on the perimeter network. Configure IIS to redirect requests to the Online Responder on the internal network.

Answer: D
Explanation:

http://www.ipsure.com/blog/2010/installation-and-configuration-of-active-directory-certificate-services-onwindows-server-2008-r2-1/
http://msdn.microsoft.com/en-us/library/cc732956.aspx

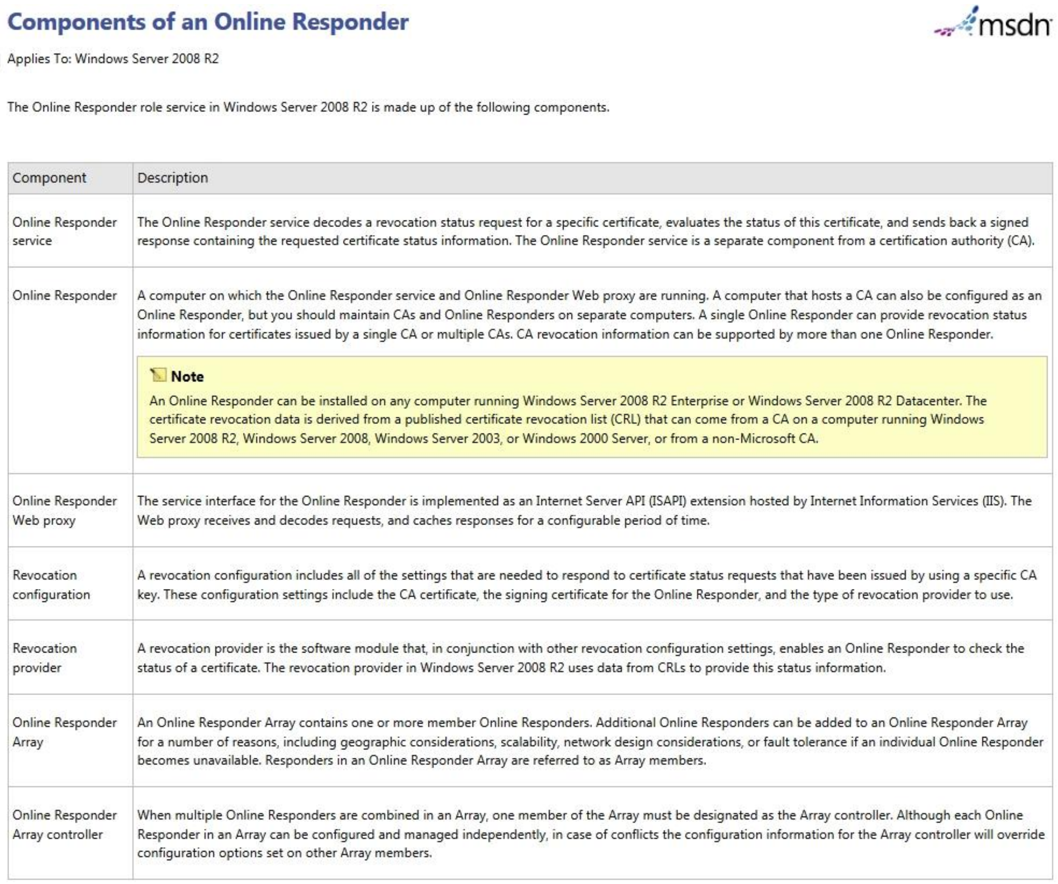

Components of an Online Responder

Applies To: Windows Server 2008 R2

The Online Responder role service in Windows Server 2008 R2 is made up of the following components.

Component	Description
Online Responder service	The Online Responder service decodes a revocation status request for a specific certificate, evaluates the status of this certificate, and sends back a signed response containing the requested certificate status information. The Online Responder service is a separate component from a certification authority (CA).
Online Responder	A computer on which the Online Responder service and Online Responder Web proxy are running. A computer that hosts a CA can also be configured as an Online Responder, but you should maintain CAs and Online Responders on separate computers. A single Online Responder can provide revocation status information for certificates issued by a single CA or multiple CAs. CA revocation information can be supported by more than one Online Responder. **Note** An Online Responder can be installed on any computer running Windows Server 2008 R2 Enterprise or Windows Server 2008 R2 Datacenter. The certificate revocation data is derived from a published certificate revocation list (CRL) that can come from a CA on a computer running Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, or Windows 2000 Server, or from a non-Microsoft CA.
Online Responder Web proxy	The service interface for the Online Responder is implemented as an Internet Server API (ISAPI) extension hosted by Internet Information Services (IIS). The Web proxy receives and decodes requests, and caches responses for a configurable period of time.
Revocation configuration	A revocation configuration includes all of the settings that are needed to respond to certificate status requests that have been issued by using a specific CA key. These configuration settings include the CA certificate, the signing certificate for the Online Responder, and the type of revocation provider to use.
Revocation provider	A revocation provider is the software module that, in conjunction with other revocation configuration settings, enables an Online Responder to check the status of a certificate. The revocation provider in Windows Server 2008 R2 uses data from CRLs to provide this status information.
Online Responder Array	An Online Responder Array contains one or more member Online Responders. Additional Online Responders can be added to an Online Responder Array for a number of reasons, including geographic considerations, scalability, network design considerations, or fault tolerance if an individual Online Responder becomes unavailable. Responders in an Online Responder Array are referred to as Array members.
Online Responder Array controller	When multiple Online Responders are combined in an Array, one member of the Array must be designated as the Array controller. Although each Online Responder in an Array can be configured and managed independently, in case of conflicts the configuration information for the Array controller will override configuration options set on other Array members.

QUESTION NO: 13

Your network contains two DHCP servers. The DHCP servers are named DHCP1 and DHCP2. The internal network contains 1,000 DHCP client computers that are located on a single subnet. A router separates the internal network from the Internet. The router has a single IP address on the internal interface.

DHCP1 has the following scope information:

- Starting IP address: 172.16.0.1
- Ending IP address: 172.16.7.255
- Subnet mask: 255.255.240.0

You need to provide a fault tolerant DHCP infrastructure that supports the client computers on the internal network. In the event that a DHCP server fails, all client computers must be able to obtain a valid IP address.

How should you configure DHCP2?

A. Create a scope for the subnet 172.16.0.0/20. Configure the scope to use a starting IP address of 172.16.8.1 and an ending IP address of 172.16.15.254.
B. Create a scope for the subnet 172.16.0.0/21. Configure the scope to use a starting IP address of 172.16.0.1 and an ending IP address of 172.16.15.254.
C. Create a scope for the subnet 172.16.8.0/21. Configure the scope to use a starting IP address of 172.16.8.1 and an ending IP address of 172.16.10.254.
D. Create a scope for the subnet 172.17.0.0/16. Configure the scope to use a starting IP address of 172.17.0.1 and an ending IP address of 172.17.255.254.

Answer: A

Explanation:

Create a scope for the subnet 172.16.0.0/20.

Configure the scope to use a starting IP address of 172.16.8.1 and an ending IP address of 172.16.15.254.

Subnet 255.255.240.0 is a /20 subnet in CIDR notation, this allows for 4096 client IPs, ranging from 172.16.0.1 all the way to 172.16.15.254 as DHCP1 only used half of the available IPs then you should configure DHCP2 to use the other half.

http://en.wikipedia.org/wiki/Classless_Inter-Domain_Routing as an aside you could consider the 80/20 design rule for balancing scope distribution of addresses where multiple DHCP servers are deployed to service the same scope.

Using more than one DHCP server on the same subnet provides increased fault tolerance for servicing DHCP clients located on it. With two DHCP servers, if one server is unavailable, the other server can take its place and continue to lease new addresses or renew existing clients.

A common practice when balancing a single network and scope range of addresses between two DHCP servers is to have 80 percent of the addresses distributed by one DHCP server and the remaining 20 percent provided by a second.

QUESTION NO: 14

Your company has a main office and three branch offices. The network consists of a single Active Directory domain. Each office contains an Active Directory domain controller.

You need to create a DNS infrastructure for the network that meets the following requirements:

- The DNS infrastructure must allow the client computers in each office to register DNS names within their respective offices.
- The client computers must be able to resolve names for hosts in all offices.

What should you do?

A. Create an Active Directory-integrated zone at the main office site.
B. Create a standard primary zone at the main office site and at each branch office site.
C. Create a standard primary zone at the main office site. Create a secondary zone at each branch office site.
D. Create a standard primary zone at the main office site. Create an Active Directory-integrated stub zone at each branch office site.

Answer: A

Explanation:

http://searchwindowsserver.techtarget.com/tip/DNS-Primer-Tips-for-understanding-Active-Directory-integratedzone-design-and-configuration

http://technet.microsoft.com/en-us/library/cc772101.aspx

In an ADI primary zone, rather than keeping the old zone file on a disk, the DNS records are stored in the AD, and Active Directory replication is used rather than the old problematic zone transfer. If all DNS servers were to die or become inaccessible, you could simply install DNS on any domain controller (DC) in the domain. The records would be automatically populated and your DNS server would be up without the messy import/export tasks of standard DNS zone files. Windows 2000 and 2003 allow you to put a standard secondary zone (read only) on a member server and use one of the ADI primary servers as the master.

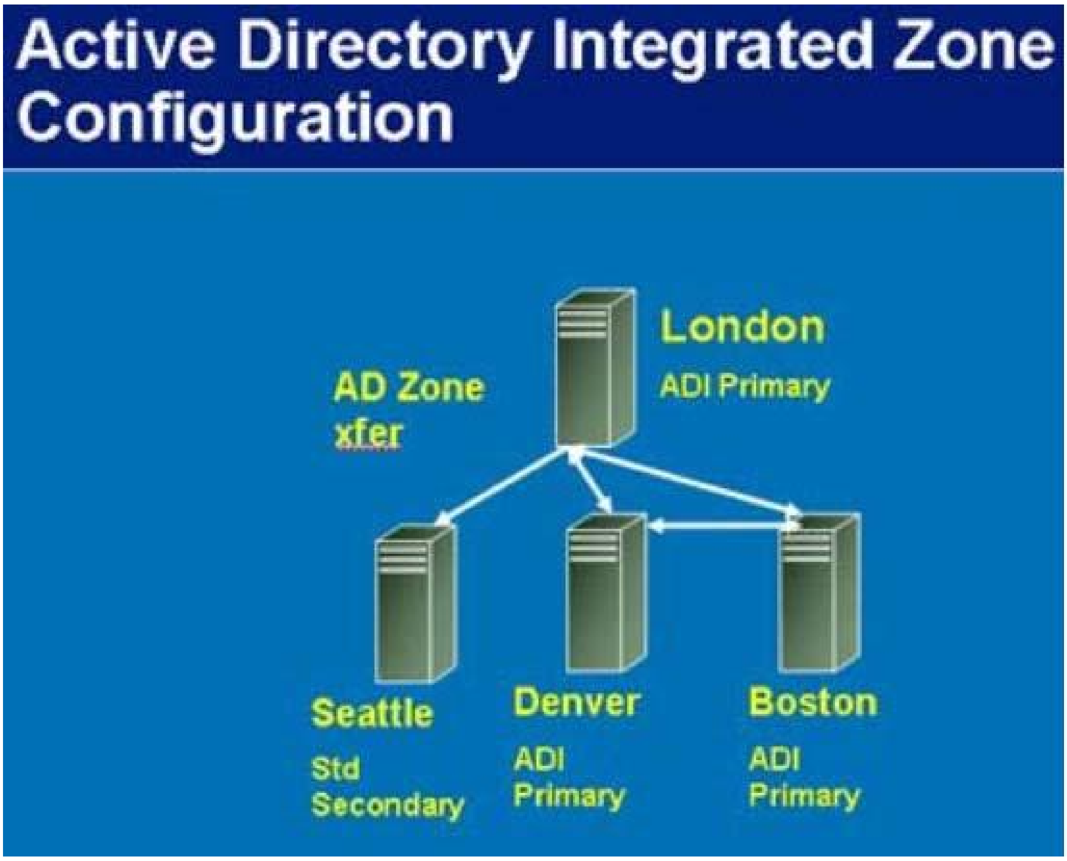

Active Directory Integrated Zone Configuration

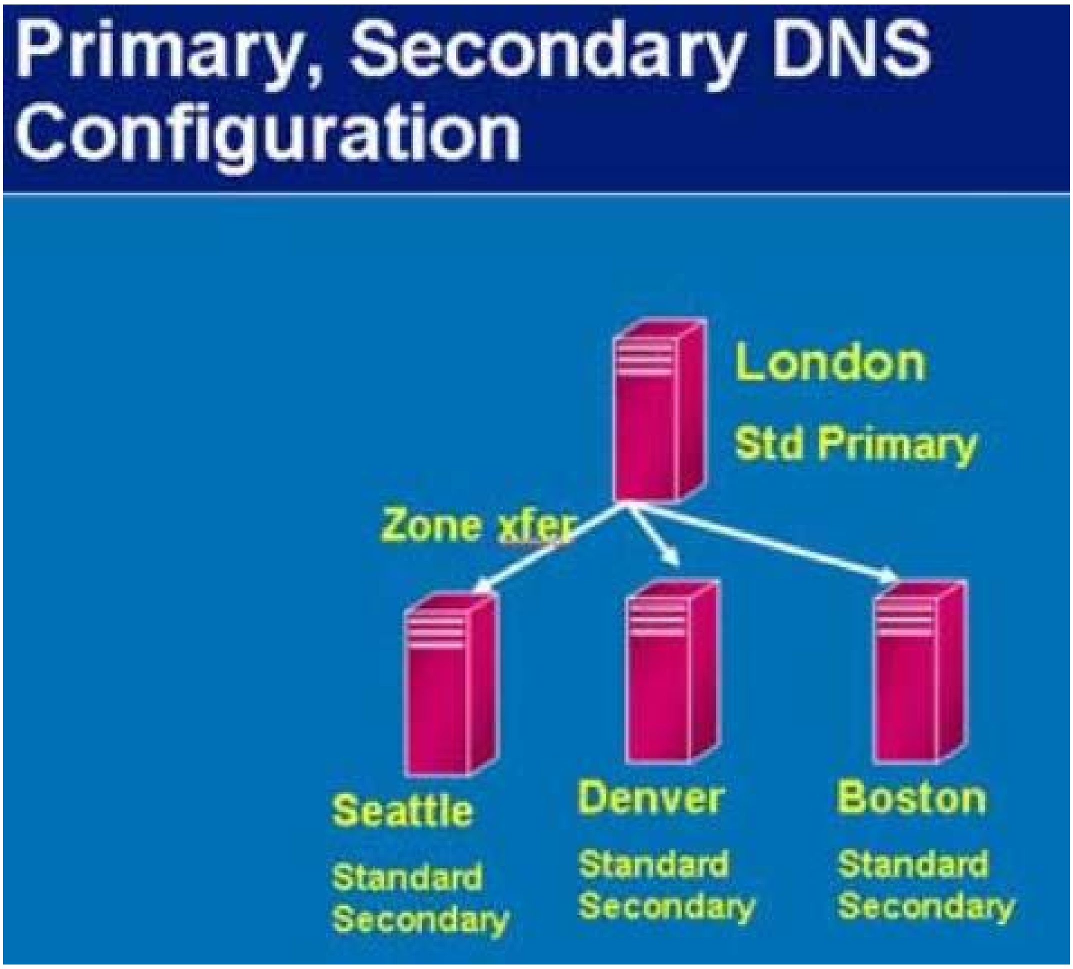

Primary, Secondary DNS Configuration

London

Std Primary

Zone xfer

Seattle

Standard Secondary

Denver

Standard Secondary

Boston

Standard Secondary

When you decide which replication scope to choose, consider that the broader the replication scope, the greater the network traffic caused by replication. For example, if you decide to have AD DS-integrated DNS zone data replicated to all DNS servers in the forest, this will produce greater network traffic than replicating the DNS zone data to all DNS servers in a single AD DS domain in that forest.

AD DS-integrated DNS zone data that is stored in an application directory partition is not replicated to the global catalog for the forest. The domain controller that contains the global catalog can also host application directory partitions, but it will not replicate this data to its global catalog.

AD DS-integrated DNS zone data that is stored in a domain partition is replicated to all domain controllers in its AD DS domain, and a portion of this data is stored in the global catalog. This setting is used to support Windows 2000.

If an application directory partition's replication scope replicates across AD DS sites, replication will occur with the same intersite replication schedule as is used for domain partition data.

By default, the Net Logon service registers domain controller locator (Locator) DNS resource records for the application directory partitions that are hosted on a domain controller in the same manner as it registers domain controller locator (Locator) DNS resource records for the domain partition that is hosted on a domain controller.

QUESTION NO: 15

Your network consists of a single Active Directory domain. The network contains two Windows Server 2008 R2 computers named Server1 and Server2. The company has two identical print devices. You plan to deploy print services.

You need to plan a print services infrastructure to meet the following requirements:

- Manage the print queue from a central location.
- Make the print services available, even if one of the print devices fails.

What should you include in your plan?

A. Install and share a printer on Server1. Enable printer pooling.
B. Install the Remote Desktop Services server role on both servers. Configure Remote Desktop Connection Broker (RD Connection Broker).
C. Install and share a printer on Server1. Install and share a printer on Server2. Use Print Management to install the printers on the client computers.
D. Add Server1 and Server2 to a Network Load Balancing cluster. Install a printer on each node of the cluster.

Answer: A
Explanation:
http://www.techrepublic.com/blog/datacenter/configure-printer-pooling-in-windows-server-2008/964

Managing printers can be the bane of a Windows administrator. One feature that may assist you with this task is the Windows printer pooling feature. Windows Server 2008 offers functionality that permits a collection of multiple like-configured printers to distribute the print workload.

Printer pooling makes one share that clients print to, and the jobs are sent to the first available printer. Configuring print pooling is rather straightforward in the Windows printer configuration applet of the Control Panel. Figure A shows two like-modeled printers being pooled.

To use pooling, the printer models need to be the same so that the driver configuration is transparent to the end device; this can also help control costs of toner and other supplies. But plan accordingly — you don’t want users essentially running track to look for their print jobs on every printer in the office.

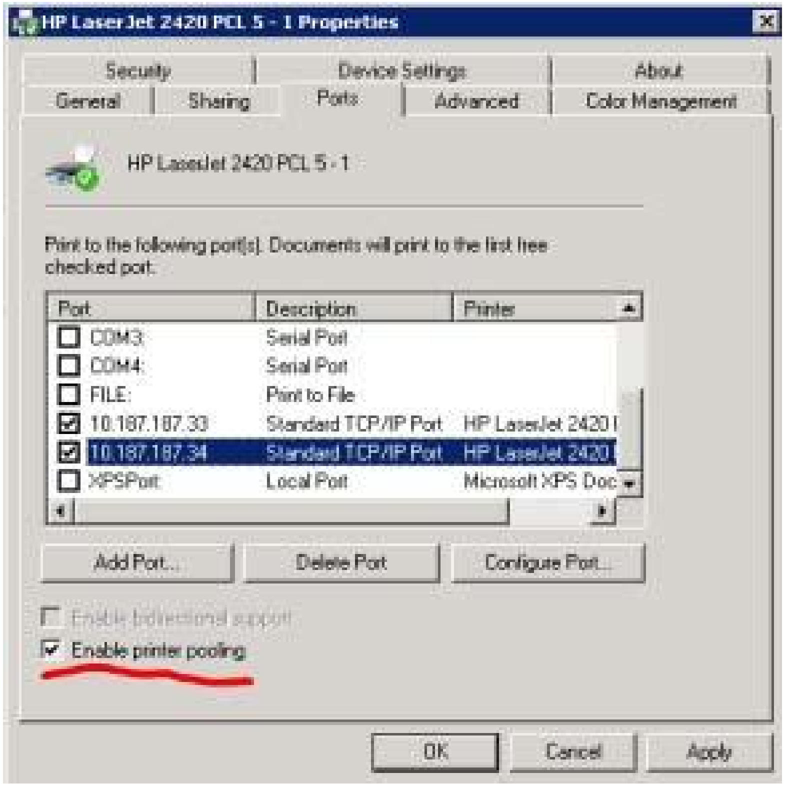

HP LaserJet 2420 PCL 5 - 1

Print to the following port(s). Documents will print to the first free checked port.

Port	Description	Printer
□ COM3:	Serial Port	
□ COM4:	Serial Port	
□ FILE:	Print to File	
☑ 10.187.187.33	Standard TCP/IP Port	HP LaserJet 2420 I
☑ 10.187.187.34	Standard TCP/IP Port	HP LaserJet 2420
□ XPSPort:	Local Port	Microsoft XPS Doc

□ Enable bidirectional support

☑ Enable printer pooling

QUESTION NO: 16

Your network contains two servers that run the Server Core installation of Windows Server 2008 R2. The two servers are part of a Network Load Balancing cluster.

The cluster hosts a Web site. Administrators use client computers that run Windows 7.

You need to recommend a strategy that allows the administrators to remotely manage the Network Load Balancing cluster. Your strategy must support automation.

What should you recommend?

A. On the servers, enable Windows Remote Management (WinRM).
B. On the servers, add the administrators to the Remote Desktop Users group.
C. On the Windows 7 client computers, enable Windows Remote Management (WinRM).
D. On the Windows 7 client computers, add the administrators to the Remote Desktop Users group.

Answer: A
Explanation:

http://support.microsoft.com/kb/968929
http://msdn.microsoft.com/en-us/library/aa384291%28VS.85%29.aspx

WinRM 2.0

WinRM is the Microsoft implementation of WS-Management Protocol, a standard Simple Object Access Protocol (SOAP)-based, firewall-friendly protocol that allows for hardware and operating systems from different vendors to interoperate. The WS-Management Protocol specification provides a common way for systems to access and exchange management information across an IT infrastructure.

WinRM 2.0 includes the following new features:

• The WinRM Client Shell API provides functionality to create and manage shells and shell operations, commands, and data streams on remote computers.
• The WinRM Plug-in API provides functionality that enables a user to write plug-ins by implementing certain APIs for supported resources and operations.
• WinRM 2.0 introduces a hosting framework. Two hosting models are supported. One is Internet Information Services (HS)-based and the other is WinRM service-based.
• Association traversal lets a user retrieve instances of Association classes by using a standard filtering mechanism.
• WinRM 2.0 supports delegating user credentials across multiple remote computers.
• Users of WinRM 2.0 can use Windows PowerShell cmdlets for system management.
• WinRM has added a specific set of quotas that provide a better quality of service and allocate server resources to concurrent users. The WinRM quota set is based on the quota infrastructure that is implemented for the IIS service.

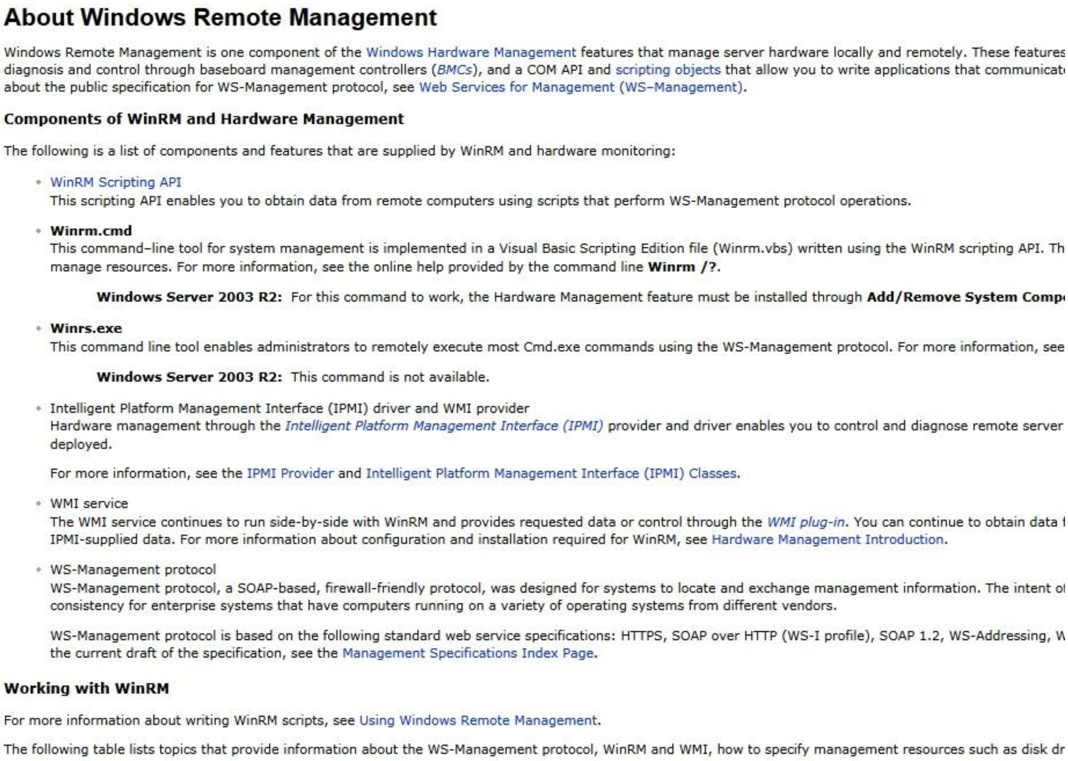

About Windows Remote Management

Windows Remote Management is one component of the Windows Hardware Management features that manage server hardware locally and remotely. These features
diagnosis and control through baseboard management controllers (*BMCs*), and a COM API and scripting objects that allow you to write applications that communicat
about the public specification for WS-Management protocol, see Web Services for Management (WS–Management).

Components of WinRM and Hardware Management

The following is a list of components and features that are supplied by WinRM and hardware monitoring:

- WinRM Scripting API
  This scripting API enables you to obtain data from remote computers using scripts that perform WS-Management protocol operations.
- **Winrm.cmd**
  This command–line tool for system management is implemented in a Visual Basic Scripting Edition file (Winrm.vbs) written using the WinRM scripting API. Th
  manage resources. For more information, see the online help provided by the command line **Winrm /?**.

  Windows Server 2003 R2: For this command to work, the Hardware Management feature must be installed through **Add/Remove System Comp**
- **Winrs.exe**
  This command line tool enables administrators to remotely execute most Cmd.exe commands using the WS-Management protocol. For more information, see

  Windows Server 2003 R2: This command is not available.
- Intelligent Platform Management Interface (IPMI) driver and WMI provider
  Hardware management through the *Intelligent Platform Management Interface (IPMI)* provider and driver enables you to control and diagnose remote server
  deployed.

  For more information, see the IPMI Provider and Intelligent Platform Management Interface (IPMI) Classes.
- WMI service
  The WMI service continues to run side-by-side with WinRM and provides requested data or control through the *WMI plug-in*. You can continue to obtain data
  IPMI-supplied data. For more information about configuration and installation required for WinRM, see Hardware Management Introduction.
- WS-Management protocol
  WS-Management protocol, a SOAP-based, firewall-friendly protocol, was designed for systems to locate and exchange management information. The intent o
  consistency for enterprise systems that have computers running on a variety of operating systems from different vendors.

  WS-Management protocol is based on the following standard web service specifications: HTTPS, SOAP over HTTP (WS-I profile), SOAP 1.2, WS-Addressing, W
  the current draft of the specification, see the Management Specifications Index Page.

Working with WinRM

For more information about writing WinRM scripts, see Using Windows Remote Management.

The following table lists topics that provide information about the WS-Management protocol, WinRM and WMI, how to specify management resources such as disk dr

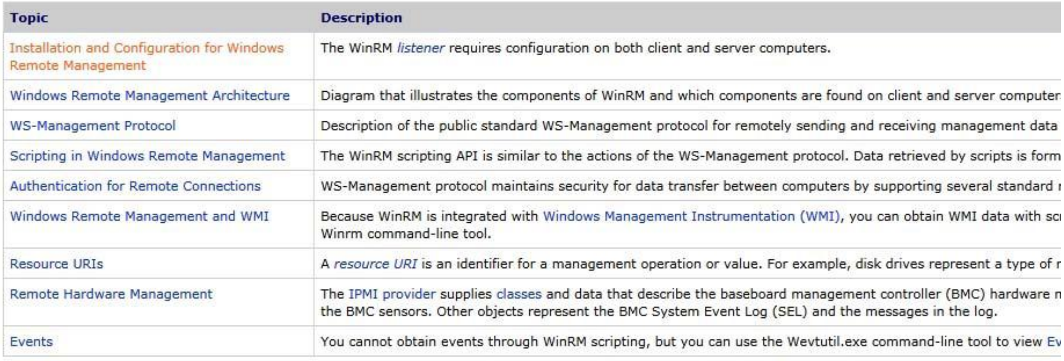

Topic	Description
Installation and Configuration for Windows Remote Management	The WinRM *listener* requires configuration on both client and server computers.
Windows Remote Management Architecture	Diagram that illustrates the components of WinRM and which components are found on client and server computer
WS-Management Protocol	Description of the public standard WS-Management protocol for remotely sending and receiving management data
Scripting in Windows Remote Management	The WinRM scripting API is similar to the actions of the WS-Management protocol. Data retrieved by scripts is form
Authentication for Remote Connections	WS-Management protocol maintains security for data transfer between computers by supporting several standard
Windows Remote Management and WMI	Because WinRM is integrated with Windows Management Instrumentation (WMI), you can obtain WMI data with sc Winrm command-line tool.
Resource URIs	A *resource URI* is an identifier for a management operation or value. For example, disk drives represent a type of
Remote Hardware Management	The IPMI provider supplies classes and data that describe the baseboard management controller (BMC) hardware the BMC sensors. Other objects represent the BMC System Event Log (SEL) and the messages in the log.
Events	You cannot obtain events through WinRM scripting, but you can use the Wevtutil.exe command-line tool to view Ev

USAGE
=====
(ALL UPPER-CASE = value that must be supplied by user.)
winrs [-/SWITCH[:VALUE]] COMMAND
COMMAND - Any string that can be executed as a command in the cmd.exe shell.

SWITCHES
========
(All switches accept both short form or long form. For example both -r and
-remote are valid.)

-r[emote]:ENDPOINT - The target endpoint using a NetBIOS name or the standard connect
ion URL: [TRANSPORT://]TARGET[:PORT]. If not specified
-r:localhost is used.

-un[encrypted] - Specify that the messages to the remote shell will not be encrypted. This is useful for troubleshooting, or when the network traffic is already encrypted using ipsec, or when physical security is enforced. By default the messages are encrypted
using Kerberos or NTLM keys. This switch is ignored when HTTPS transport is selected.
-u[sername]:USERNAME - Specify username on command line. If not specified the tool will
use Negotiate authentication or prompt for the name.
If -username is specified, -password must be as well.
-p[assword]:PASSWORD - Specify password on command line. If -password is not specified but -username is the tool will prompt for the password. If -password is specified, -user must be specified as well.
-t[imeout]:SECONDS - This option is deprecated.
-d[irectory]:PATH - Specifies starting directory for remote shell. If not specified the remote shell will start in the user's home directory defined by the environment variable %USERPROFILE%.
-env[ironment]:STRING=VALUE - Specifies a single environment variable to be set when shell starts, which allows changing default environment for shell. Multiple occurrences of this switch must be used to specify multiple environment variables.

-noe[cho] - Specifies that echo should be disabled. This may be necessary to ensure that user's answers to remote prompts are not displayed locally. By default echo is "on".

-nop[rofile] - Specifies that the user's profile should not be loaded. By default the server will attempt to load the user profile. If the remote user is not a local administrator on the target system then this option will be required (the default will result in error).

-a[llow]d[elegate] - Specifies that the user's credentials can be used to access a remote share, for example, found on a different machine than the target endpoint.

-comp[ression] - Turn on compression. Older installations on remote machines may not support compression so it is off by default.

-[use]ssl - Use an SSL connection when using a remote endpoint. Specifying this instead of the transport "https:" will use the default WinRM default port.

-? - Help

To terminate the remote command the user can type Ctrl-C or Ctrl-Break, which will be sent to the remote shell. The second Ctrl-C will force termination of winrs.exe.

To manage active remote shells or WinRS configuration, use the WinRM tool. The URI alias to manage active shells is shell/cmd. The URI alias for WinRS configuration is winrm/config/winrs. Example usage can be found in the WinRM tool by typing "WinRM -?".

Examples:

```
winrs -r:https://myserver.com command
winrs -r:myserver.com -usessl command
winrs -r:myserver command
winrs -r:http://127.0.0.1 command
winrs -r:http://169.51.2.101:80 -unencrypted command
winrs -r:https://[::FFFF:129.144.52.38] command
winrs -r:http://[1080:0:0:0:8:800:200C:417A]:80 command
winrs -r:https://myserver.com -t:600 -u:administrator -p:$%fgh7 ipconfig
winrs -r:myserver -env:PATH=^%PATH^%;c:\tools -env:TEMP=d:\temp config.cmd
winrs -r:myserver netdom join myserver /domain:testdomain /userd:johns /passwordd:$%fgh789
winrs -r:myserver -ad -u:administrator -p:$%fgh7 dir \\anotherserver\share
```

QUESTION NO: 17

Your company has a main office and a branch office. You plan to deploy a Readonly Domain Controller (RODC) in the branch office.

You need to plan a strategy to manage the RODC. Your plan must meet the following requirements:

- Allow branch office support technicians to maintain drivers and disks on the RODC

- Prevent branch office support technicians from managing domain user accounts

What should you include in your plan?

A. Configure the RODC for Administrator Role Separation.
B. Configure the RODC to replicate the password for the branch office support technicians.
C. Set NTFS permissions on the Active Directory database to Read & Execute for the branch office support technicians.
D. Set NTFS permissions on the Active Directory database to Deny Full Control for the branch office support technicians.

Answer: A

Explanation:

http://technet.microsoft.com/en-us/library/cc753170%28WS.10%29.aspx

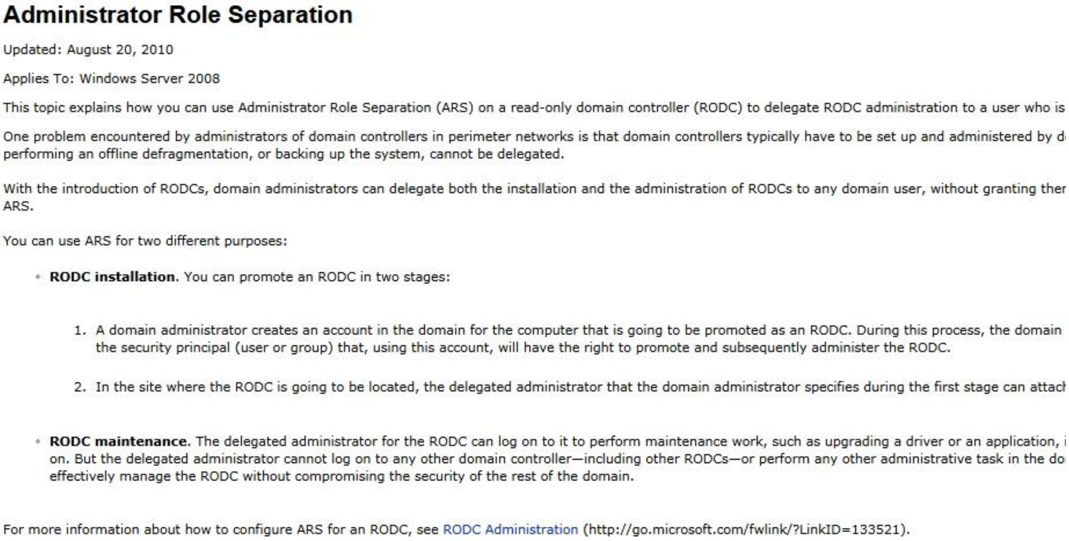

Administrator Role Separation

Updated: August 20, 2010

Applies To: Windows Server 2008

This topic explains how you can use Administrator Role Separation (ARS) on a read-only domain controller (RODC) to delegate RODC administration to a user who is

One problem encountered by administrators of domain controllers in perimeter networks is that domain controllers typically have to be set up and administered by d
performing an offline defragmentation, or backing up the system, cannot be delegated.

With the introduction of RODCs, domain administrators can delegate both the installation and the administration of RODCs to any domain user, without granting ther
ARS.

You can use ARS for two different purposes:

- **RODC installation.** You can promote an RODC in two stages:

  1. A domain administrator creates an account in the domain for the computer that is going to be promoted as an RODC. During this process, the domain
     the security principal (user or group) that, using this account, will have the right to promote and subsequently administer the RODC.

  2. In the site where the RODC is going to be located, the delegated administrator that the domain administrator specifies during the first stage can attacl

- **RODC maintenance.** The delegated administrator for the RODC can log on to it to perform maintenance work, such as upgrading a driver or an application, i
  on. But the delegated administrator cannot log on to any other domain controller—including other RODCs—or perform any other administrative task in the do
  effectively manage the RODC without compromising the security of the rest of the domain.

For more information about how to configure ARS for an RODC, see RODC Administration (http://go.microsoft.com/fwlink/?LinkID=133521).

QUESTION NO: 18

Your network consists of a single Active Directory domain. The network contains five Windows Server 2008 R2 servers that host Web Applications. You need to plan a remote management strategy to manage the Web servers.

Your plan must meet the following requirements:

- Allow Web developers to configure features on the Web sites
- Prevent Web developers from having full administrative rights on the Web servers

What should you include in your plan?

A. Configure request filtering on each Web server.
B. Configure authorization rules for Web developers on each Web server.
C. Configure the security settings in Internet Explorer for all Web developers by using a Group Policy.
D. Add the Web developers to the Account Operators group in the domain.

Answer: B
Explanation:

http://mscerts.programming4.us/windows_server/windows%20server%202008%20%20%20controlling%20access%20to%20web%20services%20%28part%205%29%20-%20managing%20url%20authorization%20rules.aspx

Managing URL Authorization Rules

Authorization is a method by which systems administrators can determine which resources and content are available to specific users Authorization relies on authentication to validate the identity of a user. Once the identity has been proven, authorization rules determine which actions a user or computer can perform IIS provides methods of securing different types of content using URL-based authorization. Because Web content is generally requested using a URL that includes a full path to the content being requested, you can configure authorization settings easily, using IIS Manager

Creating URL Authorization Rules

To enable URL authorization, the UrlAuthorizationModule must be enabled Authorization rules can be configured at the level of the Web server for specific Web sites, for specific Web applications, and for specific files (based on a complete URL path). URL authorization rules use inheritance so that lower-level objects inherit authorization settings from their parent objects (unless they are specifically overridden).

To configure authorization settings, select the appropriate object in the left pane of IIS Manager, and then select Authorization Rules in Features View. Figure 6 shows an example of multiple rules configured for a Web site.

Figure 6. Viewing authorization rules for a Web site

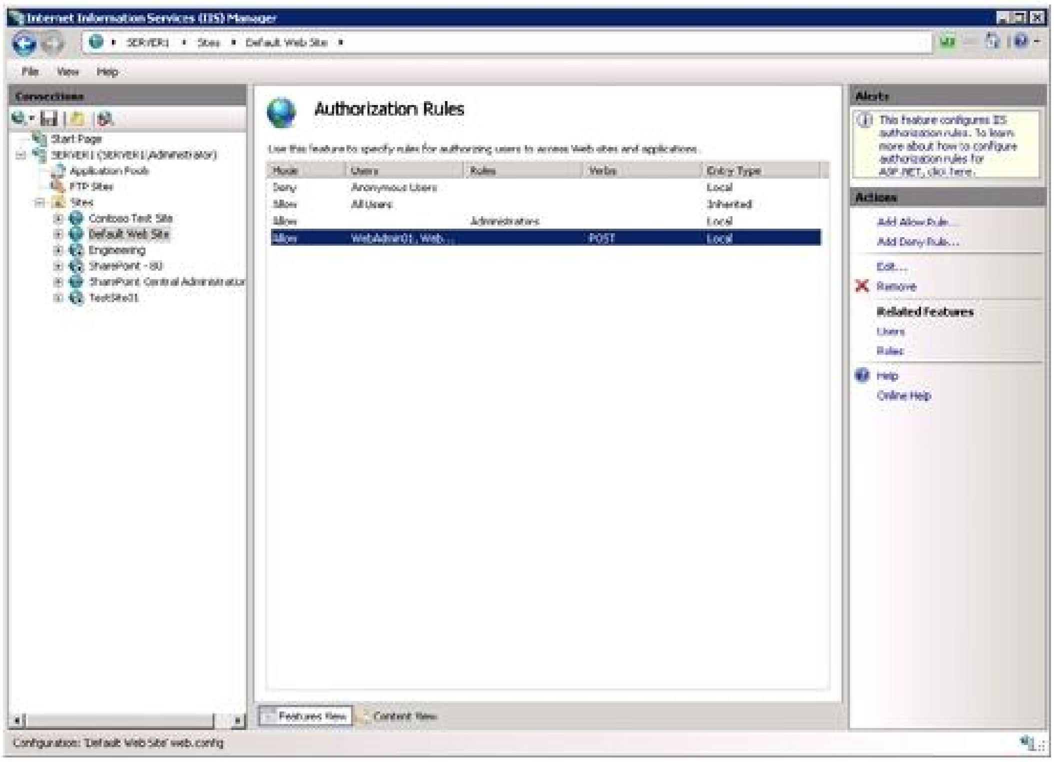

Authorization Rules

Use this feature to specify rules for authorizing users to access Web sites and applications.

Mode	Users	Roles	Verbs	Entry Type
Deny	Anonymous Users			Local
Allow	All Users			Inherited
Allow		Administrators		Local
Allow	WebAdmin01, Web...		POST	Local

This feature configures IIS authorization rules. To learn more about how to configure authorization rules for ASP.NET, click here.

Configuration: 'Default Web Site' web.config

There are two types of rules: Allow and Deny. You can create new rules by using the Add Allow Rule and Add Deny Rule commands in the Actions pane The available options for both types of rules are the same.

(See Figure 7) When creating a new rule, the main setting is to determine to which users the rule applies. The options are:

- All Users
- All Anonymous Users
- Specific Roles Or User Groups
- Specific Users

Figure 7. Creating a new Allow Rule for a Web application

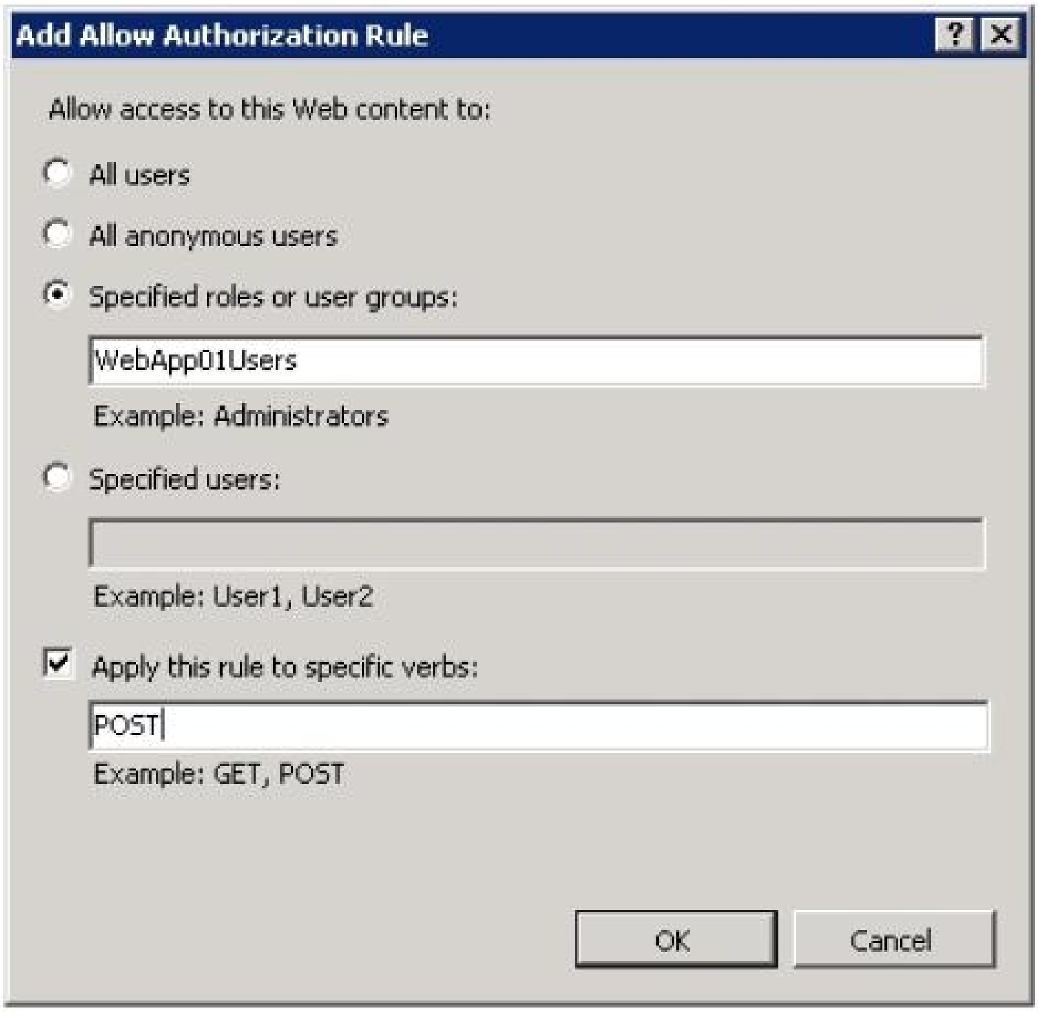

Add Allow Authorization Rule

Allow access to this Web content to:

○ All users

○ All anonymous users

● Specified roles or user groups:

WebApp01Users

Example: Administrators

○ Specified users:

Example: User1, User2

☑ Apply this rule to specific verbs:

POST

Example: GET, POST

When you choose to specify users or groups to which the rule applies, you can type the appropriate names in a command-separated list. The specific users and groups are defined using NET role providers. This is a standard feature that is available to ASP NET Web developers. Developers can create their own roles and user accounts and can define permissions within their applications. Generally, information about users and roles is stored in a relational database or relies on a directory service such as Active Directory.

In addition to user and role selections, you can further configure an authorization rule based on specific HTTP verbs. For example, if you want to apply a rule only for POST commands (which are typically used to send information from a Web browser to a Web server), add only the POST verb to the rule

Managing Rule Inheritance

As mentioned earlier in this section, authorization rules are inherited automatically by lower-level objects This is useful when your Web site and Web content is organized hierarchically based on intended users or groups The Entry Type column shows whether a rule has been inherited from a higher level or whether it has been defined locally IIS Manager automatically will prevent you from creating duplicate rules. You can remove rules at any level, including both Inherited and Local entry types

QUESTION NO: 19

Your network consists of a single Active Directory domain. The functional level of the domain is Windows Server 2008 R2. The domain contains 200 Windows Server 2008 R2 servers.

You need to plan a monitoring solution that meets the following requirements:

- Sends a notification by email to the administrator if an Application error occurs on any of the servers
- Uses the minimum amount of administrative effort

What should you include in your plan?

A. On one server, create event subscriptions for each server. On the server, attach tasks to the Application error events.
B. On one server, create an Event Trace Sessions Data Collector Set. On all servers, create a System Performance Data Collector Set.
C. On all servers, create event subscriptions for one server. On all servers, attach a task for the Application error events.
D. On all servers, create a System Performance Data Collector Set. On one server, configure the report settings for the new Data Collector set.

Answer: A
Explanation:
http://technet.microsoft.com/en-us/library/cc749183.aspx
http://technet.microsoft.com/en-us/library/cc748890.aspx
http://technet.microsoft.com/en-us/library/cc722010.aspx

Event Subscriptions
Applies To: Windows 7, Windows Server 2008 R2, Windows Vista
Event Viewer enables you to view events on a single remote computer. However, troubleshooting an issue might require you to examine a set of events stored in multiple logs on multiple computers.

Windows Vista includes the ability to collect copies of events from multiple remote computers and store them locally. To specify which events to collect, you create an event subscription. Among other details, the subscription specifies exactly which events will be collected and in which log they will be stored locally. Once a subscription is active and events are being collected, you can view and manipulate these forwarded events as you would any other locally stored events.

Using the event collecting feature requires that you configure both the forwarding and the collecting computers. The functionality depends on the Windows Remote Management (WinRM) service and the Windows Event Collector (Wecsvc) service. Both of these services must be running on computers participating in the forwarding and collecting process. To learn about the steps required to configure event collecting and forwarding computers, see Configure Computers to Forward and Collect Events.

Additional Considerations

• You can subscribe to receive events from an existing subscription on a remote computer.

Configure Computers to Forward and Collect Events

Applies To: Windows 7, Windows Server 2008 R2, Windows Vista

Before you can create a subscription to collect events on a computer, you must configure both the collecting computer collected (collector) and each computer from which events will be collected (source). Updated information about event subscriptions may be available online at Event Subscriptions.

To configure computers in a domain to forward and collect events

1. Log on to all collector and source computers. It is a best practice to use a domain account with administrative privileges.
2. On each source computer, type the following at an elevated command prompt:

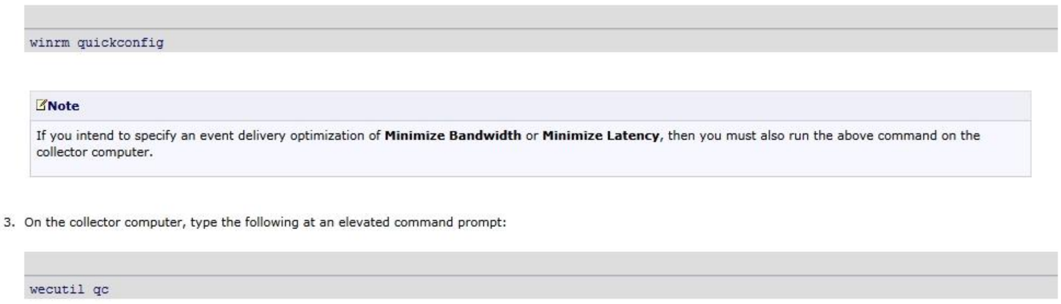

```
winrm quickconfig
```

Note

If you intend to specify an event delivery optimization of **Minimize Bandwidth** or **Minimize Latency**, then you must also run the above command on the collector computer.

3. On the collector computer, type the following at an elevated command prompt:

```
wecutil qc
```

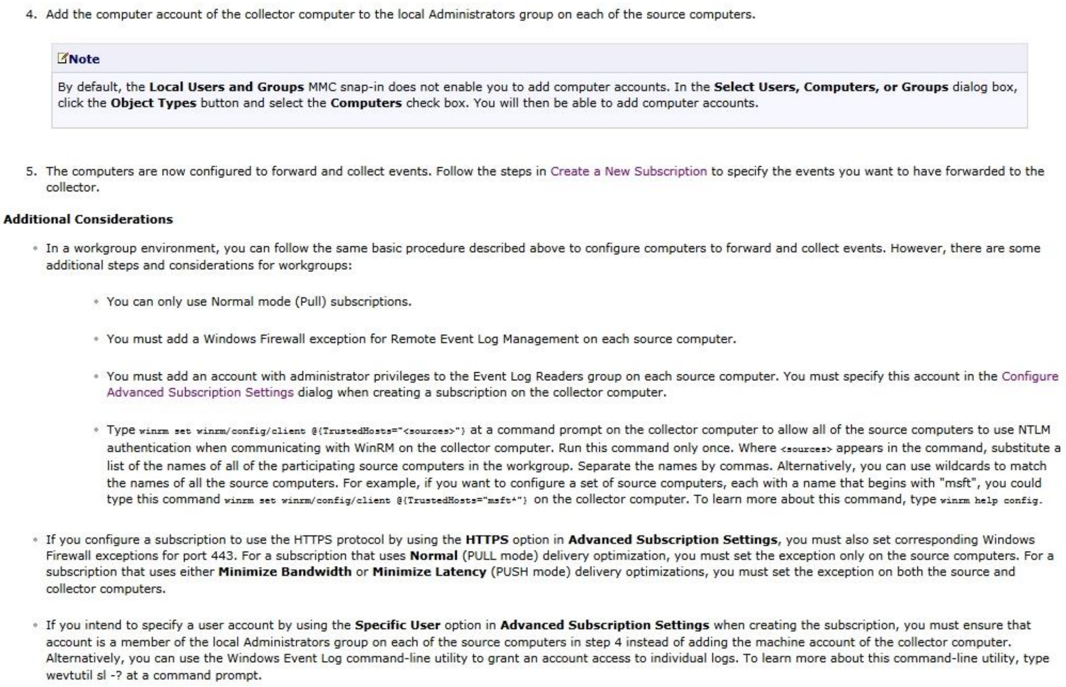

4. Add the computer account of the collector computer to the local Administrators group on each of the source computers.

> **Note**
>
> By default, the **Local Users and Groups** MMC snap-in does not enable you to add computer accounts. In the **Select Users, Computers, or Groups** dialog box, click the **Object Types** button and select the **Computers** check box. You will then be able to add computer accounts.

5. The computers are now configured to forward and collect events. Follow the steps in Create a New Subscription to specify the events you want to have forwarded to the collector.

Additional Considerations

- In a workgroup environment, you can follow the same basic procedure described above to configure computers to forward and collect events. However, there are some additional steps and considerations for workgroups:
    - You can only use Normal mode (Pull) subscriptions.
    - You must add a Windows Firewall exception for Remote Event Log Management on each source computer.
    - You must add an account with administrator privileges to the Event Log Readers group on each source computer. You must specify this account in the Configure Advanced Subscription Settings dialog when creating a subscription on the collector computer.
    - Type `winrm set winrm/config/client @{TrustedHosts="<sources>"}` at a command prompt on the collector computer to allow all of the source computers to use NTLM authentication when communicating with WinRM on the collector computer. Run this command only once. Where `<sources>` appears in the command, substitute a list of the names of all of the participating source computers in the workgroup. Separate the names by commas. Alternatively, you can use wildcards to match the names of all the source computers. For example, if you want to configure a set of source computers, each with a name that begins with "msft", you could type this command `winrm set winrm/config/client @{TrustedHosts="msft*"}` on the collector computer. To learn more about this command, type `winrm help config`.
- If you configure a subscription to use the HTTPS protocol by using the **HTTPS** option in **Advanced Subscription Settings**, you must also set corresponding Windows Firewall exceptions for port 443. For a subscription that uses **Normal** (PULL mode) delivery optimization, you must set the exception only on the source computers. For a subscription that uses either **Minimize Bandwidth** or **Minimize Latency** (PUSH mode) delivery optimizations, you must set the exception on both the source and collector computers.
- If you intend to specify a user account by using the **Specific User** option in **Advanced Subscription Settings** when creating the subscription, you must ensure that account is a member of the local Administrators group on each of the source computers in step 4 instead of adding the machine account of the collector computer. Alternatively, you can use the Windows Event Log command-line utility to grant an account access to individual logs. To learn more about this command-line utility, type wevtutil sl -? at a command prompt.

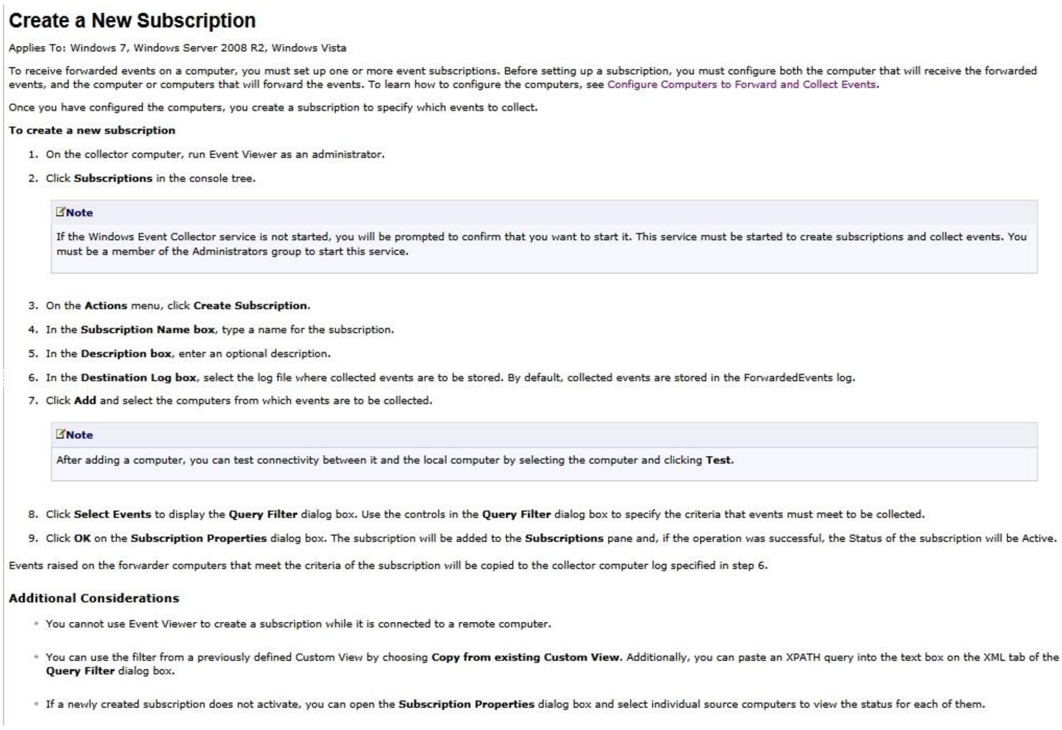

Create a New Subscription

Applies To: Windows 7, Windows Server 2008 R2, Windows Vista

To receive forwarded events on a computer, you must set up one or more event subscriptions. Before setting up a subscription, you must configure both the computer that will receive the forwarded events, and the computer or computers that will forward the events. To learn how to configure the computers, see Configure Computers to Forward and Collect Events.

Once you have configured the computers, you create a subscription to specify which events to collect.

To create a new subscription

1. On the collector computer, run Event Viewer as an administrator.
2. Click **Subscriptions** in the console tree.

> **Note**
>
> If the Windows Event Collector service is not started, you will be prompted to confirm that you want to start it. This service must be started to create subscriptions and collect events. You must be a member of the Administrators group to start this service.

3. On the **Actions** menu, click **Create Subscription.**
4. In the **Subscription Name box**, type a name for the subscription.
5. In the **Description box**, enter an optional description.
6. In the **Destination Log box**, select the log file where collected events are to be stored. By default, collected events are stored in the ForwardedEvents log.
7. Click **Add** and select the computers from which events are to be collected.

> **Note**
>
> After adding a computer, you can test connectivity between it and the local computer by selecting the computer and clicking **Test**.

8. Click **Select Events** to display the **Query Filter** dialog box. Use the controls in the **Query Filter** dialog box to specify the criteria that events must meet to be collected.
9. Click **OK** on the **Subscription Properties** dialog box. The subscription will be added to the **Subscriptions** pane and, if the operation was successful, the Status of the subscription will be Active.

Events raised on the forwarder computers that meet the criteria of the subscription will be copied to the collector computer log specified in step 6.

Additional Considerations

- You cannot use Event Viewer to create a subscription while it is connected to a remote computer.
- You can use the filter from a previously defined Custom View by choosing **Copy from existing Custom View.** Additionally, you can paste an XPATH query into the text box on the XML tab of the **Query Filter** dialog box.
- If a newly created subscription does not activate, you can open the **Subscription Properties** dialog box and select individual source computers to view the status for each of them.

QUESTION NO: 20

Your network consists of a single Active Directory domain. The network includes a branch office named Branch1. Branch1 contains 50 member servers that run Windows Server 2008 R2. An organizational unit (OU) named Branch1Servers contains the computer objects for the servers in Branch1. A global group named Branch1admins contains the user accounts for the administrators. Administrators maintain all member servers in Branch1. You need to recommend a solution that allows the members of Branch1admins group to perform the following tasks on the Branch1 member servers.

- Stop and start services
- Change registry settings

What should you recommend?

A. Add the Branch1admins group to the Power Users local group on each server in Branch1.
B. Add the Branch1admins group to the Administrators local group on each server in Branch1.
C. Assign the Branch1admins group change permissions to the Branch1Servers OU and to all child objects.
D. Assign the Branch1admins group Full Control permissions on the Branch1Servers OU and to all child objects.

Answer: B

Explanation:

Local admins have these rights.

Power Users do not

By default, members of the power users group have no more user rights or permissions than a standard user account. The Power Users group in previous versions of Windows was designed to give users specific administrator rights and permissions to perform common system tasks. In this version of Windows, standard user accounts inherently have the ability to perform most common configuration tasks, such as changing time zones. For legacy applications that require the same Power User rights and permissions that were present in previous versions of Windows, administrators can apply a security template that enables the Power Users group to assume the same rights and permissions that were present in previous versions of Windows.

QUESTION NO: 21

Your network consists of a single Active Directory domain. The network includes a branch office named Branch1. Branch1 contains a Read only Domain Controller (RODC) named Server1. A global group named Branch1admins contains the user accounts for administrators. Administrators manage the client computers and servers in Branch1.

You need to recommend a solution for delegating control of Server1.

Your solution must meet the following requirements:

- Allow the members of the Branch1admins group to administer Server1 including, change device drivers and install operating system updates by using Windows Update.
- Provide the Branch1admins group rights on Server1 only.

- Prevent Branch1admins group from modifying Active Directory objects.

What should you recommend?

A. Add the Branch1admins global group to the Server Operators builtin local group.
B. Add the members of the Branch1admins global group to the Administrators builtin local group of Server1.
C. Grant Full Control permission on the Server1 computer object in the domain to the Branch1admins group
D. Move the Server1 computer object to a new organizational unit (OU) named Branch1servers. Grant Full Control permission on the Branch1servers OU to the Branch1admins group.

Answer: B

Explanation:

http://technet.microsoft.com/en-us/library/cc753223%28WS.10%29.aspx

Administrator role separation

Administrator role separation specifies that any domain user or security group can be delegated to be the local administrator of an RODC without granting that user or group any rights for the domain or other domain controllers. Accordingly, a delegated administrator can log on to an RODC to perform maintenance work, such as upgrading a driver, on the server. But the delegated administrator is not able to log on to any other domain controller or perform any other administrative task in the domain. In this way, a security group that comprises branch users, rather than members of the Domain Admins group, can be delegated the ability to effectively manage the RODC in the branch office, without compromising the security of the rest of the domain.

QUESTION NO: 22

Your network consists of a single Active Directory forest. The forest functional level is Windows Server 2008 R2. The forest contains two domains named contoso.com and na.contoso.com. Contoso.com contains a user named User1. Na.contoso.com contains an organizational unit (OU) named Security.

You need to give User1 administrative rights so that he can manage Group Policies for the Security OU.

You want to achieve this goal while meeting the following requirements:

- User1 must be able to create and configure Group Policies in na.contoso.com.
- User1 must be able to link Group Policies to the Security OU.
- User1 must be granted the least administrative rights necessary to achieve the goal.

What should you do?

A. Add User1 to the Administrators group for na.contoso.com.
B. Add User1 to the Group Policy Creator Owners group in contoso.com. Modify the permissions on the Security OU.
C. Run the Delegation of Control Wizard on the Security OU. In the Group Policy Management Console, modify the permissions of the Group Policy Objects container in the na.contoso.com domain.
D. Run the Delegation of Control Wizard on na.contoso.com. In the Group Policy Management Console, modify the permissions of the Group Policy Objects container in the contoso.com domain.

Answer: C

Explanation:

http://technet.microsoft.com/en-us/library/dd145442.aspx

http://technet.microsoft.com/en-us/library/dd145338.aspx

http://technet.microsoft.com/en-us/library/dd145594.aspx

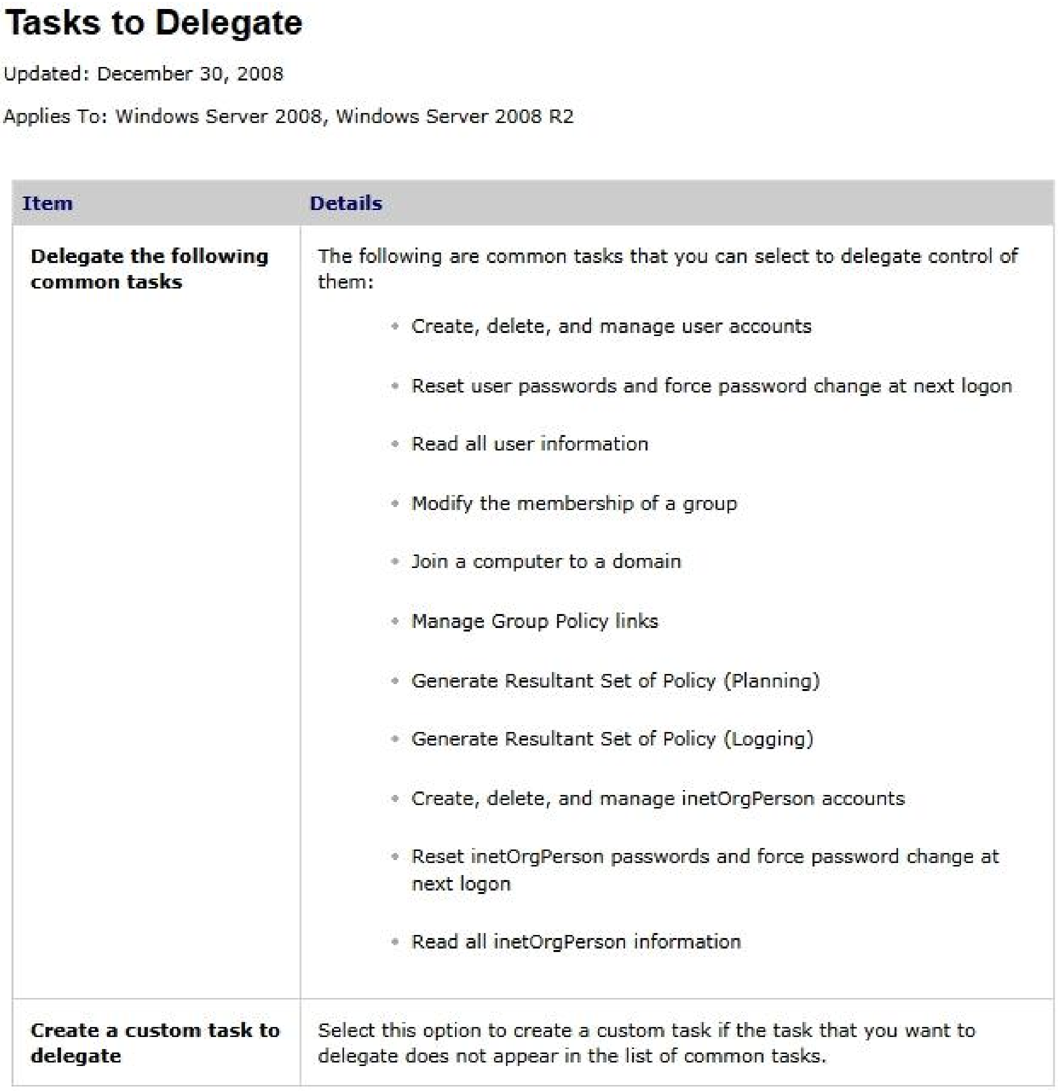

Tasks to Delegate

Updated: December 30, 2008

Applies To: Windows Server 2008, Windows Server 2008 R2

Item	Details
Delegate the following common tasks	The following are common tasks that you can select to delegate control of them: • Create, delete, and manage user accounts • Reset user passwords and force password change at next logon • Read all user information • Modify the membership of a group • Join a computer to a domain • Manage Group Policy links • Generate Resultant Set of Policy (Planning) • Generate Resultant Set of Policy (Logging) • Create, delete, and manage inetOrgPerson accounts • Reset inetOrgPerson passwords and force password change at next logon • Read all inetOrgPerson information
Create a custom task to delegate	Select this option to create a custom task if the task that you want to delegate does not appear in the list of common tasks.

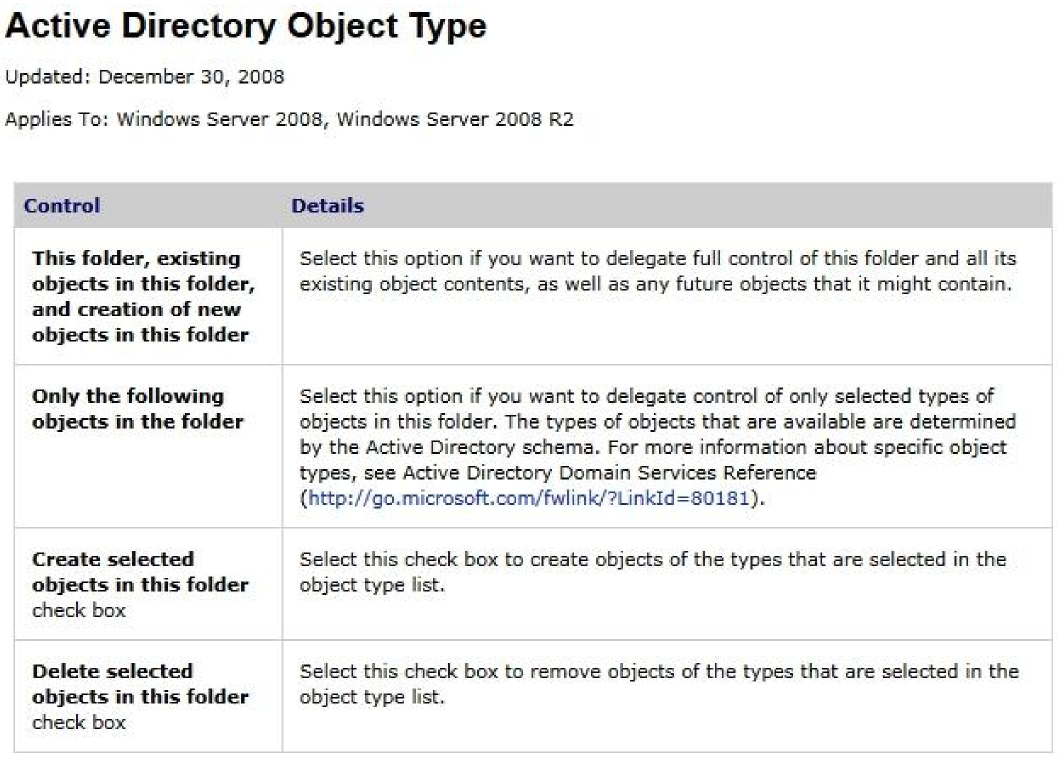

Active Directory Object Type

Updated: December 30, 2008

Applies To: Windows Server 2008, Windows Server 2008 R2

Control	Details
This folder, existing objects in this folder, and creation of new objects in this folder	Select this option if you want to delegate full control of this folder and all its existing object contents, as well as any future objects that it might contain.
Only the following objects in the folder	Select this option if you want to delegate control of only selected types of objects in this folder. The types of objects that are available are determined by the Active Directory schema. For more information about specific object types, see Active Directory Domain Services Reference (http://go.microsoft.com/fwlink/?LinkId=80181).
Create selected objects in this folder check box	Select this check box to create objects of the types that are selected in the object type list.
Delete selected objects in this folder check box	Select this check box to remove objects of the types that are selected in the object type list.

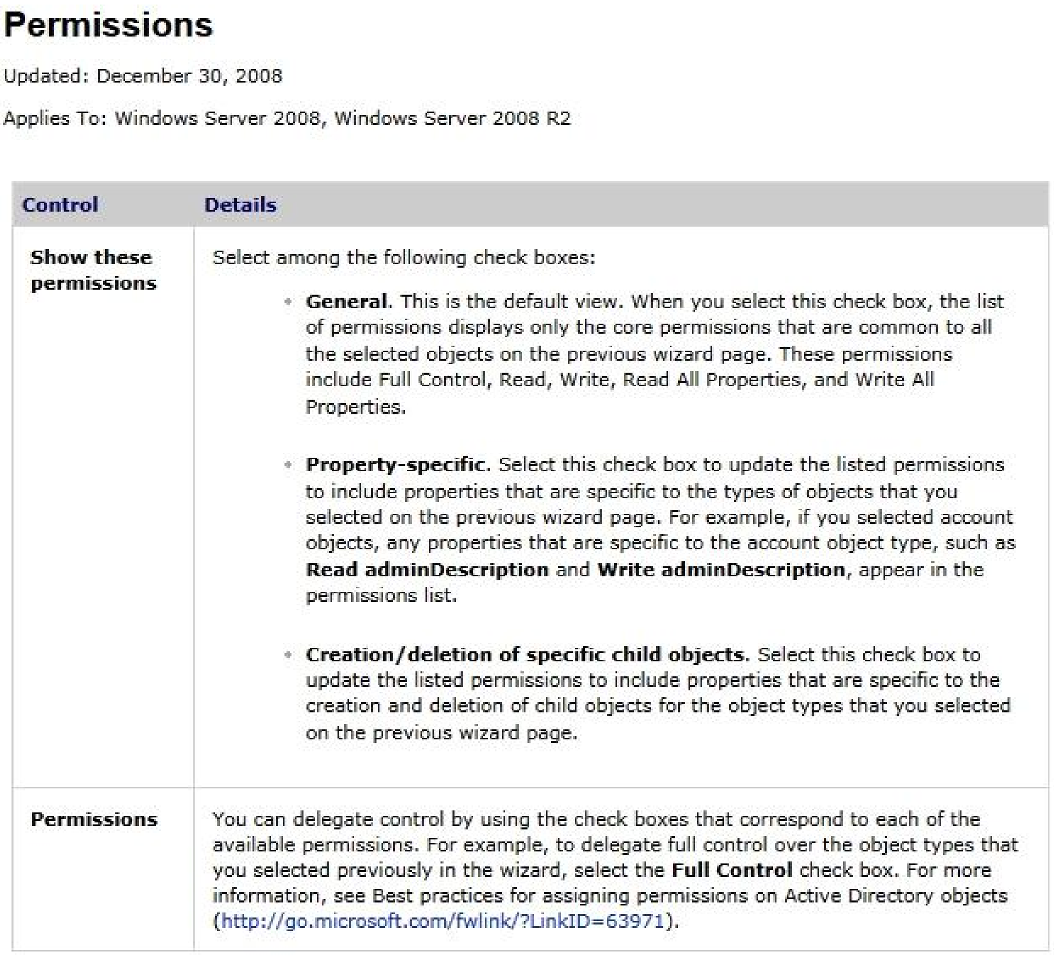

Permissions

Updated: December 30, 2008

Applies To: Windows Server 2008, Windows Server 2008 R2

Control	Details
Show these permissions	Select among the following check boxes: • **General**. This is the default view. When you select this check box, the list of permissions displays only the core permissions that are common to all the selected objects on the previous wizard page. These permissions include Full Control, Read, Write, Read All Properties, and Write All Properties. • **Property-specific**. Select this check box to update the listed permissions to include properties that are specific to the types of objects that you selected on the previous wizard page. For example, if you selected account objects, any properties that are specific to the account object type, such as **Read adminDescription** and **Write adminDescription**, appear in the permissions list. • **Creation/deletion of specific child objects.** Select this check box to update the listed permissions to include properties that are specific to the creation and deletion of child objects for the object types that you selected on the previous wizard page.
Permissions	You can delegate control by using the check boxes that correspond to each of the available permissions. For example, to delegate full control over the object types that you selected previously in the wizard, select the **Full Control** check box. For more information, see Best practices for assigning permissions on Active Directory objects (http://go.microsoft.com/fwlink/?LinkID=63971).

QUESTION NO: 23

Your network contains several branch offices. All servers run Windows Server 2008 R2. Each branch office contains a domain controller and a file server.

The DHCP Server server role is installed on the branch office domain controllers. Each office has a branch office administrator.

You need to delegate the administration of DHCP to meet the following requirements:

- Allow branch office administrators to manage DHCP scopes for their own office
- Prevent the branch office administrators from managing DHCP scopes in other offices
- Minimize administrative effort

What should you do?

A. In the Active Directory domain, add the branch office administrators to the Server Operators builtin local group.
B. In the Active Directory domain, add the branch office administrators to the Network Configuration Operators builtin local group.
C. In each branch office, migrate the DHCP Server server role to the file server. On each file server, add the branch office administrator to the DHCP Administrators local group.
D. In each branch office, migrate the DHCP Server server role to the file server. In the Active Directory domain, add the branch office administrators to the DHCP Administrators domain local group.

Answer: C

Explanation:

http://technet.microsoft.com/en-us/library/dd379494%28WS.10%29.aspx

http://technet.microsoft.com/en-us/library/dd379483%28WS.10%29.aspx

http://technet.microsoft.com/en-us/library/dd379535%28WS.10%29.aspx

http://technet.microsoft.com/en-us/library/cc737716%28WS.10%29.aspx

DHCP Server Migration: Appendix A

Updated: April 29, 2009

Applies To: Windows Server 2008 R2

Migration tools

Migration tools are provided in Windows Server 2008 R2. The tools for earlier versions of the Windows operating system are also available in Windows Server 2008 R2.

Follow these steps to access the tools on the destination server:

1. Open Server Manager.
2. Click **Action**, and then select **Add Features**. The Add Features Wizard opens.
3. On the **Select Features** page, from the **Features** list, select **Windows Server Migration Tools**, and then click **Next**.
4. Complete the steps in the wizard, and then click **Close**.

The previous steps do not work for Server Core installations. To install the migration tools on a Server Core installation, see Windows Server Migration Tools Installation, Access, and Removal (http://go.microsoft.com/fwlink/?LinkID=134763).

Installing and using Windows PowerShell with migration cmdlets
To access, download, and install migration tools (the migration toolkit), any role-specific tools, and Windows PowerShell, see Windows Server Migration Tools Installation, Access, and Removal (http://go.microsoft.com/fwlink/?LinkID=134763).

Known issues
If the DHCP installation on the source server has a database path that varies from the default, before you perform the import, provide the destination server with a volume with the same drive letter on which the DHCP server database exists on the source server. For example, if the DHCP server database on the source server is located on D:\, then the destination server should have a volume with the driver letter D.

If you cannot match the volume on the destination server that has the same driver letter as that shown for the source DHCP server database, then the DHCP database path on the source server must be changed back to the default path (%systemroot%\system32\dhcp) before you start the migration.

DHCP Server migration process
As shown in the following illustration, the pre-migration process involves the manual collection of data, followed by procedures on the destination and source servers. The migration process includes source and destination server procedures that use the **Export** and **Import** cmdlets to automatically collect, store, and then migrate server role settings. Post-migration procedures include verifying that the destination server successfully replaced the source server and then retiring or repurposing the source server. If the verification procedure indicates that the migration failed, troubleshooting begins. If troubleshooting fails, rollback instructions are provided to return to the use of the original source server.

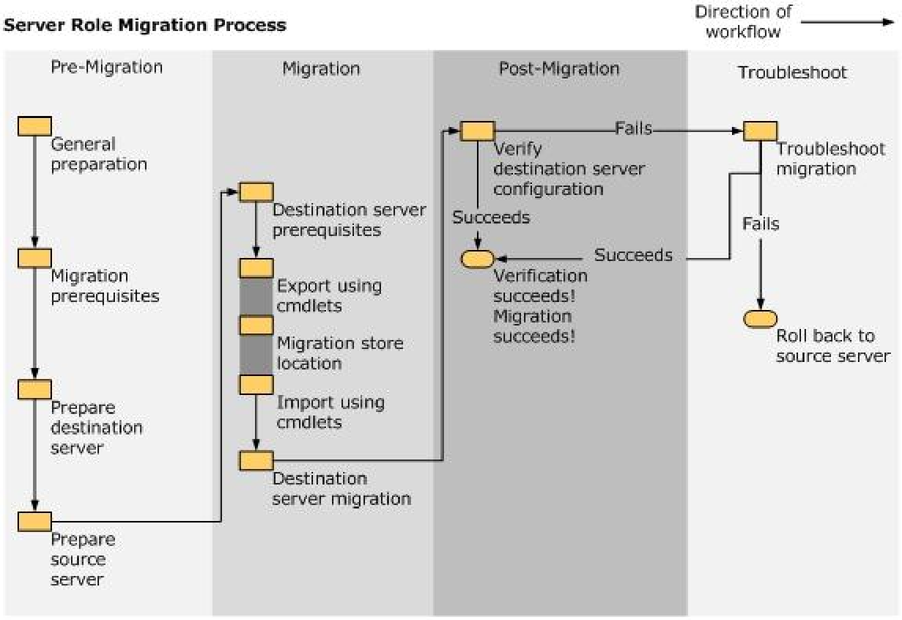

DHCP Administrators

Members of the DHCP Administrators group can view and modify any data at the DHCP server. DHCP Administrators can create and delete scopes, add reservations, change option values, create superscopes, or perform any other activity needed to administer the DHCP server, including export or import of the DHCP server configuration and database. DHCP Administrators perform these tasks using the Netsh commands for DHCP or the DHCP console. For more information, see DHCP tools.

Members of the DHCP Administrators group do not have unlimited administrative rights. For example, if a DHCP server is also configured as a DNS server, a member of the DHCP Administrators group can view and modify the DHCP configuration but cannot modify DNS server configuration on the same computer.

Because members of the DHCP Administrators group have rights on the local computer only, DHCP Administrators cannot authorize or unauthorize DHCP servers in Active Directory. Only members of the Domain Admins group can perform this task. If you want to authorize or unauthorize a DHCP server in a child domain, you must have enterprise administrator credentials for the parent domain. For more information about authorizing DHCP servers in Active Directory, see Authorizing DHCP servers and Authorize a DHCP server in Active Directory.

Using groups to administer DHCP servers in a domain

When you add a user or group to a DHCP Users or DHCP Administrators group on a DHCP server, the rights of the DHCP group member do not apply to all of the DHCP servers in the domain. The rights apply only to the DHCP service on the local computer.

QUESTION NO: 24

Your company has a single Active Directory domain. You have 30 database servers that run Windows Server 2008 R2.

The computer accounts for the database servers are stored in an organizational unit (OU) named Data. The user accounts for the database administrators are stored in an OU named Admin. The database administrators are members of a global group named D_Admins.

You must allow the database administrators to perform administrative tasks on the database servers. You must prevent the database administrators from performing administrative tasks on other servers.

What should you do?

A. Deploy a Group Policy to the Data OU.
B. Deploy a Group Policy to the Admin OU.
C. Add D_Admins to the Domain Admins global group.
D. Add D_Admins to the Server Operators built-in local group.

Answer: A
Explanation:
http://technet.microsoft.com/en-us/library/cc754948%28WS.10%29.aspx
Group Policy Planning and Deployment Guide
You can use Windows Server 2008 Group Policy to manage configurations for groups of computers and users, including options for registry-based policy settings, security settings, software deployment, scripts, folder redirection, and preferences. Group Policy preferences, new in Windows Server 2008, are more than 20 Group Policy extensions that expand the range of configurable policy settings within a Group Policy object (GPO). In contrast to Group Policy settings, preferences are not enforced. Users can change preferences after initial deployment. For information about Group Policy Preferences, see Group Policy Preferences Overview.
Using Group Policy, you can significantly reduce an organization's total cost of ownership. Various factors, such as the large number of policy settings available, the interaction between multiple policies, and inheritance options, can make Group Policy design complex. By carefully planning, designing, testing, and deploying a solution based on your organization's business requirements, you can provide the standardized functionality, security, and management control that your organization needs.

Overview of Group Policy
Group Policy enables Active Directory–based change and configuration management of user and computer settings on computers running Windows Server 2008, Windows Vista, Windows Server 2003, and Windows XP. In addition to using Group Policy to define configurations for groups of users and computers, you can also use Group Policy to help manage server computers, by configuring many server-specific operational and security settings.
By using a structure in which OUs contain homogeneous objects, such as either user or computer objects but not both, you can easily disable those sections of a GPO that do not apply to a particular type of object. This approach to OU design, illustrated in Figure 1, reduces complexity and improves the speed at which Group Policy is applied. Keep in mind that GPOs linked to the higher layers of the OU structure are inherited by default, which reduces the need to duplicate GPOs or to link a GPO to multiple containers.
When designing your Active Directory structure, the most important considerations are ease of administration and delegation.

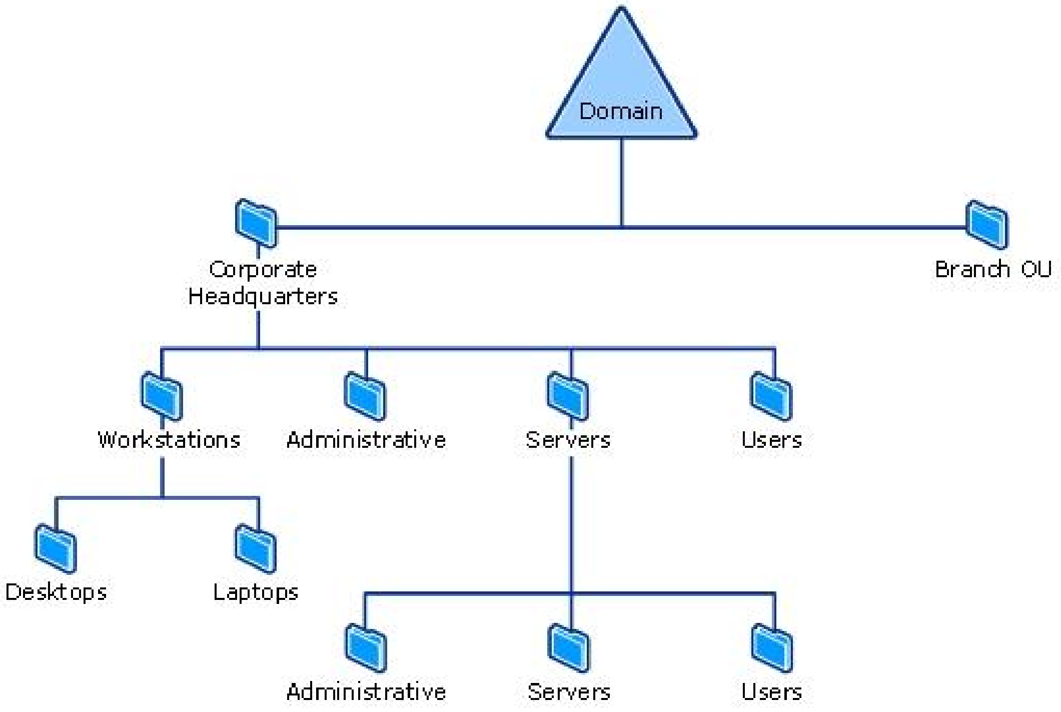

QUESTION NO: 25

Your network consists of a single Active Directory forest that contains a root domain and two child domains.

All servers run Windows Server 2008 R2. A corporate policy has the following requirements:

- All local guest accounts must be renamed and disabled.
- All local administrator accounts must be renamed.
- You need to recommend a solution that meets the requirements of the corporate policy.

What should you recommend?

A. Implement a Group Policy object (GPO) for each domain.
B. Implement a Group Policy object (GPO) for the root domain.
C. Deploy Network Policy and Access Services (NPAS) on all domain controllers in each domain
D. Deploy Active Directory Rights Management Services (AD RMS) on the root domain controllers.

Answer: A

Explanation:

http://www.windowsecurity.com/articles/protecting-administrator-account.html

http://www.pctips3000.com/enable-or-disable-group-policy-object-in-windows-server-2008/
http://blogs.technet.com/b/chenley/archive/2006/07/13/441642.aspx

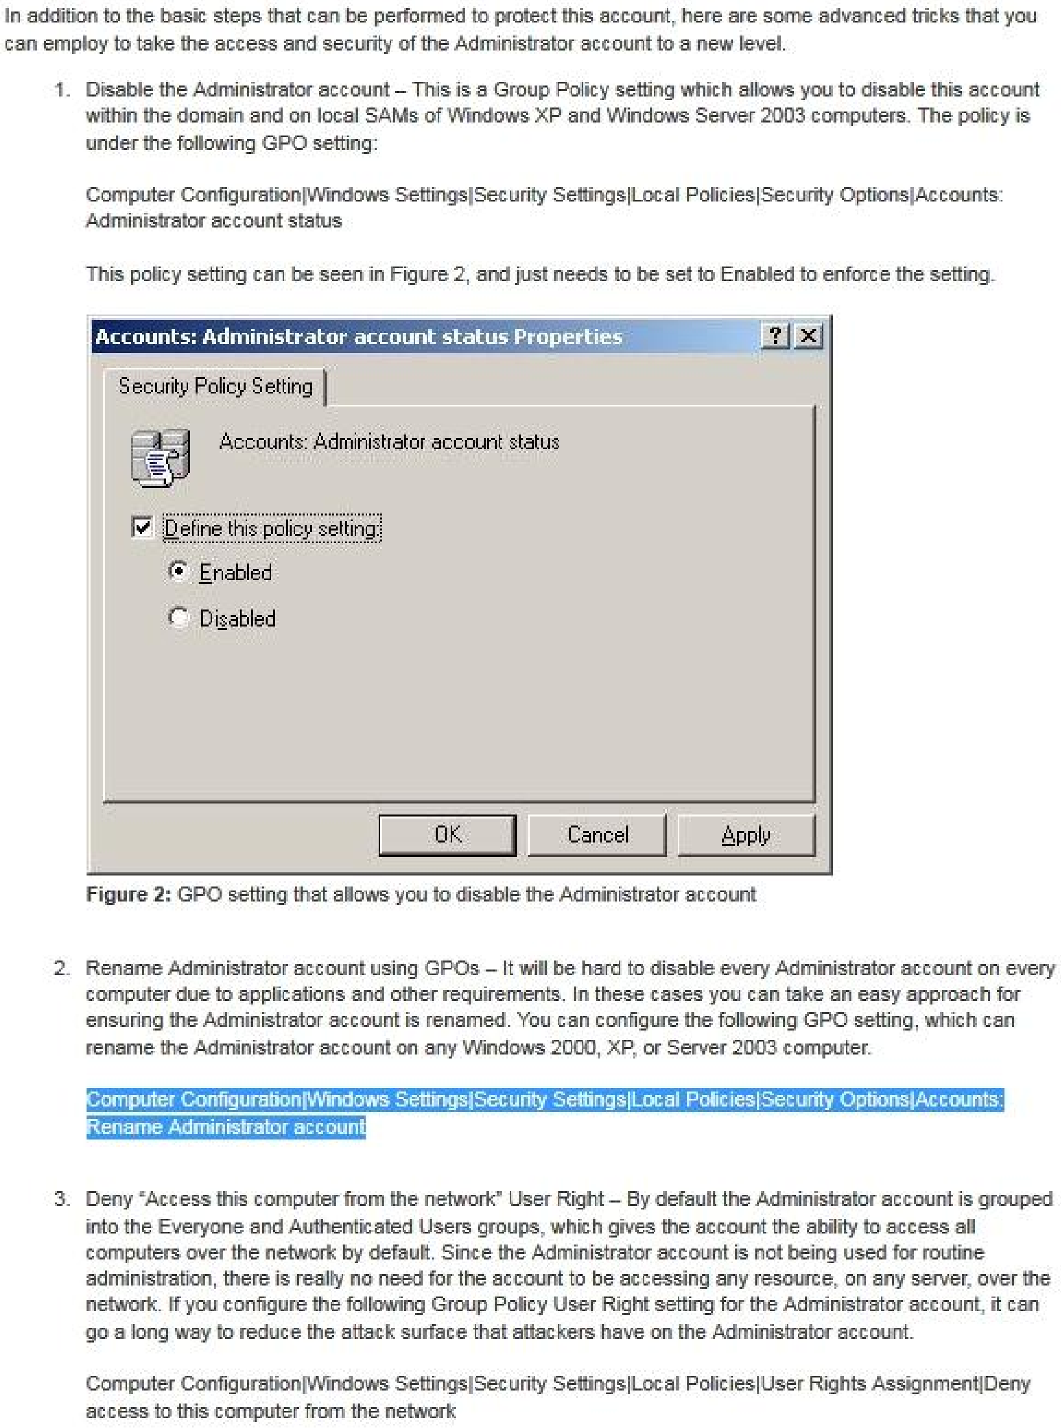

In addition to the basic steps that can be performed to protect this account, here are some advanced tricks that you can employ to take the access and security of the Administrator account to a new level.

1. Disable the Administrator account – This is a Group Policy setting which allows you to disable this account within the domain and on local SAMs of Windows XP and Windows Server 2003 computers. The policy is under the following GPO setting:

   Computer Configuration|Windows Settings|Security Settings|Local Policies|Security Options|Accounts: Administrator account status

   This policy setting can be seen in Figure 2, and just needs to be set to Enabled to enforce the setting.

   Accounts: Administrator account status Properties

   Security Policy Setting

   Accounts: Administrator account status

   ☑ Define this policy setting:

   ● Enabled

   ○ Disabled

   Figure 2: GPO setting that allows you to disable the Administrator account

2. Rename Administrator account using GPOs – It will be hard to disable every Administrator account on every computer due to applications and other requirements. In these cases you can take an easy approach for ensuring the Administrator account is renamed. You can configure the following GPO setting, which can rename the Administrator account on any Windows 2000, XP, or Server 2003 computer.

   Computer Configuration|Windows Settings|Security Settings|Local Policies|Security Options|Accounts: Rename Administrator account

3. Deny "Access this computer from the network" User Right – By default the Administrator account is grouped into the Everyone and Authenticated Users groups, which gives the account the ability to access all computers over the network by default. Since the Administrator account is not being used for routine administration, there is really no need for the account to be accessing any resource, on any server, over the network. If you configure the following Group Policy User Right setting for the Administrator account, it can go a long way to reduce the attack surface that attackers have on the Administrator account.

   Computer Configuration|Windows Settings|Security Settings|Local Policies|User Rights Assignment|Deny access to this computer from the network

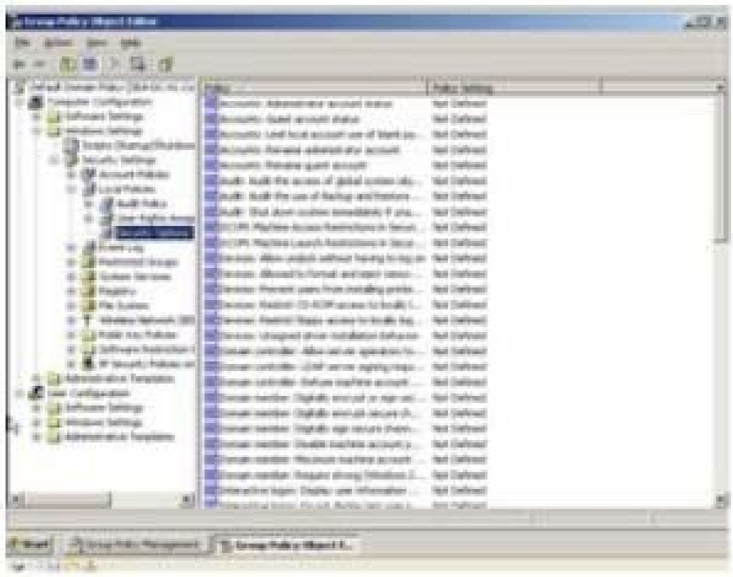

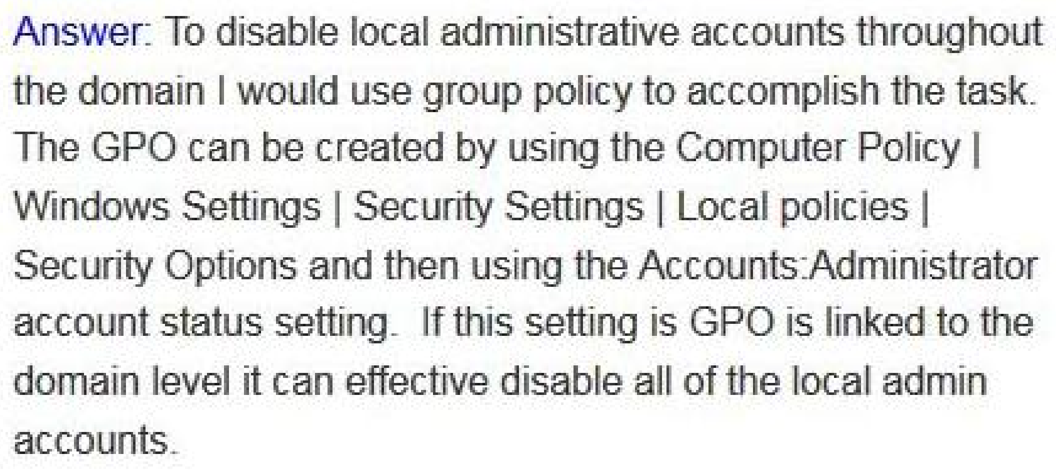

Answer: To disable local administrative accounts throughout the domain I would use group policy to accomplish the task. The GPO can be created by using the Computer Policy | Windows Settings | Security Settings | Local policies | Security Options and then using the Accounts:Administrator account status setting. If this setting is GPO is linked to the domain level it can effective disable all of the local admin accounts.

QUESTION NO: 26

Your network consists of a single Active Directory domain. The functional level of the domain is Windows Server 2008 R2.

All domain controllers run Windows Server 2008 R2. A corporate policy requires that the users from the research department have higher levels of account and password security than other users in the domain.

You need to recommend a solution that meets the requirements of the corporate policy. Your solution must minimize hardware and software costs.

What should you recommend?

A. Create a new Active Directory site. Deploy a Group Policy object (GPO) to the site.
B. Create a new Password Settings Object (PSO) for the research department's users.
C. Create a new organizational unit (OU) named Research in the existing domain. Deploy a Group Policy object (GPO) to the Research OU.
D. Create a new domain in the forest. Add the research department's user accounts to the new domain. Configure a new security policy in the new domain.

Answer: B

Explanation:

http://technet.microsoft.com/en-us/library/cc770842%28WS.10%29.aspx

http://technet.microsoft.com/en-us/library/cc754461%28WS.10%29.aspx

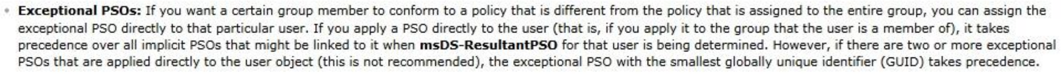

- **Exceptional PSOs:** If you want a certain group member to conform to a policy that is different from the policy that is assigned to the entire group, you can assign the exceptional PSO directly to that particular user. If you apply a PSO directly to the user (that is, if you apply it to the group that the user is a member of), it takes precedence over all implicit PSOs that might be linked to it when **msDS-ResultantPSO** for that user is being determined. However, if there are two or more exceptional PSOs that are applied directly to the user object (this is not recommended), the exceptional PSO with the smallest globally unique identifier (GUID) takes precedence.

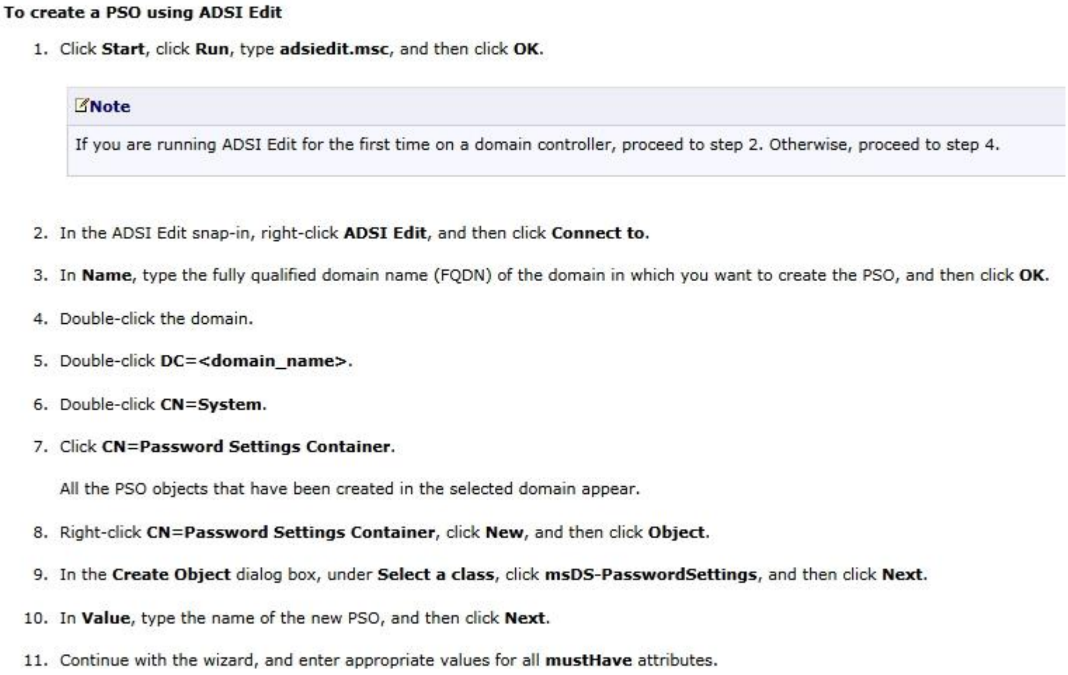

To create a PSO using ADSI Edit

1. Click **Start**, click **Run**, type **adsiedit.msc**, and then click **OK**.

   Note

   If you are running ADSI Edit for the first time on a domain controller, proceed to step 2. Otherwise, proceed to step 4.

2. In the ADSI Edit snap-in, right-click **ADSI Edit**, and then click **Connect to.**
3. In **Name**, type the fully qualified domain name (FQDN) of the domain in which you want to create the PSO, and then click **OK.**
4. Double-click the domain.
5. Double-click **DC=<domain_name>**.
6. Double-click **CN=System**.
7. Click **CN=Password Settings Container**.

   All the PSO objects that have been created in the selected domain appear.

8. Right-click **CN=Password Settings Container**, click **New**, and then click **Object**.
9. In the **Create Object** dialog box, under **Select a class**, click **msDS-PasswordSettings**, and then click **Next**.
10. In **Value**, type the name of the new PSO, and then click **Next**.
11. Continue with the wizard, and enter appropriate values for all **mustHave** attributes.

QUESTION NO: 27

Your network consists of a single Active Directory domain. The functional level of the domain is Windows Server 2008 R2. All servers run Windows Server 2008 R2. A corporate security policy requires complex passwords for user accounts that have administrator privileges.

You need to design a strategy that meets the following requirements:

- Ensures that administrators use complex passwords

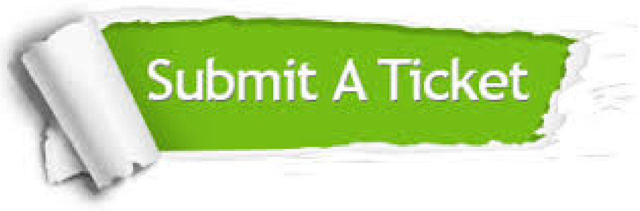
Submit A Ticket