

Microsoft

Server 2008

Enterprise Administrator

(70-647)

Microsoft Certified
IT Professional (MCITP)



**Smarter
Training**

This LearnSmart exam manual helps candidates demonstrate proficiency when taking the Microsoft 70-647 Exam, Server 2008 Enterprise Administrator. By explaining exam topics clearly and directly, the manual equips IT professionals with the knowledge and confidence necessary to pass the exam and gain credit towards earning Microsoft certification.

Topics covered in this manual include:

- Network and Application Services
- Core Identity and Access Management Components
- Support Identity and Access Management Components
- And more!

Give yourself the competitive edge necessary to further your career as an IT professional and purchase this exam manual today!

Windows Server 2008 Enterprise Administrator (70-647) LearnSmart Exam Manual

Copyright © 2011 by PrepLogic, LLC
Product ID: 012305
Production Date: July 12, 2011

All rights reserved. No part of this document shall be stored in a retrieval system or transmitted by any means, electronic, mechanical, photocopying, recording, or otherwise, without written permission from the publisher. No patent liability is assumed with respect to the use of the information contained herein.

Warning and Disclaimer

Every effort has been made to make this document as complete and as accurate as possible, but no warranty or fitness is implied. The publisher and authors assume no responsibility for errors or omissions. The information provided is on an “as is” basis. The authors and the publisher shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this document.

LearnSmart Cloud Classroom, LearnSmart Video Training, Printables, Lecture Series, Quiz Me Series, Awdeeo, PrepLogic and other PrepLogic logos are trademarks or registered trademarks of PrepLogic, LLC. All other trademarks not owned by PrepLogic that appear in the software or on the Web Site (s) are the property of their respective owners.

Volume, Corporate, and Educational Sales

Favorable discounts are offered on all products when ordered in quantity. For more information, please contact us directly:

1-800-418-6789
solutions@learnsmartsystems.com

International Contact Information

International: +1 (813) 769-0920

United Kingdom: (0) 20 8816 8036

Table of Contents

What to Know	6
Tips	6
Domain 1 – Planning Network and Application Services	6
Objectives.....	6
IP Addressing.....	6
<i>Syntax</i>	7
<i>Address Types</i>	7
<i>Global Unicast Addresses</i>	8
<i>Advantages</i>	9
<i>IPv4 Compatibility</i>	9
<i>IPv4 to IPv6 Transition</i>	10
Network Access	10
<i>Network Perimeters</i>	11
<i>Remote Access</i>	11
<i>VPN Methodologies</i>	12
<i>VPN Authentication</i>	12
<i>VPN Authentication Services</i>	13
<i>Network Access Protection</i>	13
<i>DHCP Enforcement</i>	14
<i>VPN Enforcement</i>	14
<i>802.1x Enforcement</i>	14
<i>IPsec Enforcement</i>	14
<i>Domain and Server Isolation</i>	15
Name Resolution	15
<i>WINS</i>	15
<i>DNS</i>	16
<i>Servers</i>	16
<i>Zones</i>	16
<i>Zone Replication</i>	17
<i>Forwarding</i>	17
<i>Windows Server 2008 DNS Enhancements</i>	18
<i>Command Line Tools</i>	18
Application Delivery.....	18

<i>Locally Installed Software</i>	18
<i>Terminal Services</i>	19
<i>Application Virtualization</i>	21
Domain 2 – Designing Core Identity and Access Management Components	22
Objectives	22
Forest and Domain	22
Functional Levels	24
Schema Modifications	26
Trusts	26
Physical Implementation	27
<i>Site Structure</i>	27
<i>Replication</i>	27
<i>Domain Controllers</i>	29
<i>Printer location policies</i>	32
Administration	32
Group Policy	34
<i>Policy Prioritization</i>	34
<i>Device Installation Management</i>	35
<i>Authentication and Authorization</i>	36
Domain 3 – Designing Support Identity and Access Management Components	38
Objectives	38
Updating Active Directory	38
Inter-Forest Trusts	40
Interoperability	42
Branch Office Deployment	43
<i>Organization</i>	43
<i>Infrastructure</i>	44
<i>Security</i>	46
<i>Administration</i>	47
Public Key Infrastructure	48
<i>PKI Components</i>	48
<i>Certificate Templates</i>	51
<i>Certificate Life Cycle Management</i>	51

Domain 4 – Designing for Business Continuity and Data Availability	53
Objectives.....	53
Business Continuity Planning	53
<i>Service Availability</i>	53
<i>Network Load Balancing</i>	54
<i>Failover Clustering</i>	54
<i>Windows Server Backup</i>	55
<i>Active Directory Recovery</i>	56
<i>Data Security</i>	57
<i>Encrypting File System (EFS)</i>	58
Data Sharing and Collaboration	59
<i>Distributed File System (DFS)</i>	58
<i>Windows Sharepoint Services (WSS)</i>	60
<i>Microsoft Office Sharepoint Server (MOSS)</i>	60
Server Virtualization	61
<i>Windows Server 2008 Hyper-V</i>	62
<i>Virtual Server 2005 R2</i>	63
<i>System Center Virtual Machine Manager</i>	63
<i>Virtual Server Migration Toolkit</i>	64
Patch Management.....	64
<i>Windows Update (WU) / Microsoft Update (MU)</i>	64
<i>Automatic Updates (AU)</i>	65
<i>Windows Server Update Services (WSUS)</i>	65
<i>Windows Update Agent (WUAgent)</i>	66
<i>System Center Configuration Manager (SCCM) 2007</i>	69
<i>System Center Essentials (SCE) 2007</i>	69
<i>Microsoft Baseline Security Advisor (MBSA)</i>	70
Practice Questions	71
Answers & Explanations	77

What to Know

The Enterprise Administrator exam is designed to demonstrate the candidate's knowledge and skill with enterprise-level infrastructure topics, including networking and application services, identity and access management components, data availability, and business continuity. The exam focuses on planning and design activities, and requires a thorough understanding of all topics at the conceptual level as well as the implementation level.

Tips

This Exam Manual is not intended to be the only study reference, but could be considered the "Cliff's Notes" to the material required. The Exam Manual should be used for quick reference and study after a more extensive level of preparation, involving other study guides and Microsoft product documentation, including the TechNet Library and the Windows Server 2008 Resource Kits. In addition, this document assumes that the candidate has successfully passed the related Microsoft Certified Technical Specialist (MCTS) examinations. A candidate with limited or no experience in the field is recommended to first focus on the related MCTS examinations.

Domain 1 – Planning Network and Application Services

Objectives

Planning network and application services is concerned with all aspects of developing the network infrastructure necessary to deliver functionality to the users, as well as the applications that provide that functionality. The other three domains in this exam cover the support services that provide access and control to the network and applications, but it is this domain which represents the primary purpose that an enterprise technology environment exists.

The specific objectives covered in this domain are:

- Plan for name resolution and IP addressing
- Design for network access
- Plan for Terminal Services
- Plan for application delivery

IP Addressing

The first step in any enterprise network implementation is proper planning of the addressing scheme. This is not limited to just a decision of whether to use 10-dot addresses or a 192.168.x.x network but also proper selection of subnet masks for site networks, subnets for wide-area connectivity, and where to use public IP addresses.

Adding to these historical planning activities, we now have the considerations of whether to implement IPv6 addressing, how to manage the coexistence of IPv4 and IP v6 addressing, whether to transition entirely to an IPv6-only addressing scheme, how to design and implement DHCPv6, and how to configure DNS to support IPv6 addresses. In addition, the extent to which you can use IPv6 will also be dependent on the extent to which your Internet Service Provider (ISP) has implemented IPv6 technologies.

Regardless of what your attitudes are about the implementation of IPv6 within the enterprise organization (and everybody has an opinion on this), the Enterprise Administrator exam requires you to know how to plan an IPv6 environment.

Syntax

The IPv6 address is a 128-bit address (as compared to the IPv4 address which is 32 bits). To facilitate the display of an address four times larger, a new display scheme is presented for IPv6. IPv6 uses eight blocks of four hexadecimal digits separated by colons; each block represents 16 bits of the address.

```
fe80:0000:0000:0000:059e:00b7:0006:fdfc
```

There are two display simplification rules that can be applied to the full 39 character address (32 hex digits plus seven colons):

- Leading zeros within each block of four hex characters can be dropped, but at least one digit must be retained.

```
fe80:0:0:0:59e:b7:6:fdfc
```

- Contiguous blocks of zeros can be compressed to a double-colon, but this can only be done once per IPv6 address. This ensures non-ambiguous reconstitution of the full address, as it can be easily determined how many blocks are contained within a single double-colon, by counting the number of displayed blocks and subtracting from eight.

```
fe80::59e:b7:6:fdfc
```

Like IPv4 addresses, the IPv6 address is split into two parts. The first part is the *prefix* and contains fixed values and network identification bits. The second part is the *Interface ID* and is determined by the network adapter card. Both the prefix and the interface ID are 64-bits in length.

Address Types

The three IPv6 address types are unicast, multicast, and anycast. *Unicast* addresses identify a single logical interface, which may comprise multiple physical interfaces (such as might be seen in network adapter teaming or load balancing scenarios). *Multicast* addresses identify a collection of independent network interfaces, all of which receive the traffic. An *anycast* address also identifies multiple independent interfaces, but only the nearest interface receives the traffic.

IPv6 multicast addresses are identified by the value '0xFF' in the first 8 bits, a 4-bit *flags* field in bits 8-11, and a 4-bit field called the *scope*, which is contained in bits 12-15 of the first 16 bits. Currently the only flag defined is the *Transient* flag. The scope is used by routers to determine whether the multicast packet is forwarded, and it currently defines five possible values (0 and F are Reserved; the other nine are undefined):

Value	Scope
1	Node-local
2	Link-local
5	Site-local
8	Organization-local
E	Global

As a result, multicast addresses are easily identifiable as addresses with the first four hex characters between '0xFF00' and '0xFF1F'.

The IPv6 anycast address is a methodology by which multiple interfaces can be identified by a single address and routers forward an anycast address to the closest interface based on available routing metrics. An anycast address is identified by all zeros in the interface ID and its routing is based on the subnet prefix. At this time anycast addresses are only implemented as destination addresses assigned directly to router interfaces.

Global Unicast Addresses

The unicast addresses are classified in five types. The first are the *special* addresses that mimic the same features IPv4. The *unspecified address* (in IPv4 this is 0.0.0.0) is 0:0:0:0:0:0:0, or using the address display compression rules, can be displayed simply as a double-colon (::). Like IPv4, this address indicates the absence of a network address. The other is the *loopback address*. In IPv4 this is displayed as 127.0.0.1 (or, strictly speaking, any host address in the 127.0.0.0/8 block); in IPv6 this is 0:0:0:0:0:0:0:1, or ::1 using the contiguous zero rule.

In addition to the special addresses described above, IPv6 provides four other address types, of which three are of primary interest to the MCITP: Enterprise Administrator. The other type is a cross-protocol type designed to support ATM or IPX interoperability and is not likely to appear on the examination.

The three types of IPv6 addresses the enterprise administrator should be completely familiar with are:

- The *link-local* address, which is the equivalent of the APIPA (169.254.0.0/16) address. Link-local addresses are identified with the 64-bit prefix fe80:: followed by the 64-bit interface ID.
- The *site-local* address, which is the equivalent of the IPv4 private address space. Site-local addresses are identified with the 64-bit prefix fec0::xxxx, where xxxx is a 16-bit subnet identifier used within the local organization.
- The *global* address, which is the equivalent of the IPv4 public address.

The prefix of a unicast address is broken down into multiple fields and it's necessary to understand each of these fields in order to properly recognize and configure the type of address needed for a particular scenario.

A global unicast address is identified by these attributes:

- The value '001' in the first three bits, known as the *Format Prefix*, is followed by a 13-bit *Top-Level Aggregator* (TLA). Blocks of TLAs are allocated by the IANA to Internet registries, who then reallocate one or more TLAs to Internet Service Providers.
- Bits 24 through 47 are the *Next-Level Aggregator* (NLA) and are assigned by the ISP to uniquely identify a customer or customer site.
- Bits 48 through 63 are the *Site-Level Aggregator* (SLA) and are used to facilitate routing, either to downstream ISPs (if the NLA ID represents a subordinate ISP), or to sites or subnets within a specific customer organization or site.
- Bits 64 through 127, which as noted earlier, represent the Interface ID.

Because the first three bits are '001', we can also easily identify global unicast addresses as having the first block in the range 2000 – 3FFF.

Like their IPv4 cousins, the site-local and link-local addresses are not Internet-routable. Site-local addresses can be auto-configured by the client, using *stateless address configuration*, which is based on a link-local multicast router solicitation message and ICMPv6 router discovery messages. Alternatively, they can be auto-configured using DHCPv6, which provides *stateful address configuration*. Choosing whether to use stateless or stateful address configuration is a primary planning issue for the implementation of IPv6.

While the Interface ID is, by design, determined by the network adapter card, Windows Vista and newer operating systems allow the local generation of a random 64-bit Interface ID, and this is the default behavior for auto-configured addresses. (Generating random IDs precludes the use of the hardware-specific Interface ID for itinerary tracking, which can be done with a laptop or with cellular telephones.) You can restore the use of hardware-based Interface IDs by disabling this behavior with the command:

```
netsh interface ipv6 set global randomizeidentifiers=disabled
```

Advantages

Aside from the commonly known reason for adopting IPv6—larger address space, which has been notably unsuccessful in promoting the adoption of IPv6 at the enterprise level—there are four other significant advantages that should be evaluated when considering adoption of IPv6:

- IPv6 permits the elimination of Network Address Translation (NAT) and traditional subnet masking, including the “higher mathematics” that are often associated with such functionality.
- IPv6 allows for auto-configuration of site-local addresses, thus eliminating the need to manually configure interfaces or implement DHCP.
- IPv6 requires the use of IPSec (which is only optional for IPv4 networks), effectively making IPv6 networks inherently more secure with regards to data protection.
- IPv6 provides better support for Quality of Service (QoS) and real-time data delivery requirements.

IPv4 Compatibility

Four methodologies exist to provide compatibility between IPv4 and IPv6 environments.

- The *IPv4-compatible* address replaces the last two blocks of the IPv6 address with the four-octet dotted decimal representation and is used by IPv6 nodes transmitting data over an IPv4 infrastructure. The IPv6 traffic is encapsulated in an IPv4 packet. If you recall that an IPv4 address is just a 32-bit value, and we’re replacing two 16-bit blocks, you’ll then understand that this is really nothing more than deriving the IPv6 Interface ID from the IPv4 address. The underlying 64-bit interface ID is derived from the assigned (32-bit) IPv4 address and prepended by 32 bits of zeros.

```
259e:75b7:af06:fdcf::192.168.100.250  
259e:75b7:af06:fdcf::c0a8:64fa
```

- The *IPv4-mapped* address prepends 16 bits of 0 and 16 bits of 1 (0000:ffff) in front of the 32-bit IPv4 address to produce the 64-bit interface ID and is used by an IPv4-only node in an IPv6 network infrastructure.

```
259e:75b7:af06:fdcf::ffff:192.168.100.250  
259e:75b7:af06:fdcf::ffff:c0a8:64fa
```

- *ISATAP* (Intra-site Automatic Tunnel Addressing Protocol) is a methodology for IPv6 nodes to communicate across an IPv4 infrastructure. The ISATAP address is constructed with the IPv6 unicast prefix, plus the ISATAP identifier 0:5efe, followed by the 32-bit IPv4 address. The IPv4 address can be a private or public address. The ISATAP address is automatically configured by Windows Server 2008.

```
259e:75b7:0:5efe::192.168.100.250
259e:75b7:0:5efe::c0a8:64fa
```

- *Teredo* is a methodology of mapping a public IPv4 address to an IPv6 address syntax. The Teredo IPv6 address consists of the prefix 2001::/32 plus the 32-bit IPv4 address of the Teredo server.

Note: Windows XP and Windows Server 2003 use an older prefix identifier – 3ffe:831f1::/32 – and need to be updated with MS06-064 (or the latest service pack: XP SP3, Win2003SP2) to use the 2001::/32 prefix. The Teredo client's IP Address is encoded in the final 32-bits by XORing the external client address with 0xFFFFFFFF. You need to be able to recognize Teredo addresses and encode and decode the IPv4 addresses contained therein.

```
2001:0:[Teredo Server]::Teredo Client]
```

Example:

Tunneling from corporate HQ to branch office.

Corporate HQ: 192.168.64.100 = c0a8:4064

Branch Office: 192.168.32.75 = c0a8:204b

and XOR 0xFFFFFFFF = 3f57:dfb4

IPv6 Teredo address is

```
2001:0:c0a8:4064::3f57:dfb4
```

IPv4 to IPv6 Transition

Choosing to transition to an IPv6-only environment is still probably a long-term planning consideration as there are no specific time frames mandated for transition, and until they are mandated— or Windows drops support for IPv4— there are probably very few organizations that can justify the added knowledge and complexity required to convert to an IPv6 infrastructure.

The transition methodologies you should consider, however, all involve one or more of the “compatibility” features discussed in the previous section. Either both versions co-exist in a dual-stack or dual-layer scenario, or one version is tunneled through the other, using Automatic Tunneling (via IPv4-compatible addresses), ISATAP, or Teredo.

Network Access

Ensuring protection of the enterprise network is one of the most important tasks in planning a network, but this has to be balanced against the need for access. There are many methodologies, vendors, and products available to choose from to provide protection and access services.

Network Perimeters

One of the first tasks is identifying the systems and services that need to be accessible from the public Internet, as well as the systems and services that need to be accessible to authenticated and authorized users. The publicly accessible systems will be in the *perimeter network*; the remainder of the systems will be in the *internal network*.

Methodologies used to protect the network perimeter are as basic as *Network Address Translation* (NAT) which provides security through obscurity because the outside world cannot see the actual IP addresses of protected resources, and private IP subnets are not routable on the public Internet.

NAT is not actually a security methodology, it's a routing methodology designed to support Internet access for networks using private addressing. NAT devices are typically combined with a packet-level firewall to help keep out unwanted traffic. A packet-level firewall inspects each packet for the destination IP address and port, verifies that a rule exists to allow traffic to that IP address and port, and blocks all other traffic. It knows nothing about the traffic contained in the packet, and the decision to permit the traffic to pass through is made exclusively on the destination IP address and port. NAT and packet-level firewalls are rarely used by enterprises as a network perimeter security methodology, but are quite common in home networks.

More advanced methodologies include stateful-inspection firewalls, circuit-level firewalls, proxy servers, and the application-layer firewall. Each of these methodologies represents a continuing advancement in the development of network perimeter security, but in today's enterprise, the application-layer firewall is almost exclusively the solution implemented. Some organizations also use proxy servers but more so as a performance solution than a security solution. ISA Server 2006 and Forefront Threat Management Gateway 2010 are examples of Microsoft application-layer firewalls that also provide proxy services.

ISA Server 2006 runs on Windows Server 2003 32-bit operating systems, and Forefront Threat Management Gateway 2010 runs on Windows Server 2008 R2 64-bit systems. The Forefront product is most likely too new to appear on this exam, but you should be familiar with the existence and purposes of ISA Server 2006.

Remote Access

Providing remote access used to be a once-in-a-while activity or something the I.T. department deployed for departmental use to do systems maintenance after hours. Today, though, telecommuters are everywhere, and more companies are exploring the benefits of virtual existence. As a result, providing remote access to the corporate network is a foregone conclusion. You must know how to design and configure a remote access strategy as an enterprise administrator.

The first step to designing a remote access strategy is identifying who the persons are that will require remote access. They'll likely be employees but may also be business partners, customers, or even guests. Guests, you say— why would we let guests connect to our network? Remote access is not just persons in another location accessing the network in your location; it's any person accessing your network who is not an established, authenticated member with a permanently assigned access point (such as a desktop computer). This very likely will include visitors to your facility bringing their notebooks with them to conduct work in the interest of your business.

Another decision point to be made is how the access will be provided. External users needing inbound access might use VPN connections, Remote Web Workplace, Terminal Services, or the new Direct Access feature in Windows Server 2008 R2. Local guests might use an isolated wireless network. You need to be familiar with all of the possible remote access methodologies, be able to identify the appropriate methodology in a given situation, and design an implementation plan for that remote access methodology. Terminal Services is covered in another section, and Remote Web Workplace (an SBS and EBS technology) and Direct Access (an R2-only technology) are not covered on this exam.

VPN Methodologies

Traditional VPN methodologies include *Point-to-Point Tunneling Protocol* (PPTP) which is a password-based protocol and *Layer 2 Tunneling Protocol* (L2TP) which is used in conjunction with IPSec authentication. The use of IPSec authentication, however, requires VPN client systems to have certificates from a trusted root certification authority (CA), and this limits where L2TP can be deployed.

Note: L2TP can be used with a pre-shared key rather than a certificate, but this is really intended for testing purposes only and should not be implemented as a production solution.

Both PPTP and L2TP also require special considerations when traversing a NAT boundary (i.e. the VPN client and VPN server are on opposite sides of the NAT-based device). PPTP requires a NAT editor. The Routing and Remote Access service provides a NAT editor, as well as ISA Server. L2TP requires NAT Traversal (NAT-T), and the feature must be supported by both the client and the server. However, today almost all NAT devices, as well as VPN clients and servers, provide native support for PPTP and L2TP traffic.

PPTP and L2TP require special considerations for the configuration of firewalls. In addition to access through standard TCP or UDP ports, these remote access services also require access for additional IP protocols. Without diving too deeply into Internet Protocol Theory, which is beyond the scope of this exam, you are reminded that IP is a suite of protocols, and TCP and UDP are only 2 of over 140 that are currently defined. Each protocol is identified by a decimal value 0 through 255—TCP is protocol 6 and UDP is protocol 17. Two additional protocols you must be aware of for configuring remote access are General Routing Encapsulation (GRE), protocol 47, and Encapsulating Security Payload (ESP), protocol 50. GRE was developed by Cisco in 1994 and is described in RFC 1701, 1702, 2784, and 2890. ESP was developed by the Naval Research Library in 1995 and is described in RFC 1827, 2406, and 4303.

The PPTP connection requires access through the firewall on TCP port 1723 and IP protocol 47 (GRE). L2TP requires access through the firewall on UDP port 500, IP protocol 47 (GRE), and IP protocol 50 (ESP). The VPN server also requires outbound access on UDP port 4500 (to support NAT-T). You should have sufficient practical skills with configuring VPN servers and VPN clients for both PPTP and L2TP VPN connections, as well as determining when each is an appropriate solution given the conditions. A new VPN technology available in Windows Server 2008 (and introduced in Vista SP1) is the Secure Sockets Tunneling Protocol (SSTP). The primary advantage to SSTP is that it uses SSL for encryption and is transmitted over the well-known SSL port 443. Its primary disadvantage is that it's only available on the newest Microsoft operating systems and will not be available for Windows XP systems. As with L2TP/IPSec, it also requires a certificate from a trusted root CA.

VPN Authentication

After determining how connections will be made, you will need to be able to determine who is using those connections. There are several authentication protocols available for use. You need to be able to identify each of the protocols and know their relative levels of security and the limitations of each—most notably when they cannot be used because legacy systems do not provide the functionality. The six authentication protocols supported in Windows Server 2008 are:

- PEAP-TLS
- EAP-TLS
- PEAP/EAP-MSCHAP v2
- MSCHAP v2
- MSCHAP
- PAP

As you can see the list includes several permutations of only a few technologies. PEAP and EAP are authentication mechanisms, and when they are transported via TLS, PEAP and EAP use certificate-based authentication. The primary difference between PEAP and EAP is when the exchange of certificates for authentication takes place. PEAP exchanges the certificates after the TLS tunnel has been established, thus providing an additional layer of protection to the certificates. When PEAP or EAP are used with MSCHAP v2, the authentication is password based. Not all systems support PEAP or EAP, so for legacy clients it may be necessary to provide authentication using MSCHAP v2 only. The last two, MSCHAP and PAP, are extremely weak methodologies and should only be used when no other authentication is available to the client. If you must use MSCHAP or PAP you should ensure that the use of strong passwords is enforced.

VPN Authentication Services

VPN authentication can occur by a number of means. The simplest, as when using Routing and Remote Access services, is by simple password authentication with a local user account (on the VPN/RRAS server) or a domain account (when the VPN/RRAS server is deployed in a domain environment). This works fine when an organization has a single VPN server and a small number of VPN users who are employees, but for an enterprise organization that has a large number of VPN users or VPN servers, that may span across multiple domains or non-domains, this methodology simply is not practical.

For enterprise authentication needs, RADIUS is the starting point for this solution. *Remote Authentication Dial In User Service* (RADIUS), as its name implies, dates back to the days of authenticating dial-in users on analog modems. It was developed in 1991 and implemented widely by Internet Service Providers, as well as enterprises that used it to address all sorts of access and authentication control needs.

Microsoft has supported some form of RADIUS service since Windows 2000. In Windows Server 2003, it was given its own name—Internet Authentication Service (IAS). In Windows Server 2008 it has once again been renamed to *Network Policy Server* (NPS) and support for RADIUS Proxy Servers is also introduced. This allows the client authentication point to be moved closer to the client—within the DMZ, for example—but still retain optimal security on the authentication server itself by locating it in the internal network and being able to explicitly identify the resources it must communicate with (the proxy servers). RADIUS proxy servers also provide the capability to support cross-forest authentication, which was not possible with IAS.

You should be able to identify the appropriate deployment strategies for the Windows Server 2008 Network Policy Server and appropriate RADIUS proxy servers.

Network Access Protection

Network Access Protection (NAP) is the principle of only allowing “safe” entities to access the network, enforcing those access restrictions through the validation of minimum system requirements, and providing the resources to remediate those deficiencies. The intent of NAP is not to guarantee that a resource being connected to the network is absolutely secure, but rather to minimize the threat that the resource poses by ensuring the known and predictable risks are mitigated.

There are three stages of the process of applying network access protection. The first is *health state validation*. An agent on the resource conducts an analysis of the resource configuration and reports the current health state to the NAP server. This would include that AV software is installed and status information about the AV software but not the actual performance of an AV scan. It might include information about the presence of known security updates but not whether the machine had been infiltrated prior to the update being applied.

The second stage is *health policy compliance*. Administrators create health policies—rules that establish the minimum acceptable state of a resource. A comparison of the health state against the health policies determines the compliance level. Compliance results can be simply logged and then access granted to the resource, or the resource may be granted *limited access* for the purpose of remediating the identified compliance deficiencies. Limited access, if employed, is the third stage.

There are four fundamental methodologies for implementing Network Access Protection in a Windows environment. You should be familiar with planning and designing the implementation of each of these in appropriate scenarios.

- DHCP enforcement – for IPv4 clients, by blocking access to the DHCP server for non-compliant computers
- VPN enforcement – for PPP-based VPN clients using PEAP authentication
- 802.1x enforcement – for 802.1x capable clients using EAP authentication
- IPsec enforcement – for IPsec capable clients

DHCP Enforcement

DHCP enforcement is an IPv4-only solution. It requires that all DHCP servers be running Windows Server 2008 and is easily overridden by a local administrator on the candidate computer. This methodology is considered the weakest of the four and should only be implemented when the other three are not suitable solutions.

VPN Enforcement

VPN enforcement can be implemented on PPP-based VPN connections that use PEAP authentication. It is, perhaps, the simplest NAP enforcement methodology to implement.

802.1x Enforcement

The primary issue with 802.1x enforcement is that all devices in the network infrastructure must be 802.1x compliant or be exempted from the enforcement policies. It requires the use of PEAP for authentication, and it precludes the use of PXE boot on 802.1x enabled switch ports. It can be used with wireless networks and can be an excellent solution for guest access to a small number of resources in a business facility when used with PEAP-MSCHAPv2.

IPsec Enforcement

IPsec is the most secure enforcement methodology. Its primary advantages are that local administrators cannot bypass the enforcement process on the candidate computer, and upgrades to network infrastructure devices are not necessary, as is the case with the 802.1x enforcement methodology. The primary disadvantages are the complexity of the solution, requiring additional server hardware, the creation of a public-key infrastructure if one does not already exist, and the creation of zones to separate the authorized network communications pathways during the compliance and limited access stages of the process.

The internal network infrastructure of an IPsec enforcement environment contains three logical zones. The VPN or other remote access server is contained in the perimeter network and is accessed by resources coming from the Internet (or, perhaps, a local wireless access point connected outside the internal network). The *restricted network* contains the non-compliant (or non-NAP-capable) candidate computers. The *boundary network* contains the NAP resources used to establish the compliance status and the resources necessary for facilitating remediation, including any necessary infrastructure support services. This zone would include DHCP, DNS, AD domain controllers, AV signature servers, Windows Server Update Services (WSUS) and the NAP server(s). These servers are accessible to the new resource to facilitate the evaluation of compliance, and the communication between these components and the communication is unauthenticated. The *secure network* contains all computers that have been certified as compliant with the established health policies.

Domain and Server Isolation

Domain isolation refers to a technique of controlling access to resources in the domain by requiring that all resources use IPsec. This effectively isolates untrusted (non-domain) resources because they do not have the necessary certificates and policy-based IPsec configurations to establish communication with trusted resources.

Server isolation applies the IPsec capabilities as well as other technologies to ensure that all client connections to the server are authorized. In addition to IPsec, the Windows Firewall with Advanced Security can be used to restrict access by client IP Address. On Windows Server 2003, the "Access This Computer From The Network" local security policy can restrict access to specified users or computers, and in Windows Server 2008 you can base firewall access rules on Active Directory Security Groups.

Name Resolution

Planning for a name resolution service requires an extensive understanding of the concepts of the Domain Name System (DNS), as well as a determination as to whether the Windows Internet Name Service (WINS) is required in your environment.

WINS

Perhaps the first question to dispense with when planning a name resolution strategy is determining whether WINS is necessary or not. WINS is a Microsoft developed name service designed for use on NetBIOS networks. In general, if you have legacy network technologies still active, such as Windows NT or workgroup-based networks, you likely still require WINS to some extent. If your environment is almost exclusively Active Directory-enabled using Windows 2000 or later systems, you can probably dispense with the use of WINS.

If you determine that you need WINS, you need to identify the extent to which you need to deploy the capability. At a minimum you should have two WINS-enabled resources to ensure availability of services should one resource be offline. These servers could be configured using clustering as a single fault-tolerant WINS server, or they could be configured as two independent WINS server using replication.

If your WINS requirements are more widely spread geographically, you may require additional remote site servers to be enabled with the WINS service.

In a multi-server implementation, there are three replication topologies you should be familiar with:

- Hub and spoke—a central WINS server with downstream WINS servers who replicate only with the central (hub) WINS server.
- Ring—each WINS server is identified with a neighboring partner for replication. Variations on the implementation of this topology depend on whether the replication partners are push-only, pull-only, or push-pull, and whether the replication partnerships are uni-directional around the ring or bi-directional. Each of these permutations will impact the convergence time of any given change.
- Full-mesh—every WINS server is configured to replicate with every other WINS server. This topology reduces convergence, but network traffic is significantly higher, and the complexity introduces security risks and a higher level of management support.

An alternative to WINS, introduced with Windows Server 2008, is the GlobalNames zone in DNS which creates the ability to maintain a namespace of single-label (non-domain) names.

DNS

One of the first questions to be resolved in planning DNS is determining where you need additional DNS servers and what type of DNS servers they should be. The second planning question is to determine which servers should host which zones and what type of zone they should be. Planning for the domain names, specifically, is covered in Domain 2 as a planning function of Active Directory. Let's discuss how DNS servers are deployed, what types of zones can be deployed, and the circumstances in which they would be used.

Servers

DNS Servers are typically located on Domain Controllers; however, DNS can be installed as a standalone service if required. Standalone servers would be used in the perimeter network, or when Active Directory has not been deployed, such as when the primary DNS services are BIND-based.

BIND is an acronym for Berkeley Internet Name Domain and is the original Internet naming service developed in the 1970s for the Berkeley Software Distribution (BSD), aka Berkeley Unix, derivatives of which are still available today in the form of FreeBSD and NetBSD. BIND is the DNS implementation used on Unix systems.

A DNS server can also be deployed simply to provide a centralized cache of commonly used names. These servers are known as caching-only servers and are deployed to distribute client load away from the primary servers or to minimize DNS traffic travelling across a wide-area network.

Every node in a network should have at least two DNS servers configured as resolution resources, and at least one of them should always be online. DNS server assignment priorities are typically handled through the assignment of DNS servers in DHCP scope options so that each DHCP scope is assigned to the closest, or most reliable, DNS servers for the clients in that scope.

Zones

Typically DNS zones are deployed using Active Directory Integrated zones, which require that DNS is installed on a DC. A DNS zone can be deployed as a file-based primary or secondary zone where Active Directory integration is not possible, or not desirable. When file-based zones are used, there can be one primary zone, which is the authoritative zone data and any number of secondary zones, which are read-only copies of the primary zone. With AD Integrated zones, every instance is a primary and every instance can be written to. Changes to the DNS zones are replicated using standard AD replication services.

Other types are the reverse, stub, and GlobalNames zone. A reverse zone should exist for each IP network in the enterprise. The reverse zone contains PTR records that are used to map IP addresses to the canonical hostname for each address. A stub zone is used to provide reference information for where the DNS servers are for another domain. Stub zones are typically used by parent DNS servers to identify the authoritative DNS servers for delegated subdomains. A stub zone contains an SOA record, one or more NS records identifying the subdomain's DNS servers, and an A record for each NS record identifying the IP address of each subdomain DNS server. The GlobalNames zone is new in Windows Server 2008 and provides name resolution services for single-label names such as exists in NetBIOS networks. It is intended as a replacement for WINS and facilitates the retirement of NetBIOS over TCP/IP (NetBT).

Zone Replication

Perhaps the most common failure point in DNS is the zone transfer, when a secondary server is unable to obtain zone data from the primary server. AD Integrated zones replicate DNS changes through the AD DS replication services, so this is rarely an issue with AD Integrated zones, but standalone secondary servers must pull the file-based zone data from the primary server. This requires configuration on both sides of the connection. The secondary server must be configured with the IP address of the primary server, and the primary server must be configured to allow the secondary server(s) to obtain the zone data. Since Windows Server 2003, zone transfers have been performed using incremental transfers.

Forwarding

Identifying the need to use forwarders is an important planning decision. Should a DNS server perform lookups for external addresses itself, or should it hand off that task to another server? One advantage of using forwarders is that the forwarding server may have a larger cache of resolved addresses, and the time it takes to refer the request to a forwarder and obtain the answer from that server's cache is likely less than the amount of time it would take for the local server to perform the lookups itself. Another advantage is simply reducing the workload that must be performed by the local DNS server resolving external addresses, which provides more resource availability for resolving internal addresses—or providing other services configured on that server. A typical use of forwarding is to direct all Internet name resolution requests to the DNS servers of an organization's Internet Service Provider.

In addition to generic forwarding of all requests for an external name, *conditional forwarders* can also be configured. A conditional forwarder allows the configuration of specific DNS servers to be targeted to resolve names from specific domains and is generally useful where it is not possible to configure a secondary zone. For example, a conditional forwarder is useful in scenarios where VPN connectivity to other domains is used. The local DNS server can be configured to route all name resolution requests for the remote domain to the name servers for the remote domain—names that are not likely resolvable by any local resource anyway, given that they're internal names of another private domain. A conditional forwarder is also useful in local multi-forest/multi-domain environments where there are disparate DNS services, such as might be the case when an organization is going through a merger or acquisition. Resolution requests for a selected domain can be directed to a known authoritative DNS server for that domain, within the same forest, or even cross-forest. Finally, conditional forwarding might also be implemented to integrate with an existing BIND-based DNS service. AD Integrated DNS is used to provide name resolution for the AD domain, typically a delegated subdomain of the primary organizational domain managed by BIND-based services, and a conditional forwarder is set up for the parent domain redirecting those resolution requests to the BIND servers.

Windows Server 2008 DNS Enhancements

The *Read-Only Domain Controller* (RODC) is a new feature of Windows Server 2008. When an AD Integrated DNS zone is implemented on an RODC, the net effect is the same as a standalone secondary zone – the zone is read-only. When a new resource is added to the domain at a site where an RODC is deployed, the DNS records must be updated at a writable DC/DNS server.

Another enhancement provided in Windows Server 2008 is *Background Zone Loading*. Historically, the startup process for a DNS server has been linear and a foreground process. All zone files had to be loaded before the DNS server could begin responding to queries. Background Zone Loading allows the loading of an AD Integrated zone to be deferred to permit faster initiation of responses to pending queries. If a query is received that requires data from the unloaded zone, the data can be accessed asynchronously and directly via the AD store.

Command Line Tools

The Windows certification exams seem to be really big on command line tools, particularly the *dnscmd* tool for configuring DNS services. This may be due to the absence of the GUI interfaces on a Server Core implementation of Windows Server 2008. Therefore, for DNS management, in addition to being intimately familiar with the **dnscmd.exe** command, you should also be familiar with the **ipconfig.exe** command, specifically the use of the parameters **/registerdns**, **/displaydns**, and **/flushdns** and the **nslookup.exe** command which is a very rich diagnostics tool.

Application Delivery

A significant consideration in administering an enterprise network is planning for the delivery of applications to users. Applications are the reason the network exists in the first place. The remainder of this domain covers methodologies for getting applications in the hands of the users. We'll look at three delivery options. The first is locally installed software, which is the traditional method by which applications are accessed. The second is centrally installed software. Historically centrally installed software has been accessed via a full desktop virtualization scenario such as Terminal Services. Recently, though, the concept of virtualizing the application itself has developed, and this can be done either through new features of Terminal Services or via Microsoft's Application Virtualization (App-V) technologies.

Locally Installed Software

Since the invention of the personal computer, the traditional method for accessing and using applications software was installing the software on each user's desktop computer. For many organizations, this was the most laborious part of all of their computer administration and support activities. Maintaining installed applications was just as time-consuming. Luckily time has brought us newer and more efficient technologies for delivering applications to the desktop.

Group Policy

When Microsoft developed Active Directory and Group Policies for Windows 2000 Server, one of the features built into the Group Policy capability was application software deployment services. There are three methodologies available via Group Policy for delivery of application software:

- Publish a software package to a user
- Assign a software package to a user
- Assign a software package to a computer

Publishing a software package makes it available to install via Programs and Features (or Add/Remove Programs on Windows XP/2000/2003). Assigning a package either imposes the installation of the software upon the entity or makes it available for installation via the desktop or Start Menu.

For most organizations, the most common usage is assigning a software package to a computer. This process causes the software package to be installed on the computer the next time it powers on and authenticates with the domain and makes the application available to all users of the computer. A package can be assigned to a user at logon, and the package will be installed the next time the user logs onto the computer; a package can be assigned to a user on demand and the package is advertised on the desktop and Start Menu. On first access from the desktop or Start Menu the installation is performed.

System Center

Enterprise organizations may also choose to implement application delivery through System Center Configuration Manager (SCCM) or its mid-size cousin, System Center Essentials (SCE). Application delivery is only one of a suite of services provided by the System Center product family. SCE is designed for use in small and mid-size organizations with less than 500 client systems and 30 servers and leverages WSUS for software deployment. Only one self-contained SCE server can be deployed per domain, so for organizations with more than 30 servers or more than 500 clients, SCCM is the correct solution. The key planning note to keep in mind about SCCM is that it is an agent-based product, and the agent must be deployed to the target systems before the services can be used. In most organizations using SCCM, the agent is incorporated into the OS image (which is also deployed to new systems using SCCM).

Windows Deployment Services

Finally, application delivery can be achieved by “baking it in” to the core operating system image installed to the enterprise computers by using Windows Deployment Services (WDS). WDS is a service provided by Windows Server 2008 (first available in Windows Server 2003 R2) for deploying base level operating system installations to client systems. However, WDS allows for customization of the installation image, and can be customized to include additional hardware drivers, as well as applications.

Terminal Services

Another methodology for centralized management of application delivery is the use of Terminal Services (TS). TS has existed in Windows since the days of Windows NT and is an enhanced version of the same technology that provides Remote Desktop Connection (RDC) services. Whereas RDC is limited in the number of simultaneous connections that can be made to a system, allows connections to desktop systems, and is native in the operating system requiring no additional licensing, Terminal Services is none of those. Terminal Services requires additional licensing, can only be installed on Windows Server systems, and is not limited to a fixed number of connections, but rather only by the purchased licenses and hardware capacity of the server.

Licensing

Licensing of a Terminal Server environment is achieved through implementation of a Terminal Server License Server. Terminal Server Client Access Licenses (TS-CAL) are purchased, and the licensing authorization information is recorded in the Licensing Server. The License Server keeps track of the number of devices and users using Terminal Server services and ensures that the licensing limits are not violated.

License Servers are deployed to manage a scope of client connections. That scope can be an entire forest, a specific domain in a forest, or a workgroup when the license server is not a member of an AD domain. The scope of the license server is determined based on how CALs are purchased—organizationally (at the forest level) or at a lower level, such as business units, subsidiaries, or departments (at the domain level).

In addition to purchasing CALs, the license server must also be activated with Microsoft. Activation is a process of obtaining a Microsoft-issued digital certificate which validates server ownership and identity. This certificate is used to facilitate online purchasing of additional CALs. You can activate a license server via a wizard on the license server, a web page from another computer (when the license server does not have Internet access), or by telephone. If a license server is not activated, it can only issue temporary CALs, which are valid for only 90 days. In addition, there are special considerations for planning disaster recovery or intentional deactivation of a license server. A license server can only be deactivated via the wizard or telephone call.

Finally, when mixing Terminal Servers on different server operating systems, the license server must be running on Windows Server 2008 in order to support licensing for the older operating systems. This is because Win2000 and Win2003 TS license servers cannot issue licenses for connections to Windows Server 2008 Terminal Servers. License servers can be deployed on multiple independent machines to provide load-balancing or high-availability.

TS-CALs come in two types: per-Device and per-User (essentially like Windows CALs). There are some additional differences in the two types unique to Terminal Services. The license server automatically reclaims per-Device CALs after a random period of 52-89 days. If all per-Device CALs are in use (i.e. there are some not yet reclaimed), but not all licensed devices are connected, you can revoke up to 20% of the issued per-Device CALs, based on a specific client OS system. Contrasting that, the per-User CAL limits are not strictly enforced by the license server and it is possible to exceed the actual licensing limits.

Server Infrastructure

The infrastructure of a Terminal Services deployment may consist of one or more servers providing various roles or services. First, one or more Terminal Servers may exist which provide access to client sessions. Terminal Servers may be configured as a collection of independent servers with specific servers assigned for specific clients, or they may be configured as a farm of servers to provide load balancing or high availability.

As discussed above, the TS server infrastructure must contain at least one license server. On a small single-server terminal server implementation, it's sufficient to deploy the license server on the terminal server machine. If multiple terminal servers are deployed or a terminal server farm, the license server should be on two or more separate machines, possibly a dedicated machine depending on the number of licensing transactions that will take place in a given time. If dedicated licensing servers are to be used, they may be candidates for virtualization.

Adequate preparations must be made to ensure the availability of the license server. Losing a single license server in an environment dependent upon terminal services can be catastrophic. The only time an organization should have only one license server is when it's installed on the only terminal server and that terminal server does not host mission critical applications.

The TS infrastructure may include Terminal Servers with Web Access (TS-Web) services. The TS-Web service allows connections to the Terminal Server to be made via a webpage published from the TS-Web server, rather than directly via the RDC client. The web link still uses the RDC client (as opposed to the Active X control used in Windows Server 2003 Terminal Services). In addition, the URLs of the published TS-Web services can be distributed directly to the IE Favorites list via Group Policy.

In a TS farm deployment, Windows Server 2008 provides the Terminal Server Session Broker (TS-SB) service to manage connections to the farm. The TS-SB service monitors all nodes in an NLB cluster and ensures that new client connections are connected to the terminal server with the largest amount of free resources. If the farm is implemented using DNS Round Robin, the TS-SB has the added ability to remember where a client is connected and reconnect a dropped connection to the same server upon reconnect.

Another new feature in Windows Server 2008 is Terminal Services Gateway (TS-Gateway). TS-Gateway provides services to use RDC sessions over SSL to connect Internet clients with protected Terminal Servers or to provide server isolation by implementing IPsec between the TS-Gateway and the Terminal Server(s) and requiring all connections to process through the TS-Gateway.

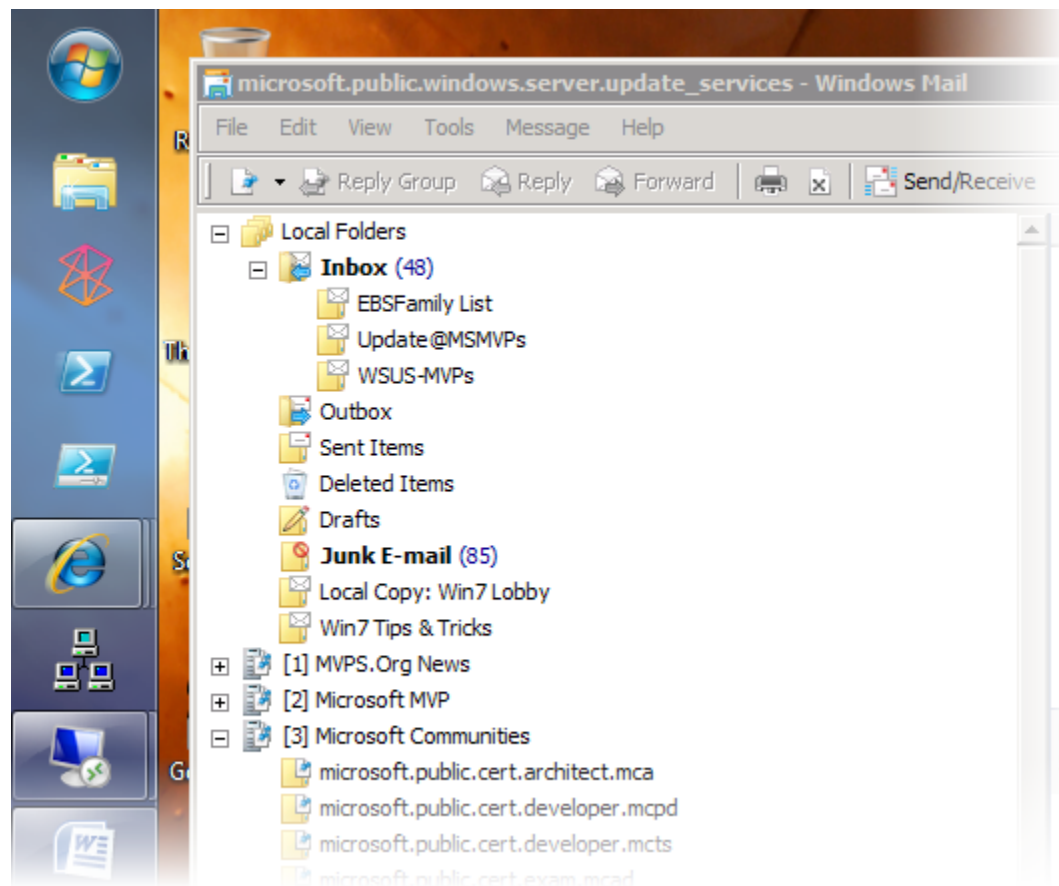
The TS-Gateway uses Connection Authorization Policies (TS CAP) and Resource Authorization Policies (TS RAP) to determine which users can access which resources. A TS CAP can specify the type of authentication required (smart card or password). NAP can be implemented to verify TS-Gateway clients pass health validation checks. A TS-Gateway user must match at least one TS CAP and one TS RAP to gain access to internal resources.

Application Virtualization

The latest developments in the delivery of applications to users are in the area of application virtualization. There are two methodologies of interest to the enterprise administrator: RemoteApp, which is a Terminal Server feature and App-V, which is Microsoft's renamed implementation of SoftGrid.

RemoteApp

With RemoteApp, instead of virtualizing the entire desktop session, an application is configured such that only the application window itself is displayed on the desktop, rather than the entire RDC desktop session. The net effect is that the application appears as if it's running on the local machine, although it does not use the desktop visualization themes that might be seen on Vista or Windows 7. For example, this is an image of Windows Mail on a Windows 7 desktop running from a Windows Server 2008 Terminal Server using RemoteApp.



App-V

App-V differs from RemoteApp in that the application does not run on the server and merely transmit the display to the client, it actually streams the application bits to the client system using the Real-Time Transport Protocol (RTP) and the application is executed on the client system's CPU using the client system's RAM. Once the bits are streamed, they're also cached on the local machine for subsequent executions.

More significantly, though, App-V runs the application in a logically separate partition isolating it from other applications running on the client. This also provides the ability to run incompatible applications or multiple versions of the same application simultaneously on the same client, as well as an additional level of isolation from the operating system resulting in better client system stability. To deploy incompatible applications or multiple versions using RemoteApp you'd actually be required to deploy a second Terminal Server. Unlike RemoteApp, though, which uses the native RDC client, App-V requires the installation of a client-side agent. The agent is essentially a VM host which provides a virtualized environment in which to run the App-V published application.

Deployment of applications via App-V also requires a bit more setup work than RemoteApp. With RemoteApp you simply use a wizard to build an RDP config file and distribute that config file to the client systems. App-V requires "sequencing" to package an application for App-V deployment, the System Center Virtual Application Server, and a database to maintain all of the metadata about the published applications.

In addition, App-V applications can be deployed to Terminal Servers and published via Terminal Services avoiding the need to deploy the desktop agent.

Domain 2 – Designing Core Identity and Access Management Components

Objectives

Designing core identity and access management components is about building the infrastructure that establishes who is authorized to access your network and what they are authorized to do once they do access it. In a practical sense, this involves planning and designing Active Directory and Group Policy. The specific objectives covered in this domain are:

- Design Active Directory forests and domains
- Design the Active Directory physical topology
- Design the Active Directory administrative model
- Design the enterprise-level group policy strategy

Forest and Domain

The first requirement of designing an Active Directory (AD) environment is understanding the distinctions between the forest and the domain. The term "domain" has been in the Windows vocabulary for many years, dating back to the days of Windows NT; however, its scope and purpose in Active Directory is different than it was in Windows NT, and some of what a domain used to provide is now provided by the forest.

A *forest* is an organizational security boundary and identifies a specific purpose for the existence of an Active Directory implementation. You may be deploying AD to serve as a network operating system simply controlling who can use your corporate network resources. You may be deploying Active Directory to function as an enterprise directory with an extended data schema providing storage for various types of enterprise data, some of which may be sensitive or confidential. Some may even represent business trade secrets or be used as an Internet directory, which might also contain publicly-accessible information related to customers or business partners. You might have requirements for more than one directory. Understanding that the forest is the primary security boundary should make it clear that separate functional purposes likely dictates separate Active Directory forests.

You might also deploy multiple forests when it is desirable to establish separate security boundaries within the same type of directory service, for instance when a separate forest is used to manage resources or when user accounts or data need to exist in isolation from the rest of the organization, particularly administrators who would otherwise have unfettered access to all resources in a forest.

A *domain* is a logical boundary that may be used to implement administration containment, data replication controls, geographical identities, or network connectivity considerations. Most organizations only require one domain; however, organizational requirements may suggest the creation of more than one domain. Some examples of scenarios which may benefit from the creation of additional domains include: delegation of administration for separate business units, multiple business locations separated by bandwidth constrained connections which would impact replication of Active Directory data, or a global organization that wishes to establish geographically-related identities, such as different countries of operation.

Two key principles related to the creation of forests and domains are autonomy and isolation. *Autonomy* is related to the independence of control of a particular entity or performance of a task. An administrator with autonomy has the authority to make the decisions about how an entity is managed or how a task is performed. However, the authority is not exclusive. Everybody has a supervisor, and supervisors usually have a wider scope of authority, including the authority to revoke the autonomy of the administrator. *Isolation* confers exclusivity of the control of an entity or the performance of a task and protection from having the entity or task interfered with by other administrators. Forests, being the security boundary, provide both autonomy and isolation in an organization. A forest can be managed in ways that are appropriate for the forest and not impacted by other forests. Domains provide only autonomy, not isolation. While the domain may be independent of other domains, the forest which encompasses a particular domain still has impact upon that domain.

Autonomy and isolation are also considered within the scope of the categories of tasks involved in administering Active Directory. *Service management* includes tasks related to the operation of the Active Directory service itself, and *data management* is concerned with the data contained in the Active Directory service.

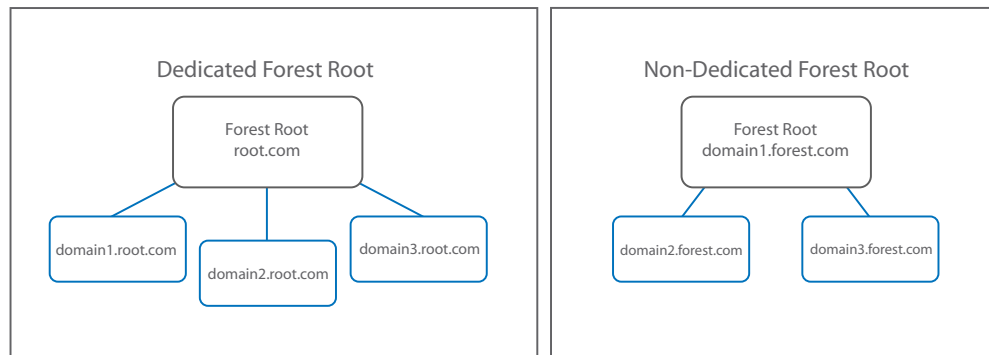
The full matrix of service autonomy, service isolation, data autonomy, and data isolation should be considered when planning for forest and domain creation. Depending on what capabilities are required, an organization might deploy a single forest: the Organizational forest model, multiple separate forests: the Restricted Access forest model, or multiple forests with trust relationships allowing cross-forest access: the Resource forest model.

Data isolation, as well as service autonomy or service isolation, can be achieved with the organizational forest model. Service isolation can also be implemented by using the resource forest model, and if data isolation is a requirement, then the restricted access forest model can be used. Finally, data autonomy can be met by simply joining an existing forest. No new forest is required.

In addition to these considerations, the availability of connectivity can also drive the choice of forest models. To provide for limited connectivity with service autonomy, use a separate organizational forest model or domain in an existing forest in combination with appropriate firewall configurations to restrict access. Limited connectivity with service isolation should use a resource forest with firewall support. If a limited connectivity scenario is desired but service autonomy and isolation are not a contributing factor, an organizational forest will suffice, as will simply reconfiguring the firewall to restrict access to the protected resources.

So, the first step in designing Active Directory forests and domains is to determine whether more than one forest and domain is needed and to document the specific reasons why multiple forests or domains are being implemented.

If multiple domains are required, the next step is to determine the appropriate configuration for the forest root domain. This is a decision to make carefully because unlike simple domain creation, the forest root domain is not a reversible decision. There are two approaches available, and the decision to use one or the other depends upon whether there is a need for isolation between forest service administration and domain service administration. If separation is required, then a dedicated forest root domain should be created. A dedicated forest root domain does not contain user accounts and continues to exist unaffected by any other organizational changes that might affect domain restructuring or renaming, but it does generate additional work to maintain the separate domain and forest resources. If a dedicated forest root is not created, then one of the domains must be selected to function as the forest root domain. The forest root domain is, simply put, the first domain created in a new forest. It's the domain created when the first domain controller in the forest is created. The image depicts the differences between a three domain environment created with a dedicated forest root domain and without.



Functional Levels

Another important decision revolves around the functional level of a forest and a domain. The *functional level* defines the feature set available to the organization and is driven by the operating systems installed on the domain controllers in a forest. When a new forest is deployed, the first domain controller installation will prompt the administrator to specify the forest level and the domain level. Functional levels can be increased but not decreased, so it's important to know what functional levels are available and what functional levels are appropriate for the needs of the organization. In addition, domain functional levels cannot be initialized lower than the forest functional level, and domain functional levels must be raised before the forest functional level can be raised.

Some of the important features available at each functional level include, but are not limited to, the following:

- Windows Server 2003 domain functional level
 - Renaming domain controllers
 - Use of constrained delegation
 - Selective authentication
- Windows Server 2008 domain functional level
 - DFS Replication for the SYSVOL
 - AES 128 and AES 256 support for Kerberos
 - Fine-grained password policies
- Windows Server 2003 forest functional level
 - Forest trusts
 - Domain renaming
 - Windows Server 2008 Read-only domain controllers

Make special note that even though Read-only domain controllers can only be installed on Windows Server 2008 systems, they can be deployed in an environment with a Windows Server 2003 *forest* functional level. The *domain* functional level, of course, will need to be Windows Server 2008.

For a comprehensive list of the features supported at each domain and forest functional level, refer to Understanding AD DS Functional Levels in the TechNet Library. The key point to understand here is that the ability to use the enhanced features is directly driven by the operating systems deployed on your domain controllers.

You need to be familiar with the key feature enhancements of each forest and domain functional level so that you can properly recognize when domain controllers need to be upgraded or domain and forest functional levels can be upgraded. Here are the basic rules to determine the minimum functional levels available and required:

- If a Windows 2000 Server Domain Controller is deployed in any domain, the domain functional level must be set to *Windows 2000 Native*, and the forest functional level must be set to Windows 2000 Native. None of the above features are available for use.
- If a Windows Server 2003 Domain Controller is deployed in any domain, the domain functional level should be set to *Windows Server 2003*, and the forest functional level should be set to Windows Server 2003. Fine-grained password policies are not available and neither are AES and DFS sysvol replication.
 - If it is anticipated that a Windows 2000 Server Domain Controller may be deployed, then the domain and forest functional levels must be left at Windows 2000 Native. Once the domain functional level is set to Windows Server 2003, a Windows 2000 Server Domain Controller cannot be deployed in that domain.
- If only Windows Server 2008 Domain Controllers are deployed, the domain functional level should be set to *Windows Server 2008*, and the forest functional level should be set to Windows Server 2008.
 - If it is anticipated that a Windows Server 2003 Domain Controller may be deployed, then the domain and forest functional levels must be left at Windows Server 2003. Once the domain functional level is set to Windows Server 2008, a Windows Server 2003 Domain Controller cannot be deployed in that domain.
 - The above rule related to Windows 2000 Server Domain Controllers also applies here.

Schema Modifications

Knowing when schema modifications are required or desirable is a key design skill that every enterprise administrator should be proficient with. There are basically two scenarios that may drive the need for a schema modification.

- The first is to implement customizations to Active Directory to support applications or enterprise directory services. A decision to modify the AD DS schema is a significant one and should be weighed against the needs of the organization, as well as the impact on the forest. Schema modifications are not reversible. If the schema modification is targeted at the needs of a small set of users or will be short-term in duration, a separate forest may be an appropriate solution combined with forest trusts.
- The second is to support the introduction of domain controllers running operating systems newer than the current domain and functional level will support. For example, you cannot join a Windows Server 2008 domain controller to a domain running at the Windows Server 2003 (or Windows 2000 Native) functional level.

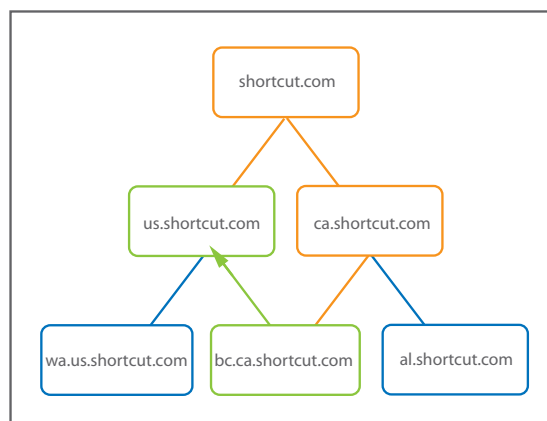
Schema modification to support a Windows Server 2008 domain controller requires two schema modification tasks:

1. Running the command `adprep /forestprep` on the domain controller holding the Schema Master operations master role.
2. Running the command `adprep /domainprep /gpprep` on the domain controller holding the Infrastructure operations master role.

In addition, if a Windows Server 2008 Read-only domain controller will be deployed, you must also run the command `adprep /rodcprep` on any domain controller in the domain.

Trusts

Intra-forest authentication is another issue to be considered in a multi-domain forest. By default, all inter-domain authentication traffic must pass up to the forest root domain and back down to the destination domain. Depending on the communication links, physical topology, and load levels of domain controllers, this can take a substantial amount of time (in terms of how long domain authentication normally takes). This intra-forest / inter-domain authentication traffic can be optimized by the creation of a *shortcut trust*. A shortcut trust is a one-way or two-way explicitly defined trust between two domains of the same forest. The creation of the shortcut trust allows the domain controllers of subject domains to communicate directly, thus bypassing the forest root domain and every parent domain in between the subject domains and the forest root domain. This image depicts a one-way trust between `bc.ca.shortcut.com` and `us.shortcut.com`. The endpoints of the shortcut trust are depicted in green, and the bypassed domains in orange.



Other trust relationships that an enterprise administrator should be familiar with include forest trusts, external trusts, and realm trusts. The *external trust* is used to establish a trust relationship between a forest and a non-forest based domain, such as would be encountered in a Windows NT domain and a *realm trust* is used to establish a trust with a Unix realm that uses Kerberos authentication. Because the forest is a security boundary, accessing resources in another forest requires a trust relationship to be designed. In the resource forest model, for example, a one-way trust is required to permit user accounts from the primary forest to access resources contained in the resource forest. A forest trust permits all domains in one forest to access all domains in another forest, so it's also important to be aware of the implications of such universal access. You can also establish forest trusts with organizations other than your own, such as partners, customers, or other organizations of interest. A common scenario is one that results from merger and acquisition activities of an organization. Forests and forest trusts can also be established for the creation of testing environments. A testing forest can be created as a clone of the production forest and established with a one-way forest trust permitting the users and resources of the production forest to access the testing forest but preventing the testing forest users or resources from accessing the production resources.

Physical Implementation

Once you've designed the forests and domains, determined the necessary functional levels, made determinations related to any schema modifications, and identified required trusts, it's time to design the physical implementation of your Active Directory service.

Site Structure

The first requirement in designing a physical implementation is identifying where AD DS sites must be established. A *site* is a logical construct that may or may not mirror physical locations and is based on network connectivity, available bandwidth, and resources (servers, workstations, users) that exist at the physical location. The purpose of the site is to establish data replication boundaries and authentication boundaries. Authentication boundaries define the collection of domain controllers that any particular computer or user account can contact for account authentication so as to preclude the transmission of authentication traffic over slow-speed WAN connections. At least one domain controller must be contained in every site. Replication boundaries isolate whether data replication between domain controllers occurs at high data-rate speeds around the clock or is restricted by time of day or bandwidth usage.

Each physical location should be evaluated based on the cost of bandwidth, the impact if the bandwidth is impeded or unavailable (i.e. the WAN link is dysfunctional or down), whether any applications installed at the location are "AD DS site aware", and the expected load of clients and machines at that physical location.


Within the IP addressing infrastructure of an organization, each IP subnet is then assigned to an AD DS site if the need for more than one site is identified. Ergo, an AD DS site contains one or more IP subnets, which may consist of one, part of one, or more than one physical location.

Replication

Once you've identified your sites, you then need to identify a replication topology to transport AD DS data between the sites. There are three approaches to replication:

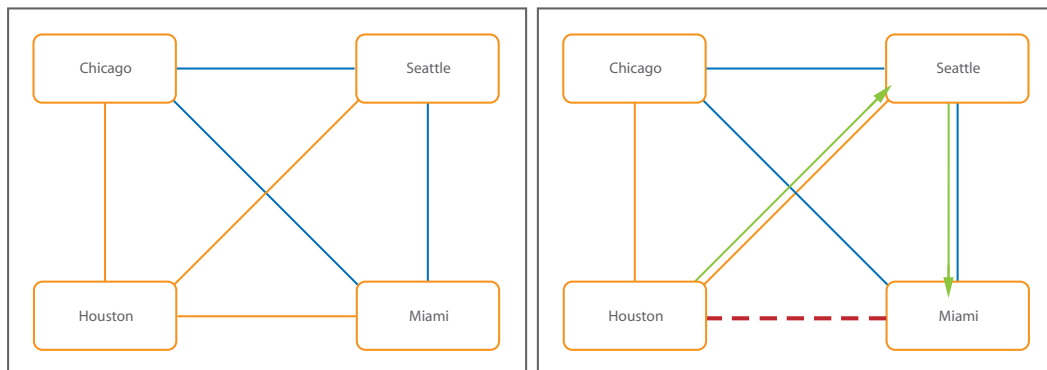
- Hub and spoke (a central office with remote sites)
- Full mesh (every site is connected with every site)
- Hybrid (a combination of these two)

The first two are exactly as described for WINS in Domain 1. The hybrid topology may consist of a full mesh central site with spokes into full mesh remote sites (which with full mesh replication may include sites that cover entire metropolitan areas) or may consist of a full mesh of hubs, and each hub has a set of individual spoke sites.

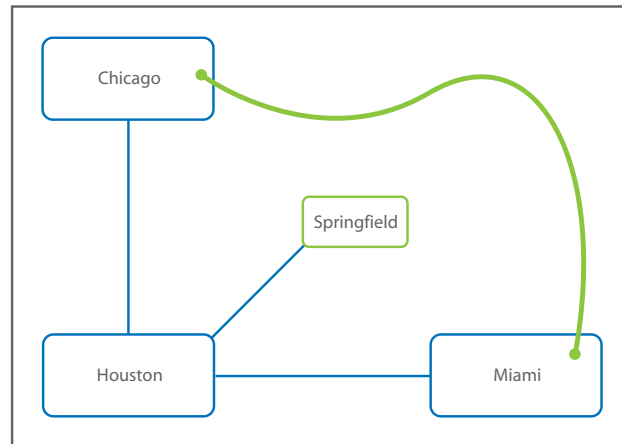


Replication between two sites is implemented by defining a site link and assigning values to the site link properties. Site link properties include the relative cost of the link (based on available bandwidth), the schedule of link availability, and how often replication should occur. These values should be set according to the specifications obtained during the site structure analysis.

To enhance replication and provide multiple pathways for replication, site replication utilizes *site link bridging*. Site link bridging connects two site links together and permits transitivity. So, in the provided full-mesh diagram, Houston is the corporate office and has direct replication links with Seattle, Chicago, and Miami (depicted in orange). Additional site links are defined between Seattle and Chicago, Chicago and Miami, and Seattle and Miami (depicted in blue). Site link bridging permits traffic to any of those three sites to transit through another site. So if the Houston-Miami WAN connection goes down, traffic can still be replicated to Miami by transiting through the Chicago link, as depicted in the diagram with the green arrows.



Site link bridging is enabled by default for all site links; however, it can be disabled to permit selective bridging. For example, imagine that the office were a small office in Springfield, Missouri instead of a large regional office in Chicago. It might not be desirable to allow a scenario in which Seattle or Miami traffic is transited through the Springfield office. In this case, you would disable the Bridge All Site Links option and manually create a site link bridge object for Seattle-Miami so that Seattle traffic can only transit through Miami and Miami traffic can only transit through Seattle if the primary link to Houston is not working.



Domain Controllers

Now that the sites and site links are designed, it's time to identify specifically where domain controllers will be deployed. For the most part this is simply a matter of following some basic "best practices" rules.

Domain controllers for the forest root should be placed in the following locations and each location should have at least two domain controllers in each data center: every primary hub site and any other location which needs to support a domain trust across a low-speed or unreliable WAN connection.

Every other domain should have at least one domain controller in each defined site. Additional domain controllers may be deployed within a site to support larger numbers of users or special location connectivity scenarios. For example, within any site the data center should have at least one domain controller, and the physical location with the most users might have a local domain controller. Any mission-critical satellite location at risk for loss of WAN connectivity would also benefit from a local domain controller in the event the WAN link fails.

For satellite locations, such as retail locations and sales offices, where physical security of the server hardware is not reliable, it may be desirable to deploy a read-only domain controller (RODC). However, deploying a read-only domain controller is done primarily to provide local authentication and AD DS data not to overcome connectivity deficiencies. The RODC requires a solid replication connection with a writable DC, preferably within the same site, but if not, in the closest site (based on site-link cost).

Global catalog (GC) servers are sometimes a confusing issue for enterprise administrators. It's important to understand the distinction between a domain controller and a global catalog. The GC is a service that resides on selected domain controllers. Its purpose is to provide a searchable data repository of every object in a multi-domain forest. This assists in the location of objects in domains other than the local domain. The question becomes: Should there be a GC on a particular domain controller or not? A few simple rules can simplify this question:

- In a single domain model—every domain controller should be a Global Catalog server.
- In a multi-domain model:
 - ▶ locations with more than 100 users should have a local GC.
 - ▶ physical locations within a site that have high-speed connections do not need to have a GC on the local domain controller. They can use the GC contained on the site's primary domain controller.
 - ▶ AD DS-aware applications may require a local GC. Some examples are Exchange, MSMQ, and DCOM.

An alternative to deploying a global catalog is the use of universal group caching, which is a site-specific setting and must be explicitly enabled. Universal group caching can be implemented on Windows Server 2003 and higher domain controllers. In most scenarios, the primary use of the global catalog is to assist in the determination of members of universal groups—recall that universal groups can have members from multiple domains. Implementing universal group caching serves this need but eliminates the replication of all of the rest of the domain objects that are carried along with a global catalog. This can be a significant reduction in unnecessary replication traffic passed across WAN connections. The non-GC domain controller must still contact a global catalog server to obtain this information, but once cached, no further contact with the global catalog server is necessary.

Finally, it is necessary to ensure the Flexible Single Master Operation (FSMO) roles are properly assigned. As a review, there are five FSMO roles in an AD DS environment:

- Schema master—responsible for managing the forest schema (which includes the domain schemas); the schema is what defines the objects contained in the directory, and the attributes of those objects. There is one schema master in the forest.
- Domain naming master—manages the naming of all domains in the forest and group memberships; must be online to add or remove domains. There is one domain naming master in the forest.
- RID (Relative ID) master—the source for all security RIDs, which are allocated to DCs to assign to new security principals; manages inter-domain moving of objects. The RID is the portion of the SID that uniquely identifies the object within the domain. There is one RID master per domain.
- PDC (Primary DC) master—processes the password changes in a domain; is queried anytime a DC fails an authentication attempt to ensure recently changed passwords are properly replicated. There is one PDC master per domain.
- Infrastructure master—maintains inter-domain identifiers (SIDs, GUIDs, and DNs), most commonly the links between users and groups. There is one Infrastructure master per domain. In a single-domain forest, this role has no work.

A detailed review of the functions of each FSMO role is available in the *Operations Masters Technical Reference*, and with an MCTS in Active Directory Configuration, you should already be intimately familiar with the functions of these roles.

As an enterprise administration, however, your focus shifts to where they should be deployed, rather than what they do. Here is some guidance to follow when designing the placement of FSMO roles.

In a single-forest/single-domain environment, every domain controller should be a GC, and all of the roles should be centrally placed on the most powerful domain controller in the forest root, although there are some considerations for splitting roles across multiple domain controllers. In a multi-domain environment, however, the placement of FSMO needs a bit more consideration.

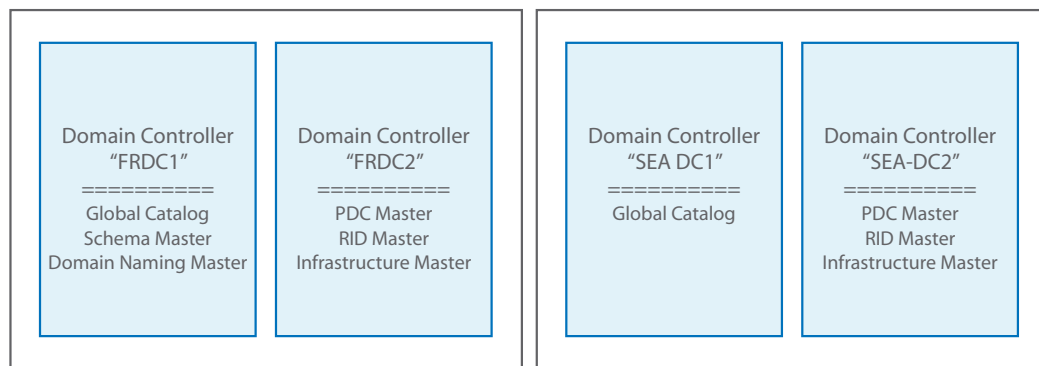
The schema master is a forest-wide role, so there is only one schema master in a forest, and initially it resides in the forest root domain (having been created on the first domain controller created in the forest). The schema master has two requirements: One is accessibility to the Enterprise Schema Administrator(s) to implement forest schema modifications; the other is reliable replication capabilities. The schema master should be placed so that both requirements are optimized. The schema master should be retained in the forest root domain, and most likely physically located in the primary data center.

The domain naming master is the other forest-wide role and its primary requirement is reliable replication to other domains. It should also be centrally located, initially exist in the forest root domain (created at the same time as the schema master), and should be retained in the forest root domain.

In addition the schema master and domain naming master should reside on a DC that is also a global catalog server.

The other three roles exist in each domain contained in the forest, and should be located on the same domain controller which is not a global catalog server. For each domain, the PDC master role should be assigned to a domain controller which is located in the location supporting the largest number of clients. The RID master's only requirement is placement near another DC replication partner. The Infrastructure master needs direct access to a global catalog server but should not be installed on a GC.

Thus, combining all of those rules, we see here a diagram showing the deployment of the domain controllers in the forest root domain in the primary data center (on the left), and the domain controllers in a standard domain in that forest.



Printer Location Policies

Finally, as an enterprise administrator designing an AD DS implementation, you need to know a bit about publishing printers and printer locations via AD so that users can easily find those resources.

To facilitate this, a standard location schema has been developed to identify what printers in an enterprise are available to a user at a particular location. The naming syntax is a string of location names separated by slashes. Each location name must be no more than 32 characters, the entire location string must be no more than 260 characters, and the maximum number of location levels is limited to 256.

The schema is a hierarchical definition of locations from largest to smallest, so for a global organization the hierarchy would look like this:

[Continent] / [Country] / [State or Prov] / [City] / [Dept or Bldg] / [Floor]

If an organization were not global in scope, this might be simplified to something like:

[Country] / [State or Prov] / [City] / [Dept or Bldg] / [Floor]

And a domestic company in a single building in a single city might be as simple as:

[Dept] / [Floor] / [Room or Cube or Desk]

Then each printer's properties are configured with the defined location, each location is associated with an IP subnet, the printer is published to Active Directory, and the group policy setting "Pre-populate printer search location text" (aka Physical Location Tracking) must be enabled.

Administration

With the environment now designed, we can turn to the question of who will perform the administrative duties to keep the environment running well. As discussed previously, there are two major scopes of administrative activities: service management and data management. Those who perform service management tasks are concerned with replication, backups, server security, functional levels, operations masters, and managing domain controllers. Data management tasks focus on computer accounts, user accounts, groups and group membership, application attributes, and other resources.

How an organization delegates administrative tasks in Active Directory is generally dependent upon its size and management model. A small organization (a few servers and a few dozen desktops) typically has a centralized management structure with a very small I.T. staff. There's not a lot of delegation of tasks in such an organization; usually every I.T. staff member is a Domain Administrator and responsible for all levels of functionality. A medium size organization (a few dozen servers and a few hundred desktops) might have centralized management or decentralized management. Business units might have authority to create and remove user accounts, and this task may be delegated outside of the I.T. department. There may also be levels of responsibility within the I.T. department separating user and desktop administration activities from the data center. It may be desirable to delegate authority for functionality within the I.T. department. Large organizations, usually with several dozen servers and several hundred desktops, rarely use a centralized management structure, making delegation a requirement. Thus it becomes necessary for the enterprise administrator to understand how to design a delegation model for Active Directory.

Delegation must first adhere to the basic security principle of “least privilege.” An administrator should be given no more authority within AD DS administration than is necessary to perform the tasks assigned. A formal job description with a well written scope of authority can be a very useful tool for determining where these privileges need or do not need to be granted. The purpose of delegating administrative tasks is to increase efficiency and reduce costs. It simply makes no sense to assign a task to another person if the end result is higher cost or it takes longer to complete a task. Also, do not forget the use of automation and AD DS as a source for increasing efficiency.

Implementation of delegation is best handled through the effective use of security groups. Be familiar with the built-in domain local security groups (e.g. Server Operators, Backup Operators) and the scope of authority granted to those groups. Do not create custom groups when membership in a built-in group will achieve the objective. The standard principle of group assignments also applies for AD delegation: Assign users to global groups; place global groups in domain local groups; assign rights to the domain local group. You should also know when to use universal groups and when not to use universal groups. Universal groups are used to provide permissions to objects in any domain or in other forests where forest trust relationships exist, but because universal groups require a search of the global catalog, they should only be used when specifically needed to provide cross-domain or cross-forest access. Universal groups can have individual accounts as members but should be restricted to Domain Local groups as members. Finally, take advantage of management roles to group administrative rights. For more information on management roles and AD administration delegation, review the Best Practices for Delegating Active Directory Administration white paper in the Technet Library.

Auditing has been enhanced significantly in Windows Server 2008 and you should be familiar with the new auditing capabilities so that you can properly design an auditing strategy for your AD environment. Auditing has more granular control than previous versions and now provides the ability to audit access, changes, and replication as individual targets. Each target is further subdivided into two levels: normal and detailed. An example of using this more granular capability is to configure auditing for normal on access and replication but detailed on changes. Changes in auditing include log entries for object creation, movement, undeletion, and modification. In addition to having a design to implement auditing, you should also have a plan to actually use the collected audit data in some type of recurring review process. An audit data collection reporting continued security issues is useless if nobody is actually looking at the collected data.

Your design initiative must also include the domain organizational structure—the creation of organizational units (OUs). Organizational units should be created to reflect the business needs of the organization and the Active Directory service and may not necessarily reflect the divisional or departmental structure of the business. A key concern when designing OUs is support for delegation of administration. The OU is typically the point at which task authority is delegated. Also, keep in mind that user accounts and computer accounts can only be members of a single OU. Finally, the OU is the primary link point for Group Policy objects and deploying Group Policy should also be considered when designing the organizational structure.

Group Policy

Group Policy is the methodology by which common configuration settings are distributed to, and imposed upon, computers and users who are members of a domain. Group policy is implemented through Group Policy Objects (GPOs) which are linked to one or more of four levels of the AD organizational structure: the individual domain-member computer, the OU, the site, and the domain. Strictly speaking, policies applied at the computer level are called “Local Group Policy” or just “Local Policy”. Other than understanding the impact of Local Policy on the application of GPOs, Local Policy is beyond the scope of this discussion.

GPOs should be created with specific purposes and appropriately named to reflect that purpose. The combination of settings in a GPO to achieve multiple purposes is discouraged, as this complicates administration and diagnosing policy malfunctions. Generally speaking, you should not modify the settings in the Default Domain Policy or the Default Domain Controller Policy; rather, you should create a new GPO containing the custom settings.

A new feature in Windows Server 2008 is Starter GPOs. This feature allows you to define baseline templates that can be used as a starting point to make a customized GPO for use in your environment. The Starter GPOs are found in the Starter GPOs node of the domain, and this node is initially empty. It must be populated with Starter GPO templates created by the organization or imported from another source, such as community content. Part of your planning activities should include the definition of baseline policy settings for your sites and domains, and implement those baseline settings in a starter GPO. When a new site or domain is created, an initial (custom) GPO for that site or domain would be built from the defined template, thus ensuring that all planned policies are properly implemented from the start and eliminating errors that might result from administrators missing or misconfiguring policies that are otherwise only documented only on paper.

Policy Prioritization

It is important when designing a policy structure for your organization that you are aware of the prioritization of the application of policies and the options available for modifying that behavior. Policies are functionally applied by starting at the local machine (using Local Policy) and then subsequently applying policies at each additional level in the order of site, domain, and OU. What this means is that policies at the OU level have the highest priority for configuration, and policies at the local (computer) level have the lowest priority. Domain policies supersede site policies. A useful acronym for remembering the order of *application* of policies is LSDO (local, site, domain, OU). Within a single level, GPOs are applied in order of listing in the Group Policy Management Console (GPMC).

The results of the application of policy can be modified through four methodologies: inheritance blocking, enforcement, filtering, and loopback policy. As a designer of Group Policy you should be familiar with all four tools and when it's appropriate to use them. Generally speaking you should only use them when absolutely necessary. Given the option between creating additional GPOs and employing a modification tool, the choice should always be made to avoid the use of the modification tools. Modification tools add complexity to the group policy design and are difficult to keep track of because their implementation is buried inside the policy object, whereas an appropriately named and linked GPO is readily apparent as to its purpose and impact.

Inheritance blocking is configured on a domain or OU and blocks *all* policy settings from earlier applied policies effectively isolating that domain or OU as its own independent policy island. This is a useful tool for creating OUs for testing of individual policies and policy interaction or for isolating users or computers from normal production policies. Because the impact is so significant, it is almost never applied at the domain level and only rarely applied at the OU level, and even then, with a very specifically defined objective to achieve.

Enforcement allows a specific GPO to prevent overriding of its settings by a subsequent policy object and also overrides any inheritance blocking configured on a domain or OU. Enforcement is typically employed at the site or domain level, where certain settings must apply to all resources in that site or domain level without exception. One place where enforcement is particularly useful is in ensuring delegated administrators do not override policies set at higher levels of the hierarchy.

Filtering can be applied for either of two purposes: To apply a GPO to specifically identified groups or to exclude application of a policy to specifically identified groups. Groups can be defined either by AD security groups created in the container being filtered or by WMI object collections, such as installed operating system or system type (workstation, server, domain controller). Common uses of filtering are exclusion of Domain Administrators from the application of user policies, such as restricting access to Microsoft Update or inclusion of certain classes of computers, such as applying a desktop or application restriction to servers and domain controllers in order to enhance security. Generally, though, the appearance of the requirement to use filtering is an indication that the OU structure may not be granular enough for the needs of the organization. In many cases, a more appropriate solution is to create additional levels in the OU structure and moving objects to more specifically defined OUs.

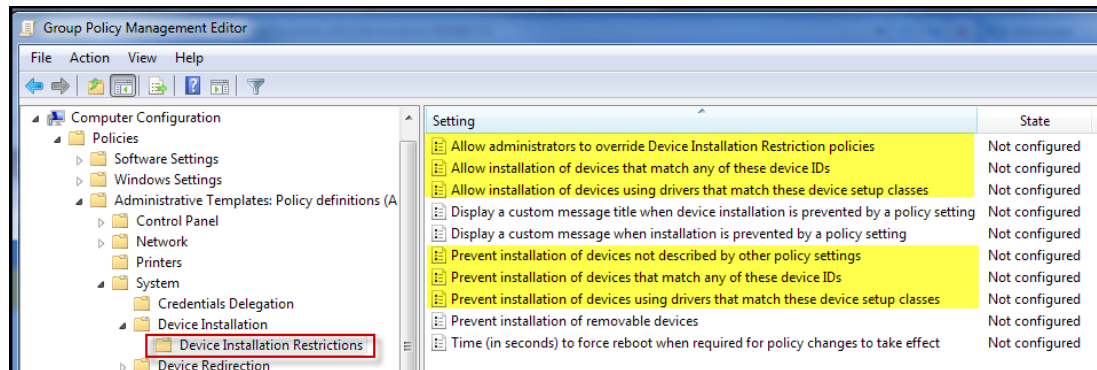
Loopback processing allows the imposition of user configuration policies based on GPOs that are targeted at computers. This tool can be useful for locking down desktop settings on shared-use or publicly-available computers, such as those located in conference rooms, kiosks, classrooms, and reception areas. Implementation of loopback processing should be combined with dedicated OUs to house such systems and separate GPOs targeted at only those OUs.

Device Installation Management

In addition to controlling user and computer settings, Group Policies can also be used to control authorizations for installation of devices. They can be used to permit the installation of only approved devices, to prevent the installation of expressly prohibited devices, or to prevent the installation of any device by non-administrators. GPOs can also be used to restrict the ability to read from or write to removable devices or removable media, for example restricting the use of the DVD-write capabilities of an installed DVD-RW drive, limiting the use of USB flash drives as read-only devices, or limiting external USB hard drives for system backups as write-only drives (thus requiring administrator assistance and approval to invoke a system-level restore).

In addition to the standard device policies that have existed, Windows Server 2008 introduces six new policy settings that can be applied to Windows Vista and newer systems. You should be familiar with the existence and potential uses of these policies:

- Prevent installation of devices not described by other policy settings
- Allow Administrators to override device installation restriction policies
- Prevent installation of devices that match any of these Device IDs
- Prevent installation of devices using drivers that match these device setup classes
- Allow installation of devices that match any of these Device IDs
- Allow installation of devices using drivers that match these device setup classes



A functional understanding of Hardware IDs, Compatible IDs, and GUIDs is necessary in order to properly design the use of group policy for device installation management. Hardware IDs provide matches between a device and the driver package required for that device. The most detailed Hardware ID is called the Device ID and matches the exact make, model, and version of a device. Compatible IDs are used to match a driver package to a device when the driver is unable to locate the actual Device ID or Hardware ID it applies to. GUIDs are used to identify a device setup class which is used to group devices according to how they are installed and configured by the device driver package.

Authentication and Authorization

The final step in designing the core components of an Active Directory implementation is understanding how authentication and authorization are implemented, as well as the distinctions between the two. Authentication and authorization are often confused, and the terms are used interchangeably, quite inappropriately so.

Authentication is the process of identifying an individual or resource to the authentication entity. The sole purpose of authentication is to establish that a user is who the user claims to be or that a computer is the computer it claims to be. *Authorization* is the process by which an authenticated computer or user is granted access to see or use a specified resource.

Authentication is implemented through the use of one of three categories of authentication *factors*: The computer or user knows something that is presumed unknown by other entities, such as a password or a Personal Identification Number. The computer or user has something that identifies it, such as a card, key, or digital certificate. The computer or user is something identified by an attribute unique to that entity, typically through biometrics such as a fingerprint, voice print, or retinal scan.

In the Microsoft Windows environment, authentication is implemented through the use of passwords or smart cards; authorization is implemented through the use of ACLs, group policies, and user rights policies. *Multifactor authentication* is a technique which uses at least two of these three categories of factors. For example, a keycard and a thumbprint to access a data center or a debit card and PIN to withdraw money from an ATM. Multifactor authentication does not include redundant imposition of the same category of factor, even if the entity or values used to achieve the authentication are different. For example, having two different keys (building and room), using two different fingerprints (thumb and forefinger) on two different biometric locks, or having two different accounts and passwords is not multifactor authentication.

Related to multifactor authentication is *multifactor authorization*, which is even more confused with multifactor authentication, than are their simpler parents. Multifactor authorization is the use of two distinct authenticated entities to achieve access to a resource or task. For example, requiring the supervisor to enter a PIN code to authorize a return or transaction cancellation at a retail cash register is a use of multifactor authorization. Turning keys in the 1980s to launch a nuclear missile was one of the earliest forms of multifactor authorization, as each officer had to independently possess (authentication) and act (authorization) in order for the task (launch) to initiate.

A significant enhancement to the authentication capabilities in Windows Server 2008 is the introduction of fine-grained password policies. Historically, the password policies in an Active Directory environment could only be applied at the domain level, and only one policy could exist. All user accounts, regardless of type or purpose, had to have the same policies. This resulted in a reluctance of organizations to impose strong password requirements on administrator or service accounts because it impacted the standard user's ability to generate passwords they could remember, or worse, the practice of writing them on sticky notes attached to the bottom of the monitor or not implementing frequent password resets because of the overhead imposed on service accounts.

With Windows Server 2008 the ability to create Password Settings Objects (PSOs) is introduced. PSOs are applied to global security groups (rather than containers of the AD hierarchy). There is no limit on the number of different PSOs you can create, but typically this number would only be a few to several at most. One common scenario involves three PSOs separated by usage for standard users, administrative users, and service accounts.

Standard users typically are characterized by challenges selecting and remembering passwords and volatility in their authorizations as a result of task requirements, team membership changes, promotions, transfers, and terminations. A PSO minimizing or eliminating the complexity requirements (thus eliminating the need to create and remember complex passwords, as well as the undesirable practice of writing them on the ubiquitous sticky notes), combined with very short aging and a long history list can be beneficial. If a user is changing a simple password every 60-90 days, this significantly reduces the likelihood of forgetting the password or password compromise, as well as mitigates any potential long-term impact as a result of a password compromise. The short aging period also assists in ensuring access to the environment is eventually blocked if the account is not appropriately disabled upon the user's departure.

Administrative users are considered to be a bit more "password savvy," but also have higher levels of access to sensitive resources. Password complexity is desirable to mitigate the risk of password attacks, and super short aging with history can assist in mitigating issues with password compromise or changed in the administrator's employment status or duty assignments.

Service accounts should be heavily protected as they generally carry access to the highest levels of the organizations application infrastructure, although they should also be configured with "least privilege" principles. Service account passwords, as a collection, should only be known by a very few persons, and individual service account passwords should only be divulged to individuals directly responsible for managing those service configurations on a "need to know" basis. A service account should be highly complex and sufficient to virtually eliminate the risk of password attacks, which can be combined with longer-term aging (minimizing the need to create such super-complex passwords). Service account passwords can be managed with an annual review and password reset event and should also be changed anytime administrative staff reassignments occur.

Domain 3 – Designing Support Identity and Access Management Components

Objectives

Designing support identity and access management components is about building, maintaining, and enhancing the support infrastructure that establishes who is authorized to access your network. Whereas Domain 2 focuses on the core process of designing the basic Active Directory services, Domain 3 focuses on growing beyond the core AD DS service.

The specific objectives covered in this domain are:

- Plan for domain or forest migration, upgrade, and restructuring
- Plan for interoperability
- Design the branch office deployment
- Design and implement a public key infrastructure

Updating Active Directory

Since the creation of Active Directory nine years ago, the question of converting the existing directory structure to a newer structure has been at the forefront of the enterprise administrator's list of responsibilities. Not much has changed since the introduction of Windows 2000, though, in terms of the methodologies—they're still pretty much the same:

- Create a new domain environment and migrate the old to the new, known as *Domain Migration*.
- Upgrade the existing domain controllers to a new operating system, or install new domain controllers running the new operating system, known as *Domain Upgrade*.
- A combination of the above two, generally involving upgrading immediately to get benefits of the functionality, and then migrate to a cleaner (new) environment to get rid of the dead weight in the old, known as *Upgrade-Then-Restructure*. (Interestingly enough, nobody has ever recommended cleaning the old environment and then upgrading.)

The upgrade-then-restructure methodology is best implemented as follows:

1. Run `adprep` to upgrade the existing forest and domain.
2. Install one or more new domain controllers using the new operating system on new systems.
3. Demote the older domain controllers and reinstall with the new operating system (or decommission them).
4. Restructure or migrate the domain (if necessary), or simply clean up the existing environment.

A key question in this process revolves around how necessary the *migration* or *restructuring* part of this process actually is. For many organizations, it may not be necessary to selectively migrate AD DS objects or engage in any restructuring of the existing domains. In addition, significant enhancements have been built into recent versions of AD DS, and situations in the past that might have required a domain migration or restructuring can now be handled by simple domain administration tasks. Recognizing whether you need to perform selective migration or restructuring is the first step in updating your Active Directory environment, and determining whether to conduct that migration/restructuring before upgrading, or after, is a key implementation decision. Being able to plan using any of these methodologies is a requirement for the certification examination.

Also related to this scenario is recognizing incompatibilities between Windows Server 2008 domain controllers and downgraded functional levels. Windows Server 2008 domain controllers can only be deployed in domains and forests operating at the Windows 2000 or Windows Server 2003 functional levels. If the environment is still operating in Windows 2000 Mixed or Windows Server 2003 Interim mode because of the presence of Windows NT 4.0 domain controllers, it will first be necessary to retire or upgrade the NT4 DCs so that the functional levels can be raised to a native mode functional level. To deploy a Windows Server 2008 domain controller into an existing domain, the following requirements must be met:

1. All Windows 2000 Server domain controllers must have Service Pack 4 installed, as well as the KB265089 update.
2. All Windows Server 2003 domain controllers must have Service Pack 1 installed or be Windows Server 2003 R2 servers.

Note: While SP1/R2 are the minimum requirements to deploy a Win2008 domain controller, Microsoft has not released any updates since April, 2009, for those platforms, as they expired Mainstream Support at that time. So while the official “answer” for examination purposes is SP1 or R2, the functional requirement as of today is that Windows Server 2003 domain controllers need to have Service Pack 2 installed.

3. Schema modification must be implemented to support deployment of a Windows Server 2008 domain controller in a Windows Server 2003 or Windows 2000 Native environment. This requires you to:
 1. Run the command **adprep /forestprep** on the domain controller holding the Schema Master role. This command must be executed by a user who is a member of Schema Admins, Domain Admins, and Enterprise Admins.
 2. Run the command **adprep /rodcprep** on any domain controller in the forest if a Windows Server 2008 Read-only domain controller will be deployed. Since you’re already running forestprep on the Schema Master, it is recommended to run rodcprep as well, even if there are no immediate plans to deploy a read-only domain controller. Having properly prepared the forest for such an eventuality will be a benefit to future enterprise or domain administrators who may or may not be aware of the status of the schema modifications.
 3. Run the command **adprep /domainprep /gpprep** on the domain controller holding the Infrastructure operations master role in each domain that will host a Windows Server 2008 domain controller. This command can be run by any user who is a member of the Domain Admins group for the domain in which the schema modification is being implemented.

When engaging in restructuring or selective migration, it will be useful to be familiar with the use of the Active Directory Migration Tool (ADMT) v3.1. The ADMT v3.1 Guide, “Migrating and Restructuring Active Directory Domains” can be downloaded from the Microsoft Download Center. This is a comprehensive, 227-page manual that covers all aspects of domain migration and restructuring.

Inter-Forest Trusts

When migrating or restructuring forests, it may be necessary to provide cross-forest authentication between the old and the new trust to allow existing users access to resources as they are migrated into the new forest. This is achieved with the forest trust and can be one-way or two-way. There are several other scenarios in which it may be necessary to establish authentication between forests. In the resource forest model, for example, a one-way trust is required to permit user accounts from the primary forest to access resources contained in the resource forest. You can also establish forest trusts with organizations other than your own, such as partners, customers, or other organizations of interest. A common scenario is one that results from merger and acquisition activities of an organization. Forests and forest trusts can also be established for the creation of testing environments. A testing forest can be created as a clone of the production forest and established with a one-way forest trust permitting the users and resources of the production forest to access the testing forest, but preventing the testing forest users or resources from accessing the production resources.

A forest trust permits all domains in one forest to access all domains in another forest, so it's also important to be aware of the implications of such universal access. As a result, it may be desirable to implement selective authentication within a forest trust relationship. Selective authentication allows for the identification of specific resources in the trusting forest that can be accessed only by specific users or groups in the trusted forest.

Another security consideration when implementing forest trusts is the use of the SIDHistory attribute. When a user or group account is migrated to a new domain, it is assigned a new Security Identifier (SID) in the new domain. In order to facilitate continued access to resources in the original domain that have not yet been migrated, the SIDHistory attribute allows the retention of that object's previous SIDs. SIDHistory filtering prevents the ability to use the SIDHistory attribute and thus, blocks access to resources in the original domain. SIDHistory filtering is enabled by default on any trust created from a Windows Server 2008 computer; therefore, SIDHistory filtering should be disabled to support forest or domain migrations in such scenarios and re-enabled when the migration of all resources is completed.

Related to SIDHistory filtering is the Quarantine option. Quarantine allows the trusting domain to restrict the use of the SIDHistory attribute to only those SIDs that originate directly from the trusted domain. This effectively breaks the transitive nature of forest trusts but allows your original (trusting) domain to be sure that only those accounts from the new (trusted) domain are able to access resources.

To determine the current settings for SID History on the new domain, run this command from a domain controller in the original domain using a domain administrator account:

```
Netdom trust OriginalDomainName /domain:NewDomainName /
EnableSIDHistory /userD:newDomainAdministratorAcct /passwordD:
newDomainAdminPwd
```

If SIDHistory is not enabled, you can enable SIDHistory by running this command on a domain controller in the original domain using a domain administrator account:

```
Netdom trust OriginalDomainName /domain:NewDomainName /
EnableSIDHistory:YES /userD:newDomainAdministratorAcct /passwordO:
newDomainAdminPwd
```

To determine the current settings for Quarantine on the original domain, run this command from a domain controller in the original domain using a domain administrator account:

```
Netdom trust OriginalDomainName /domain:NewDomainName /Quarantine /
userD:newDomainAdministratorAcct /passwordD:newDomainAdminPwd
```


If Quarantine is not enabled and you wish to enable it, run this command on a domain controller in the original domain using a domain administrator account:

```
Netdom trust OriginalDomainName /domain:NewDomainName /Quarantine:YES  
/userD:newDomainAdministratorAcct /passwordD:newDomainAdminPwd
```

Other trust relationships that an enterprise administrator should be familiar with are external trusts and realm trusts. The *external trust* is used to establish a trust relationship between a forest and a non-forest based domain, such as would be encountered in a Windows NT domain. A *realm trust* is used to establish a trust with a Unix realm that uses Kerberos authentication.

The external trust is also a solution to addressing the presence of NT4 domain controllers which may be impeding the deployment of Windows Server 2008 domain controllers in a domain still at the Windows 2000 Mixed or Windows Server 2003 Interim functional levels. The NT4 resources can be retained in their existing domain, the Active Directory resources can be migrated into a new forest, and an external trust set up between the AD DS forest and the NT4 domain.

In addition to forest trusts, Active Directory Federated Services (AD FS) is another tool that can be deployed to facilitate inter-organizational authentication and authorization. AD FS was originally provided in Windows Server 2003 R2, but the version provided in Windows Server 2008 provides enhanced integration with Microsoft Office SharePoint Server 2007 (a product particularly involved with extranet functionality) and AD Rights Management Services (AD RMS), which helps organizations to protect intellectual property and digital content from inadvertent modification as well as intentional theft. MOSS 2007 and AD RMS are discussed further in Domain 4.

The principle of AD FS is to keep account management within the organization where the user exists but to leverage the existence of those accounts and establish a trust relationship with the organization to provide authentication and authorization to use resources within the trusting organization. The principle is very similar to a forest trust with the resource forest trust model except in this case it involves distinct organizational entities referred to as the resource partner and the account partner, and it is targeted at browser-based clients and web-published resources. AD FS is based on the WS-* architecture and consists of five distinct components:

- **The Resource Federation Server (FS-R)** — AD FS installed on Windows Server 2008 in the resource partner's network that provides authentication for accounts contained in the account partner's LDAP-compliant directory service.
- **The Federation Service Proxy** — a role installed on Internet-facing connections in the resource partner's network that acts as an intermediary between the external account authentication requests from the account partner, and the resource partner's internal (secured) Federation Server.
- **The Account Federation Server (FS-A)** — AD FS installed on Windows Server 2008 in the account partner's network that issues security tokens to the account objects of the account partner and which are used to authenticate with the FS-R. This is achieved by deploying an AD repository on the Account Federation server, typically AD Lightweight Directory Services (AD LDS) or Active Directory Application Mode (ADAM).
- **The AD FS Web Agent** — a software package installed on the local web server that recognizes and uses security tokens issued by an Account Federation Server and authorized by the Resource Federation Server.
- **AD FS-enabled Web Server** — a web server with an AD FS web agent installed and an established relationship to a Federation Server to permit authentication based on the security tokens issued by the Account Federation server.

Interoperability

In today's highly heterogeneous network environments, interoperability with other vendor's operating systems and applications is a necessary capability. Windows Server 2008 provides several tools and components to support interoperability. As an enterprise administrator you need to be familiar with the uses of each of these:

- Identity Lifecycle Manager (ILM) / Forefront Identity Manager (FIM)
- Subsystem for UNIX-Based Applications (SUA)
- Identity Management for UNIX: Password Synchronization
- Identity Management for UNIX: Server for Network Information Services (NIS)
- Services for Network File System (NFS)

Microsoft Identity Lifecycle Manager (ILM) 2007 Feature Pack 1 is a resource for synchronizing and managing accounts, certificates, and smart cards among several disparate environments which maintain authentication credentials. ILM can be used to interface identity data between other AD DS forests, Active Directory Application Mode on Windows Server 2003 R2, Windows NT4 domains, all contemporary versions of Microsoft Exchange (v5.5 and newer) and Microsoft SQL Server (v7 and newer), non-AD directory services such as IBM Tivoli, Novell eDirectory, and Sun Directory Server, and non-MS applications such as Lotus Notes, Oracle, DB2, and SAP.

ILM 2007 is an enhancement of Microsoft Identity Integration Server 2003 (MIIS 2003) by adding a new feature set to the product known as Microsoft Certificate Lifecycle Manager 2007 (CLM 2007) and is currently undergoing rebranding again for the next revision, which will be known as Forefront Identity Manager 2010.

The **Subsystem for Unix-Based Applications (SUA)** provides access to POSIX-compliant functionality in the Windows environment. It has existed through several versions now and can be installed on both desktop and server operating systems. SUA provides shells, shell scripting, scripting utilities, software development kits, compilers and compilation tools, commands, utilities, and support for job control and case-sensitive filenames. It's installed as a feature of the operating system which installs a help file and a Start Menu pointer to the latest download of a platform-specific bundle of content.

Identity Management for UNIX: Password Synchronization is an implementation of the specification outlined in RFC 2307, "An Approach for Using LDAP as a Network Information Service", and supports account and password synchronization between a Windows Server 2008 AD DS environment and several versions of mainstream UNIX/Linux. The functionality must be installed on all domain controllers and requires that the password policies are consistent on both Windows and Unix systems—such as minimum length, expiration, history retention, and complexity requirements. The functionality restricts replication of sensitive accounts, such as server and domain administrator accounts, so it's an excellent solution for synchronizing user accounts in a heterogeneous environment. UNIX systems that are supported by this functionality include:

- HP-UX 11i v1
- AIX v5L 5.2 and 5.3
- Novell SUSE Linux Enterprise Server 10
- Red Hat Enterprise Linux 4 Server
- Solaris 9 and 10 (SPARC)
- Solaris 8 (x86 and SPARC)

Network Information Services (NIS) is a legacy UNIX client-server directory and authentication service developed by Sun Microsystems and specified in RFC 1302 in February, 1992. It is, for all practical purposes, the oldest known network directory service implementation. Essentially it consists of human-readable text files in a centrally accessible network-based location that individual client machines can access to establish authentication of a local user or obtain identification information about network resources. It is deployed with a master server and one or more slave servers, quite similar in architecture and replication methodology to the Windows NT4 PDC/BDC environment. Today, NIS has been generally replaced by other LDAP-based directory services.

However, NIS-based environments do still exist, and the **Identity Management for UNIX: Server for Network Information Services (NIS)** is designed to support the deployment of a Master NIS Server on a Windows Server 2008 platform with one or more subordinate NIS servers on Windows domain controllers or UNIX systems. The primary objective of this tool is to facilitate inclusion of an existing UNIX NIS infrastructure into a Windows-based AD DS environment and manage all identity information in the Active Directory environment, rather than on the UNIX systems.

Also related to interoperability with UNIX networks is **Services for Network File System (NFS)**. NFS is also a Sun-developed technology, which quite often was deployed hand-in-hand with NIS to provide file-sharing services among UNIX systems. In the 1990s a technology known as SAMBA was developed to support the Windows Server Message Block (SMB) protocol on UNIX/Linux systems and generally contributed to the demise of NFS in heterogeneous environments. Both NIS and NFS are exceptionally tedious to configure and maintain and are generally only found, today, in UNIX environments that do not have Windows system integration. Services for NFS also require deployment of Identity Management for UNIX.

Branch Office Deployment

One of the characteristics that almost certainly separates the enterprise from a simple business entity is the presence of the “branch office.” The “branch office” may be as simple as a retail outlet with a cash register and manager’s PC; it may be a full business unit or division of the organization with dozens or hundreds of people; it might even be an entire regional headquarters of a global conglomerate. Regardless of the size or functionality of the “branch office,” the role of the enterprise administrator is essentially the same—establish and maintain network connectivity and application functionality with minimal or no downtime at the lowest possible cost.

When deploying to a branch office, there are four key areas that an enterprise administrator should focus efforts on: organization, infrastructure, security, and administration.

Organization

Perhaps the first decision that must be made when contemplating the creation of a branch office, or incorporating a branch office into an existing environment, is how that branch office will plug into the organization of the existing environment:

- Does the branch office require classification as an Active Directory “site”?
- Does the branch office have administrative requirements that necessitate the creation of a separate domain?
- Does the branch office have security requirements that mandate the creation of a separate forest?

Generally speaking, a typical deployment strategy will identify a branch office as a separate Organizational Unit (OU) in an existing site, or a new site (if a separate site definition is warranted), unless other considerations indicate an alternate strategy is preferable. A subdomain, domain, or forest are options to consider but should only be implemented when specific administrative or security requirements require it.

Infrastructure

The network infrastructure that a branch office requires will vary with the size, purpose, and connectivity available to that office. As an enterprise administrator you must be familiar with the various network infrastructure services available and when it is appropriate to implement or not implement those services at a branch office.

We can divide these services into two lists of services or roles. The primary list would include critical services such as domain controllers, global catalog servers, file and print servers, DHCP servers, DNS servers, and firewall/routing services. The secondary list would include advanced services, not typically needed at smaller offices, but perhaps appropriate for business units or regional/national offices, such as Distributed File System (DFS) services, Windows Deployment Services (WDS), Windows Server Update Services (WSUS), and for branch offices which are domain entities, possibly FSMO roles.

Let's evaluate these infrastructure services and when they might be beneficial for deployment in a branch office.

File and Print Services

Almost every branch office has, as its first infrastructure service need, support for file storage and sharing services and print services. Traditionally this has been achieved by deploying File & Print Servers at branch offices because the latency of WAN connections simply does not provide the level of service needed for everyday file and print operations.

New services in Windows Server 2008 R2, which is not material contained on the 70-647 exam (at the time of this writing), include the new BranchCache feature, which is available in both a server-based mode as well as a peer-based mode. BranchCache is likely to have a significant impact on the need to deploy File and Print Services in branch offices in the future.

Domain Controller/Global Catalog

For a branch office that requires minimal latency in authentication response or has connectivity sensitivities, such as when disconnection of the WAN simply cannot be tolerated, deployment of a branch office domain controller may be appropriate. An added advantage of the additional DC is enhancement of the domain's fault tolerance for failure of any DC; but disadvantages include the hardware and licensing costs of the server, the replication traffic overhead on the WAN connection, administrative overhead, and security concerns related to the physical accessibility of the domain controller in a typical branch office scenario. A domain controller is recommended when a branch office has more than 100 users or is using an application that is dependent on AD DS replication. Additionally, a global catalog should be deployed when the domain controller is not enabled for universal group membership caching (which requires a Windows Server 2008 domain controller).

Flexible Single Master Operations (FSMO) Roles

When a branch office represents an entire domain, the branch office should have two domain controllers deployed and the operations master roles appropriately assigned, as depicted in Domain #2 for a domain deployed in a multi-domain forest. However, another alternative which should be considered is placing the second domain controller with the FSMO roles at the central site and a simple DC/GC server at the branch office. This has the added advantage of offsite fault-tolerance for the AD database and protection of the FSMO roles by locating them at a central, secured facility.

Read-Only Domain Controller (RODC)

Another service option variant for domain controllers is to deploy the DC/GC services on a Read-Only Domain Controller. The RODC does not cache passwords, though can be enabled to selectively cache credentials using a Password Replication Policy where such functionality is beneficial. Active Directory database object attributes can be filtered to remove sensitive data from the branch office systems. Replication to an RODC is one-way from another Windows Server 2008 domain controller in an adjoining site.

Domain Name Service (DNS)

Any time a domain controller is deployed, DNS should also be deployed. In addition, you may also want to deploy a DNS-only server to a branch office to facilitate local name resolution if significant off-site network activity is present. The DNS server can also enhance accessibility to resolution of SRV records, which are necessary to locate domain controllers for domain authentication. Depending on the size and needs of the branch office, you should also consider the need for and benefits of secondary zones, forwarders, conditional forwarding, subdomain delegations, and stub zones to support subdomain delegation.

Dynamic Host Configuration Protocol (DHCP)

Deploying a DHCP server to a branch office eliminates the need for a DHCP Relay Agent, eliminates the DHCP traffic transmitted across the WAN, and allows for the customization of branch office scope attributes best suited for the branch office systems.

Routing and Remote Access Service (RRAS)

For a small branch office, RRAS can be a viable alternative to the implementation of a hardware router, though it should be noted that this is rarely done in actual practice, as branch office routers are exceptionally less expensive than Windows Server licenses and significantly easier to administer. For the Enterprise Administrator examination, however, you have to be willing to suspend practicality and accept that a Windows Server 2008 system using RRAS can be deployed to act as a branch office router. Other touted reasons for using RRAS as a branch office router include support for NAT, DHCP Relay, and VPN connectivity (though VPN connections are rarely routed through a branch office to a central office. Generally they go from the central office to the branch office).

Server Core

In addition to distribution of FSMO roles to a central site domain controller or the use of a read-only domain controller, branch office infrastructure functionality for AD DS, DHCP, DNS, and File & Print Services can be deployed on a Windows Server 2008 Server Core implementation, significantly reducing the system resource footprint of the installation and enhancing the security footprint of the branch office server.

Hyper-V

An up-and-coming service finding use in branch office scenarios is a Windows Server 2008 Hyper-V server. A nominally equipped Hyper-V server can support a half-dozen or so Standard Edition server images, providing all types of server-based functionality to a branch office at a fraction of the cost that would exist if several physical servers were deployed. Additionally, it eliminates the all-too-common practice of deploying all-in-one branch office servers. Finally, combining the best of both worlds, a Windows Server 2008 Enterprise Edition Server Core Hyper-V server can be deployed with up to four Windows Server 2008 virtual machines, to provide isolated services for:

1. RODC/GC/DNS
2. DHCP services
3. File & Print Services
4. a local application server

the first three of which can also be deployed as Server Core implementations, significantly minimizing the resource footprint required by those services and providing an excellent branch office solution on a single physical computer. Additional discussion on server virtualization topics, and specifically Hyper-V, is contained in Domain 4.

Advanced Services

In addition to the primary services, other options for branch office services include:

- Deployment of Cluster Service nodes to provide local access to centralized services, with the side-benefit of providing offsite redundancy to the central office services.
- Using Distributed File System (DFS) to provide local copies of centrally stored file repositories (an option to BranchCache that does not require a Windows Server 2008 server at the branch office).
- Windows Deployment Services (WDS) to assist in the deployment of operating systems and applications to branch office computers, eliminating the need for shipment of machines or onsite visits by desktop technicians to perform those tasks.
- Windows Server Update Services (WSUS) to ensure timely installation of security and critical system updates.

Security

Branch office server security is a significant issue and should be given every bit of consideration it deserves. We've already discussed options for using read-only domain controllers or server core installations to minimize risk of compromise of server-based data, but those two options barely scratch the surface of the effort appropriate to the task. Simply stated the effort made with regard to a branch office deployment should be no less than that made for the central office and, in many cases, needs to be significantly more.

The first consideration in any branch office deployment scenario, with regards to security, should be the organization's creation of effective and enforceable organizational security policies. Explicit enumeration of authorized and unauthorized actions by branch office staff; identification of penalties for violations; and a comprehensive security awareness training program, which includes regularly recurring training that all employees are required to participate in, are minimum requirements for a branch office deployment.

Electronic security methodologies should be evaluated and implemented, as appropriate, including network and host-based firewall technologies, intrusion detection and intrusion prevention services, server hardening procedures using "best practices" guidance, encryption technologies (EFS, Bitlocker), and Network Access Protection (NAP).

Physical security must be continually evaluated, including access to the facility by the general public, access to restricted areas by unauthorized persons, and risk of burglary, fire, and other external causes of facility damage.

Finally, all of the above should be wrapped in a functional auditing and audit review process that ensures violations can be identified in a timely manner, mitigation and remediation implemented to prevent further collateral damage, and appropriate modifications made to preclude any future violation of the same type.

Administration

Administration of a branch office will exist in one of two scenarios. Either the branch office is wholly administered by the organization's central staff using remote access and onsite visitation as required or an onsite branch office administrator will be identified. If centralized administration is implemented, the delegation of responsibilities will likely be consistent with the delegations implemented within the central office.

When an onsite administrator is identified, however, it's necessary to ensure appropriate administrative configurations are in place to facilitate the performance of the assigned administrator's duties but also prevent capabilities beyond the scope of those duties. Generally a branch administrator should be enabled through the use of AD Delegation to the Organizational Unit that comprises the branch office, rather than membership in the Domain Admins group. Even if the branch office is implemented as a subdomain, domain, or forest, it may be prudent to retain administrative rights within the central staff and delegate necessary rights to the branch office administrator via AD Delegation.

In addition, the effective use of "Enforced" site-level Group Policy Objects can be used in conjunction with delegated rights to provide the appropriate level of security and mandated configuration for the branch office Active Directory environment.

A new feature in Windows Server 2008, Administrator Role Separation, allows for the delegation of local Administrator privileges on a Windows Server 2008 domain controller without granting any access to the Active Directory services itself. This is especially useful in branch office scenarios where the onsite administrator may require typical "server operator" functionality, but should not have full Domain Admin rights. Administrator Role Separation is implemented on a per-server basis using the DSMgmt.exe utility. To implement Administrator Role Separation, granting local Administrator privileges to a domain user, execute this series of commands:

```
C:\>dsmgmt
dsmgmt: local roles
local roles: list roles
No local roles are defined.

Available roles:
    Distributed COM Users
    Guests
    Server Operators
    Print Operators
    Backup Operators
    Account Operators
    Windows Authorization Access Group
    Replicator
    Remote Desktop Users
    Pre-Windows 2000 Compatible Access
    Performance Log Users
    Network Configuration Operators
    Incoming Forest Trust Builders
    Administrators
    Terminal Server License Servers
    Performance Monitor Users
    Users
local roles: add corporate\branchadmin administrators
Successfully updated local role.
local roles: show role administrators
    CORPORATE\branchadmin
local roles: quit
dsmgmt: quit

C:\>
```

As an enterprise administrator, you should be familiar with the DSMgmt.exe utility, not just because of its ability to implement Administrator Role Separation, but also because of its several other capabilities. Quoting the help screen: “DSMgmt facilitates managing ... application partitions, management and control of [FSMO roles], and cleaning up of metadata left behind by abandoned [DCs].” In addition, you should be familiar with the various roles that can be assigned to domain users using this feature, as some cases may not even warrant granting of full local Administrator privileges!

Public Key Infrastructure

What is a *public key infrastructure* (PKI)? Before you can plan one, or design one, or even implement one, it is useful to know what a PKI is. I suspect that PKIs are, perhaps, one of the most underused features of the Windows Server operating system. Many businesses and organizations purchase third-party certificates for public websites that require SSL, and I’m certain a fair amount of self-signed server certificates are used inside organizations for internal-only use, but how many organizations actually have full-blown public key infrastructures. Regardless of what that number is, it’s something you need to be able to design and implement as an Enterprise Administrator, so let’s just jump in with both feet.

First thing you may find useful is a TechNet Library article “Understanding Public Key Cryptography” written for the Exchange Server 2003 documentation. In that article is an excellent section that describes the basic principles of public key cryptography as developed by Diffie and Hellman in the mid-1970s.

The basic premise of public key cryptography is that instead of a shared secret key, which was the mainstay of cryptographic algorithms until this time, the algorithm would use two keys—a private key known only to the owner of the key and a public key shared with anybody who wished to exchange encrypted data; the two keys together are called a key-pair and have a some type of special mathematical relationship. Neither key can be used to undo its own encryption; the matching key of the key pair must be used to decrypt the work of the other and is the only key that can do so. Based on these fundamental principles, it then follows that anything that can be decrypted with the public key must have been encrypted with the private key, and anything encrypted with the public key can only be decrypted with the private key. Thus we have a method of data privacy when encrypting with the public key as only the holder of the private key can decrypt the data, and we have a method of authenticating the source of the data when decrypting with the public key, as the data must have been encrypted by the holder of the private key. This latter functionality is the basis of what we know as *digital signatures*.

PKI Components

There are five primary components of a public key infrastructure that you need to be familiar with. Four are technology components; the fifth is written.

Digital Certificate—the reason why the PKI exists in the first place; an electronic credential manifested as a collection of bits that contains the public key of a key pair and the identity attributes of the holder of the private key of that key pair.

Certificate Repository—a place where digital certificates are stored and published. Active Directory is used as a certificate repository.

Certificate Authority (CA)—an entity that issues digital certificates in response to requests submitted from authorized requestors.

Certificate Revocation List (CRL)—a list of digital certificates that have been revoked by a certificate authority.

Certificate Policy and Practice Statement—a written document specifying the need for digital certificates, how they are to be used, and how the Certificate Authority is to be implemented.

When making a determination as to whether to deploy one's own PKI, or use a third-party vendor, it would seem at first, that this is a choice; but in practicality, all organizations need both. The purpose of a digital certificate for public web use is to establish identity and trustworthiness, and for that a third-party certificate authority (such as Verisign, Thawte, or GoDaddy) is necessary. However, the cost of digital certificates can be high if a commercial certificate is purchased for every place one is needed, and for internal use one can usually trust oneself to be able to authoritatively identify one's self, thus the use of internal PKIs is usually a less expensive alternative, albeit still more complex to implement than simply buying a digital certificate and installing it on a machine.

Designing and implementing a public key infrastructure is about 90% planning and about 10% doing; the greatest amount of planning work is wrapped up in analyzing and deciding what is needed. The public key infrastructure in a Microsoft Windows environment is implemented with Active Directory Certificate Services (AD CS). Let's look at some of the key decisions that need to be made when designing and implementing AD CS.

Certificate Authority Type

The first decision point is whether the PKI will be integrated with the Active Directory environment, or not, or perhaps the recognition that both capabilities are needed. Using an Enterprise CA requires the use of Active Directory, which requires that consumers of that service have access to the Active Directory implementation. Not all possible consumers will have access to Active Directory, so a Standalone CA can be created to issue certificates through means other than the use of Active Directory, such as web-based methods or email-attached requests processed manually. There are several other advantages and disadvantages to each methodology, but the fundamental point to apply is the use of Active Directory. A key secondary point is that Standalone CAs cannot make use of certificate templates; that feature requires the use of an Enterprise CA.

Certificate Authority Roles

There are three possible roles that a CA may be configured to provide.

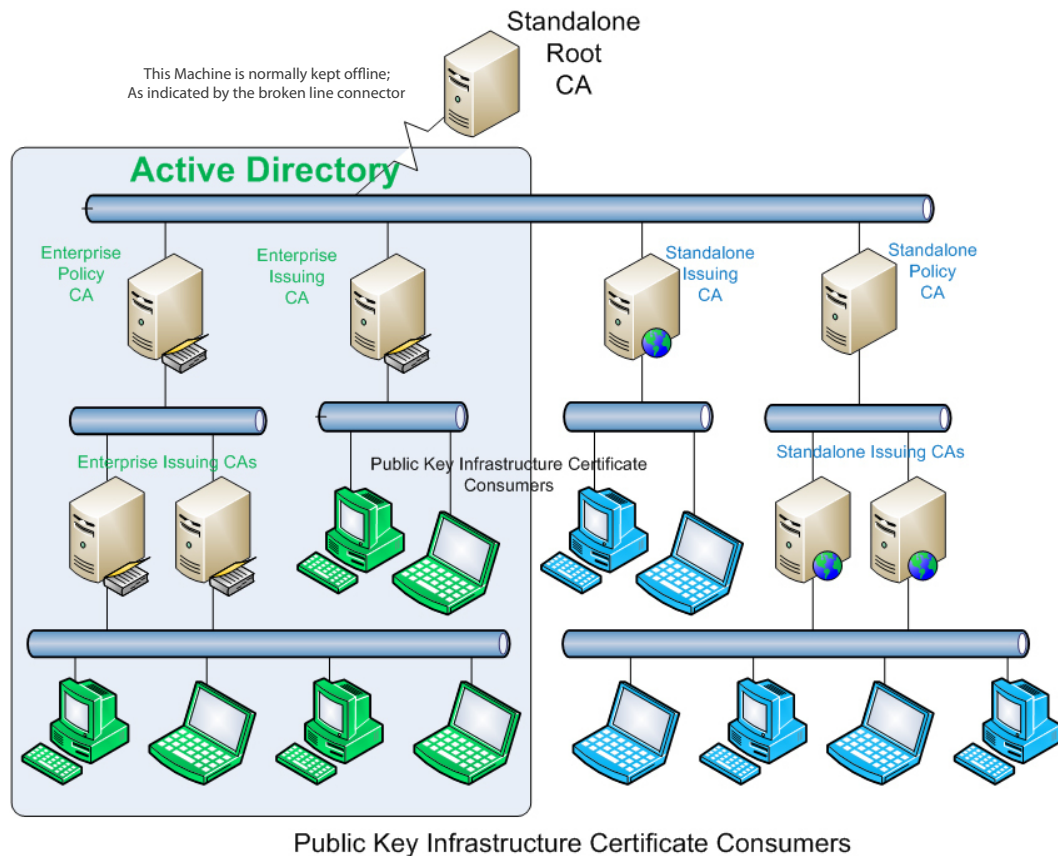
The *Root CA* is the top of the food chain. It is authoritative for all other CAs in the PKI; it issues certificates to other CAs authorizing them to publish and manage certificates on behalf of the Root CA. A Root CA should be a Standalone CA. A CA authorized by the Root CA is known as a *Subordinate CA*.

A Subordinate CA can take on one of two roles. An *Intermediate CA* is a subordinate CA that authorizes other subordinate CAs to publish and manage certificates. Normally an Intermediate CA does not handle certificates for consumers but is a control point for establishing variances in certificate policies. For this reason the Intermediate CA is often referred to as the *Policy CA*. The *Issuing CA* is a subordinate CA that publishes and manages certificates directly for the consumers of certificates. Such consumers may be users, computers, services, or network devices.

To summarize the possible types and roles of Certificate Authority Roles, you should be familiar with these five:

- Standalone Root CA
- Enterprise Policy CA
- Standalone Policy CA
- Enterprise Issuing CA
- Standalone Issuing CA

It's important to emphasize that the concept of an "Enterprise Root CA," while technologically possible, is highly irregular and strongly discouraged for security reasons. This diagram represents examples of the typical types of CA deployments that might be encountered in an Active Directory environment.



Because the Root CA is the top of the food chain, it should be given special treatment in a PKI. The Root CA should be a standalone CA and not a member of an AD domain. The machine hosting the Root CA should generally be kept offline, either by disabling the CA service, disconnecting the machine from the network, or ideally, powering the machine off and locking it in a closet. The Root CA must be brought online periodically to generate new CRLs and process certificate requests for subordinate CAs. A domain member system will encounter difficulties with domain computer account passwords, which are automatically changed every 30 days when the machine is kept offline and sporadically brought online. This complication is avoided by making the Root CA a WORKGROUP-based machine.

Recent enhancements in virtual machine technology have made the use of VMs ideal for hosting Root CA. A Root CA in a VM can be shutdown; the virtual machine and virtual hard disk can be burned to a DVD (or two, for redundancy) and removed from the virtual server. The DVDs can be stored in a lockbox, safe, fireproof vault, or bank safe deposit box. When it's needed to bring the Root CA online, simply copy the files from the DVD to the virtual server and fire the machine up. After the new CRL is generated and any pending subordinate CA requests or renewals are processed, shutdown the Root CA and burn the virtual machine and virtual hard disk files to a new DVD. Archive or destroy the previous DVD, depending on your organizational preferences for maintaining history.

To implement the Enterprise CA type, deploy a subordinate CA from the root CA as an enterprise CA; if a Standalone CA is needed to support non-AD consumers, create a second subordinate CA as the standalone CA.

In addition to determining the roles required it's also necessary to determine the number of CAs required. An initial determination might be made as to whether one CA can meet the organization's needs.

It's entirely possible that a single Enterprise Issuing CA can provide all of the services needed by a small or midsize organization. If one Enterprise Issuing CA is not sufficient, the question becomes: How many CAs are needed.

First, does your enterprise require more than one Root CA? Normally it's desirable to have a single Root CA from which all else is derived; however, organizations with highly decentralized management and no central technology administration may find it necessary to deploy multiple Root CAs at the business unit level.

One approach to answering the question of how many subordinate CAs are needed is to count from the bottom up. Start with identifying the number of Issuing CAs that is required. This will depend on a number of factors including the different policies required for issuance of certificates; the organizational hierarchy or geographical distribution of the organization that may demand issuing CAs in a division, business unit, or regional office; the requirements for redundancy; the actual work load expected; and whether one or both of enterprise or standalone CAs are required to meet the needs of the consumers.

The policy issues that affect the number of Issuing CAs required include things like the length of the private key, the lifetime of the certificate, and whether any particular certificate storage and management requirements are being imposed, such as would be the case for government entities or contractors subject to the requirements of Federal Information Processing Standards (FIPS).

Once the number of Issuing CAs required has been identified, determine the number of Policy CAs that is needed to bridge between the Root CA and the Issuing CAs. As suggested by the name, to a great extent this will be driven by the diversity of issuing policies but also by the need for redundancy at the Policy CA level and whether Standalone Policy CAs are required. It should also be noted that Policy CAs are not a strict requirement of a PKI. It's quite conceivable that an Issuing CA might be a direct child of the Root CA, particularly in a small organization that only needs one Enterprise Issuing CA; however skipping the use of the intermediate CAs could have ramifications affecting the growth and expansion of the PKI, so such decisions should be given great consideration.

Certificate Templates

Certificate templates are used in an Enterprise CA to facilitate certificate enrollment. A certificate template defines the type of certificate required and the specific consumer information that must be collected in order to issue the certificate. Certificate templates have evolved significantly over the lifetime of AD CS. The first certificate templates, introduced with Windows 2000, were static; the second generation allowed for customization of the template, and these were distributed with Windows Server 2003, known as version 2 templates.

Even richer customization capabilities exist in the third version of templates which are included with Windows Server 2008. These enhanced capabilities include properties related to Cryptography Next Generation (CNG) algorithms, which are the replacements for the existing CryptoAPI currently contained in Microsoft Windows operating systems. However, v3 certificates can only be issued to v6 operating systems—Vista, Windows 7, and Windows Server 2008.

Certificate Life Cycle Management

The life cycle of digital certificates covers the processes from initial certificate requests until the termination of the existence of the certificate, either as a result of expiration or revocation. Several methodologies exist for certificate enrollment (request) and approval (issuance), but all of them can be classified as either automatic or manual, and which methodology is used is decided by whether automatic or manual methods are preferred. Quite often both are used, depending on the nature or type of certificate. Standard types of certificates distributed to all computers or a large number of users would probably benefit from automatic enrollment and approval, whereas certificates that represent specialized or restricted use scenarios or high-security access might require manual approval, possibly even manual enrollment if the expected number of requests does not justify the effort involved in deploying an automatic enrollment capability.

Certificate enrollments and approvals, particularly for computer certificates, can be fully automated through Group Policy using the Automatic Certificate Request settings option. Certificate enrollments can also be processed through the Certificate Request Wizard found in the MMC Certificates console or through Web Enrollment pages, which run in conjunction with IIS installed on the certificate server. A new feature in Windows Server 2008 is the Network Device Enrollment Service (NDES) which is the Microsoft implementation of the Simple Certificate Enrollment Protocol (SCEP). SCEP/NDES allows network devices to enroll for X.509 certificates from a standalone CA.

Certificate enrollments can also be restricted or limited by using a restricted enrollment agent. The restricted agent is assigned one or more certificate templates it is authorized to process and users or groups it is authorized to enroll on behalf of. Access to the template can be restricted directly by modifying the DACL for the template. Those who are granted access will have Enroll and Read permissions; those who do not have those permissions will not be able to enroll with that certificate template. All certificates have a specified valid period and must be renewed with the Issuing CA prior to their expiration. Expired certificates cannot encrypt or sign content but can still be used to decrypt digital signatures created by the certificate holder. The specified valid period for a certificate should be determined based on the key length and the risk of compromise of that key prior to the end of the certificate lifetime. Additionally, the expiration of parent CAs also imposes immediate expiration of any child CAs certificates; therefore, it's critically important for the continuity of a PKI to ensure that all root and subordinate CAs are renewed in a timely manner.

Finally, certificates may be subject to revocation. Revocation is simply the forced expiration of a certificate prior to its regularly scheduled expiration date. Revocation information is distributed in one of two fundamental ways. The first is the Certificate Revocation List. Each CA must maintain and regularly publish a Certificate Revocation List (CRL) which itemizes every certificate that has been revoked by the CA by certificate serial number and the reason for the revocation. CRLs are published in two types: a *base CRL* and a *delta CRL*. The base CRL contains the complete list of revoked certificates. For a mature or busy CA, this list may be quite long. Moving a large CRL around the network can be a significant endeavor, so delta CRLs were created. A delta CRL contains only the serial number and revocation reason for certificates that have been revoked since the last base CRL was published. A delta CRL also requires the most recent base CRL as a reference document in order to possess the complete list of certificate revocations. CRLs can be published via HTTP or LDAP.

One of the significant disadvantages of CRLs, however, other than the size of such lists, is the latency between the actual revocation, the publication of the CRL, and the distribution of the CRL to the CRL publication point. In addition, the CryptoAPI supports caching of CRLs, and if the local cache contains a valid CRL, even if it's not the latest CRL, it will be used. This has the impact of potentially allowing a revoked certificate to be used when it shouldn't be.

A new technology introduced in Windows Server 2008 is the Online Certificate Status Protocol (OCSP) service. An online HTTP-based service, known as the OSCP responder, allows for real-time inquiries as to a certificate's validity. If the Issuing CA supports the use of OCSP, the URL of the OSCP responder will be contained in the certificate's Authority Information Access extension.

Finally, as with all other "new services" introduced in Windows Server 2008, use of the new features and services in the PKI requires the updating of the forest schema using 'adprep /forestprep' and the deployment of one or more Windows Server 2008 systems to host the new services. For additional study materials on implementing AD CS, see the TechNet Library article "Best Practices for Implementing a Microsoft Windows Server 2003 Public Key Infrastructure" at [<http://www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/security/ws3pkibp.mspx>].

Domain 4 – Designing for Business Continuity and Data Availability

Objectives

Business continuity is a relatively new buzzword, but the principal need for the capability has existed for eternity. Every business has the need to continue operations in the face of unforeseen circumstances whether it's something as simple and seemingly mundane as a senior executive retiring or resigning or as catastrophic as a hurricane or wildfire that makes physical access to the place of business impossible for days or weeks. Data availability is the technological implementation of the capability to provide business continuity and ranges from the simple restoration of a file when accidentally deleted to the complete relocation of the data center when physical access is impossible.

The specific objectives covered in this domain are:

- Plan for business continuity
- Design for data management and data access
- Design the operating system virtualization strategy
- Design for software updates and compliance management

Business Continuity Planning

Having a *plan* for how to maintain business operations is the key ingredient of ensuring that a business can continue to function in the face of unforeseen circumstances. The plan, however, cannot merely consider catastrophic disasters; it must also consider every day minor scenarios that have the potential to interrupt business operations (even if measured only in hours, rather than days or weeks). There are three areas of business continuity you should consider in your planning efforts: availability of services, backup (and restoration) of data, and recovery.

Service Availability

Every business continuity plan should include consideration for how existing and daily operational services will be configured to ensure their availability when needed. Much of the design efforts in Active Directory are related to ensuring continued service availability. These efforts include redundant domain controllers (with DNS), DHCP servers with overlapping scopes, firewalls, and other security tools to prevent denial of services that may result from unauthorized access.

Individual servers also implement features to ensure continued service availability in the form of RAID arrays, redundant disk controllers, network interface cards, power supplies, and hot swappable memory to name a few. In addition to simple placement of redundant independent servers and implementation of component redundancy, there are two Windows “features” that you must be familiar with—*Network Load Balancing (NLB)* and *Failover Clustering*. Identifying when their use is appropriate, as well as being able to design the implementation of each is a necessary skill.

Network Load Balancing

Network load balancing is a means by which multiple servers can be collected together and presented to the consumer of their service as if they were a single machine. A primary advantage to the use of NLB is that a workload many times greater than the capabilities of a single machine can be handled; a secondary advantage is that when a collection of machines is working together to share the work, the failure of a single machine does not remove the availability of the service, although it does have an impact on the overall performance of the environment because other nodes must pick up the extra workload. NLB is typically used for web servers, remote access servers, terminal services, firewalls, and other applications. The primary consideration for implementing NLB is that the client can connect to any NLB node at any time and receive exactly the same functionality and data. NLB is not appropriate for services that depend on local datastores or need to maintain client state on the server across client sessions or where clients are submitting data or modifying data stored on the server.

NLB can, however, be used in some scenarios which require the client to maintain state. This is done by configuring *port rule affinity*—essentially assigning a client session to a particular node of the NLB permanently, somewhat akin to using CPU affinity to assign applications to one or more CPU cores.

An NLB *port rule*, is a methodology used to assign behavior rules of the NLB cluster to a specific IP Address: Port identity. Every NLB cluster is configured with a *Default Port Rule* which essentially allows traffic for any service of the cluster to be directed to any node of the cluster. This rule is satisfactory for RRAS/VPN servers and Terminal Server farms but is insufficient for ISA Server NLB clusters and may not be sufficient for certain web applications, such as ASP.NET applications that require session state to be maintained.

A Custom Port Rule can be created and applied to specific services or applications. The custom port rule defines an IP Address:Port Range specification that the rule applies to and is configured with one of three behavior options: Multiple Hosts, Single Host and Disabled. The multiple hosts option allows all cluster nodes to respond to client requests, and the single host option restricts the response to a specified node. The disabled option prevents the cluster from responding to the traffic at all.

In addition to the node behavior, the custom port rule can also specify affinity when a rule is configured to use the multiple hosts behavior. Affinity identifies how a specific client is assigned to a particular node of the cluster for services. The options are none, single, and Class C. None is used to implement the standard NLB cluster behavior, making all nodes available at any time to any client, and is appropriate for services that are stateless, such as static web pages. Single is used to ensure a single client’s traffic is always sent to the same host where the session was originally initialized. This option is appropriate when state must be maintained, such as with SSL-based sessions, or certain ASP.NET applications that depend on session state. The Class C option is used when the client requests are coming from multiple IP identities, such as when multiple proxy servers are being used by a single client or if site-level assignments to a particular node are desirable.

Using affinity, however, somewhat minimizes the benefits of using NLB and, in some instances, would be no better than having individual standalone servers assigned to groups of clients. If the assigned server is offline, the service is unavailable to that assigned client.

Implementation of NLB in Windows Server 2008 is fairly straightforward. Install the NLB service on the first node, create a NLB cluster identity, join the first node to the cluster, and then install the service on each additional server and join that server to the existing cluster. No special hardware requirements are needed, although it is recommended that the nodes of the NLB cluster have a second, private, LAN segment to be used for cluster communications and management.

Failover Clustering

Failover clustering is more concerned with ensuring the continual availability of a service than it is with sharing the workload. Unlike the NLB environment where each server has its own independent local storage, a failover cluster uses shared storage that each node has access to. All service and data information is stored on the shared storage, such that if a node fails, another node can immediately take over, using the state of the failed node as stored on the shared storage.

Failover clustering is an appropriate solution for database servers, mail servers, file and print servers, and any other service that's dependent upon data consistency or session state and requires uninterrupted availability.

The most common scenario for failover clustering to be implemented is the two-node Active-Passive cluster. In this scenario, one server (the Active node) provides all of the services as it would if it were a single server. The second server (the Passive node) remains powered on in a "standby" state, monitoring and ready to take over should the Active node experience any failure in the ability to provide the service. The other variation of the two-node failover cluster is the Active-Active configuration. In this scenario, some of the benefits of NLB are available in that load-sharing can be implemented, but should one node fail, just as with NLB, the other node has to take on 100% of the workload demand. The primary distinction with NLB here is that the two nodes of an Active-Active configuration have shared storage, so the second node can also continue providing uninterrupted services to sessions already active on the first node. On an NLB node, the sessions would be dropped and the client would have to reconnect to resume using the service. You can also have multi-node failover clusters, which would consist of two or more active nodes and one or more passive nodes.

Installation and configuration of failover clusters are significantly more complex than NLB clusters and generally require more stringent hardware requirements. Servers ideally will be of the same manufacturer and model, however, this requirement has been relaxed in recent years with the advancement of the technology. Special device controllers are required (e.g. iSCSI, Fibre Channel) and they must be identical, including matching firmware revisions. Storage services are required and must be compatible with the device controllers as well as with clustering services.

Windows Server Backup

Probably the most significant topic in this domain is the complete redesign of the native backup services found in Microsoft Windows. There are two considerations in dealing with this topic. One, that you need to know how to use the utility to pass the exam; two, that today there are significant other resources available for data backup and recovery that may be better alternatives to using Windows Server Backup (WSB).

In terms of the alternatives, you might consider implementing System Center Data Protection Manager 2007 and backing up to a central data store. You might still be using a tape library and a third-party utility to manage the tape library and make backups. You might be on the cutting-edge using "backup to the cloud" with a vendor provided agent that continually sends data modifications to a backup store in the cloud. But, you might still have needs for local backups of servers, services, or data collections that are not practical to implement using DPM, Tape Libraries, or cloud services, which brings us back to the new Windows Server Backup in Windows Server 2008.

This utility is covered (and tested) extensively in the MCITP: Server Administration examination (70-646) and, even in the scenarios described here, will much more likely fall under the purview of individual system administrators rather than an enterprise administrator. However, as an enterprise administrator you need to be aware of this tool, and its capabilities and limitations to properly plan for its use when appropriate in the enterprise business continuity plan.

Perhaps the most significant change in this new tool is that it does not support backup to tape. This may be a good thing for those who have struggled through dealing with the Windows Server 2003 “Removable Storage” feature but, practically speaking, given the cost and speed of external USB 2.0 drives and the voluminous amounts of NAS and SAN storage available, the only place tape is still a viable solution is enterprises with established tape libraries and tape servers or isolated remote servers with built-in legacy tape drives that don’t have access to network bandwidth or USB connectivity.

The new WSB makes extensive use of Shadow Copies, so it’s also necessary to be familiar with the implementation of, and impacts of using, shadow copies on a server. Shadow copies are a great feature for use on file servers and providing consumers the ability to manage recovery of inadvertently modified or deleted user files. However, the use of shadow copy services on other types of servers, particularly where integration and consistency of file modules in applications and services is of concern, needs to be managed.

WSB cannot write to optical media (CD/DVD) or network locations during a scheduled backup. This means that you must have adequate locally attached storage (internal disk, USB external disk) or a SAN that appears to the system as locally attached storage. WSB does not back up FAT file systems; it only backs up NTFS file systems at the *volume* level, not at the individual folder or file level, and WSB cannot read legacy backups made with the Windows Server 2003 NTBackup.exe utility.

Finally, perhaps the one good feature of the new WSB: It writes backup files as a Virtual Hard Disk (VHD) file, which can be mounted and read by any of a number of utilities that support reading/writing VHD files. So while folder/file level backups cannot be taken, it becomes a trivial exercise to mount a VHD backup volume and copy folders and files during direct access technologies. Combining the planned use of Windows Server Backup with appropriate volume configurations on servers can also help streamline backup and restore operations. For example, it once again becomes reasonable to build multiple volumes on a server and segment the storage uses of those folders if the use of WSB is a consideration.

Active Directory Recovery

No discussion of service availability, redundancy, or backup/restore operations would be completed without also discussing the specialized recovery procedures for the Active Directory database. The good news, however, is that it’s become easier to do AD database backup and recovery in Windows Server 2008. First, in consideration of the new volume-based backup capabilities of Windows Server Backup, you should plan the volume structure of your domain controllers accordingly. One recommendation is to place the AD database (ntds.dit) and the SYSVOL folder on volumes separate from the operating system installation. This facilitates selective backup of the AD database or the SYSVOL folder without having to involve backing up gigabytes of unrelated operating system files.

Second, whereas in earlier versions backing up a domain controller was as simple as doing a System State backup, in Windows Server 2008 you need to specify individual volumes that must be backed up. The domain controller critical volumes include (and presumably should be separate volumes):

- The boot files volume (bootmgr and the BCD store)
- The operating system volume (including the registry hives)
- The SYSVOL folder volume (should be on its own dedicated volume)
- The AD database volume (ntds.dit)
- The AD database log files volume

Now that you've designed and implemented a functional backup strategy for Active Directory, we can turn our attention to the (hopefully never necessary) recovery of that data.

Restoration of data in the AD database is classified in two types, depending on whether the reason for the restoration is due to an issue with the *server* or an issue with the *data*. When a domain controller server fails (e.g. a volume containing AD data is corrupted or deleted), a restoration of the entire AD database on the server is necessary. In this scenario, our primary purpose is to get the server functioning again; we're not immediately concerned with the data itself—the data, if out of date or incorrect, will be replicated from another DC once we restore the database. For this type of restoration, we perform a *nonauthoritative restore*. Simply stated, nonauthoritative in this context means we know that this data is not the actual data; we'll just consider it pseudo-data for the sake of restoring services to the server.

In the scenario where one or more individual data items are deleted from a fully functioning AD environment, the issue turns to the restoration of the *data* not the restoration of the service on a particular server. This is known as an *authoritative restore*, as this particular domain controller will then possess the only valid copy of this data (the other DCs having also deleted it pursuant to replication a few minutes after the original deletion). Typical examples here are the deletion of an OU with a large number of objects, complex hierarchy, or other changes/deletion to data where the effort of restoration would be less significant than recreating the data and its associations.

Not much has changed over the years with the procedure to perform restores using Directory Services Restore Mode (DSRM). DSRM requires you to have access to the DSRM password created when the domain controller was originally initialized and requires a complete reboot of the server to enter DSRM. Restoration of the AD volume(s) with the issues will be performed using Windows Server Backup, or the command line utility `wbadmin.exe`, on Windows Server 2008. Authoritative restores require the additional step of using `ntdsutil.exe` to mark the objects as authoritative.

However, in addition to performing restores using DSRM, Windows Server 2008 provides a new capability in AD that exposes the directory service as a stoppable/restartable service in the Services console. This facilitates certain types of AD maintenance that previously required booting into DSRM, such as restores to the database, and also allows maintenance activities to be performed without interrupting other services on the domain controller, such as DNS, DHCP, or Certificate Services.

Data Security

Another area of consideration related to business continuity, service availability, as well as general user data availability is data security. Data security can be viewed as preventing or controlling access to specified data based on authentication and authorization. There are three techniques for employing data security that need to be considered and planned for.

The first is Active Directory Rights Management Services (AD RMS). AD RMS allows for the control of confidential or sensitive data by controlling what an intended recipient of that data can do with the data, such as read, copy, edit, print, forward, or save. AD RMS is implemented as a service accessible to internal and external users, and authenticates users of the data and authorizes the behaviors they can perform on that data.

The primary limitation of AD RMS is that the application that created the data must be AD RMS enabled. The short list of such applications includes Office 2003, Office 2007, Exchange 2007, MOSS 2007, IE6, and anything written using the XPS document format. For Windows Server 2003 and earlier systems, the RMS client was also required to be installed; this client software is built-in to the v6 operating systems. The second is Bitlocker. Bitlocker has been significantly enhanced for Windows Server 2008 R2 and Windows 7, but these enhancements are not likely to appear on the 70-647 exam in the near future. The basic Bitlocker functionality provides encryption services for system volumes and data volumes. Data volume functionality is not available in Windows Vista. Windows Server 2008 R2 and Windows 7 expand Bitlocker capability to external drives (USB drives; flash drives), but you need to be cognizant that this capability does not exist in the original editions of Windows Server 2008, which is what the certification exams are currently based on.

Bitlocker generally requires the system hardware to be Trusted Platform Module (TPM) v1.2 capable. There is an operational mode that allows for the use of a USB flash device to store the encryption key, but using this method requires that appropriate considerations are given to the physical security of the USB flash device. The benefits of Bitlocker are useless if the USB flash device is kept in the bag with the stolen notebook. Bitlocker authentication is provided on TPM-capable systems in one of three modes:

- **TPM only.** This is an easy to use mode as it requires no user intervention and provides protections against rootkits and malware, but is not appropriate for portable computers, as it provides no protections at all if the computer is lost or stolen.
- **TPM with PIN.** The advantage here is that the PIN is something the user *knows* rather than something the user has, which the thief might also have if both are stolen at the same time, but like passwords, PINs are inherently insecure, easily attacked by brute-force methods, and users have a tendency to write them on sticky notes and tape them to the computer or the inside of the computer bag.
- **TPM with USB flash.** This mode suffers from the same potential risk as the non-TPM mode using only the USB flash device, but if the USB flash device is kept separated from the computer when powered off, it has the ability to protect the data if the computer is stolen. With proper management of the USB flash device, this is the most secure of these three authentication modes.

Encrypting File System (EFS)

For machines where Bitlocker is not a viable encryption methodology, its predecessor, EFS, is available. EFS combines symmetric key cryptography with public key cryptography to encrypt data; it does not require TPM or a USB flash drive for implementation and is supported on Windows 2000 and later systems (but not Home Editions of XP/Vista/Win7) and other non-Microsoft operating systems. The files are encrypted with a symmetric key, and the symmetric key is then encrypted with a public key and stored. In a domain environment, the private key (and certificate) can be issued by an Enterprise CA and maintained in the Active Directory database. In addition, a domain user can be designated as a Data Recovery Agent (DRA) in the event the original key/certificate is lost or destroyed. On a non-domain machine, the certificate and private key are stored on the local computer, minimizing the actual value in the file encryption effort.

It's also noteworthy that EFS only encrypts data while it's stored on the local disk, but during transmission, and once transmitted elsewhere, it is no longer encrypted. As an advantage, though, and unlike Bitlocker which only encrypts at the volume level, EFS has the ability to encrypt at the folder or file level. Evaluation of the use of EFS should be part of the planning scenario when Bitlocker is not a viable solution or not an appropriate solution based on the actual encryption needs of the organization or individual user.

Data Sharing and Collaboration

Distributed File System (DFS)

Distributed file sharing is not a new feature in Windows Server. The concept actually dates all the way back to Windows NT v3, but with each release of a new server operating system the capability has been enhanced and sometimes renamed. The immediate predecessor of the Windows Server 2008 Distributed File System (DFS) was known as the File Replication Service (FRS), and it dates back to Windows 2000.

The principle behind DFS is to create a single namespace that is used to reference shared file-based resources without requiring the user to be aware of the file's actual physical location. This capability provides two benefits: One, files can be physically located anywhere in an enterprise's global network; and two, files can exist in multiple writeable locations, are accessed from the nearest storage location, and are synchronized through a background file replication service.

The namespace created to support DFS can be either a standalone namespace which is stored in the registry of a single server and is used to simply consolidate individual independent file stores into a single namespace or a domain-based namespace, which is stored in and published from AD DS and supports file replication and fault tolerance capabilities. Two enhancements of the Windows Server 2008 DFS implementation is the capacity for larger file stores—over 5,000 folders—and access-based enumeration of files, which simply means that users cannot see files they don't have permissions to access. To take advantage of these enhancements, all servers hosting DFS resources in the domain-based namespace must be running Windows Server 2008, and the functional level of the domain must be Windows Server 2008.

Related to the implementation of DFS is the methodology for providing the file replication. The new DFS file replication capability uses Remote Differential Compression (RDC), which is optimized to support delta-based file replication over bandwidth constrained links. RDC makes use of AD site links to optimize replication scheduling. RDC can be selectively enabled or disabled and may be disabled on LAN-based replication connections where bandwidth is not an issue. Also, RDC is not implemented on files smaller than 64kb; these files are simply compressed and replicated *in toto*. In addition, cross-file RDC is used when the target of a file replication effort already contains portions of a file's content, perhaps even contained in multiple unrelated files. For the portions of the source file that can be found on the target server, that content is merely copied from the existing files on the target server, rather than replicated across the network.

DFS Replication can be specified in time windows of 15-minute increments across a seven day period and allows for throttling of bandwidth availability from 16kb/sec to 256mb/sec or no throttling at all. Each connection can have its own customized replication configuration or be configured to use an enterprise default schedule. All schedules are based on the time zone of the target (receiving) server and can be based on either local time or Coordinated Universal Time (UTC), as specified in the connection's replication configuration settings.

Two special folders are used in conjunction with DFS file replication services. The *Staging Folder* is used to establish a cache of files that are pending replication to a target server. The staging folder has a configurable quota that governs housekeeping of the folder contents and must be coordinated with the replication schedules and available disk space to ensure only successfully replicated file content is purged. Each folder configured for replication has its own dedicated staging folder. The *Conflict and Deleted Folder* is used to handle conflict management on file changes, as well as copies of deleted files, and works from a "most recent change" wins philosophy. The file that doesn't get its changes replicated is written to the Conflict and Deleted Folder where the files can be reviewed and any missed changes can be re-implemented. This folder also uses quota management for housekeeping. There is one Conflict and Deleted Folder per DFS file server.

DFS is best used in a geographical distributed organization that relies on file-based documents and data that need to be accessed at multiple sites in the geographic distribution but can also be implemented in a centralized organization with a file server farm in conjunction with network load balancing.

Customizable features of DFS that you must be familiar with include:

- **Referral Ordering.** A prioritized list of resources for accessing the contents specified by a namespace location based on cost, resource exclusion based on site, or random ordering (to facilitate load balancing in a file-server farm, for example).
- **Target Priority.** Enhancing the referral ordering by specifying additional prioritizations such as first or last among all or first or last among equals. Target Priority establishes a resource's permanent presence in a referral ordering list, even if the resource has been otherwise excluded by site from the list.
- **Failover and Failback.** Automatic selection of a secondary resource for accessing a namespace location when the primary is unavailable and the automatic reversion to the primary resource when it comes back online.
- **Namespace Redundancy.** Namespace servers can be deployed as multiple instances to mitigate the need to cross network connections to perform namespace queries for available resources.
- **Namespace Scalability Mode.** Used in organization with more than 16 namespace servers to provide a more efficient namespace replication scheme based on "nearest neighbor" rather than the default replication methodology which is based on a full-mesh replication and hourly polling of the PDC emulator in the domain.
- **Disabled Memberships.** Allows selective replication of a folder to specified targets by revoking the memberships of undesirable partners in the replication partnership.
- **Replication Filtering.** Allows selected folders and files to be excluded from file replication completely.

Window Sharepoint Services (WSS)

While the mainstay of document collaboration for many years has been file sharing and file-based services, the advent of Windows Sharepoint Services (WSS) and its document libraries a few years ago is rapidly replacing file-sharing and file distribution and replication technologies.

WSS is a database-based storage environment, which supports document collaboration tools, granular access permissions, document check-in and check-out, workflow and document routing capabilities, document metadata, and version control. The WSS environment is accessible via web browser or standard file access dialogs of Microsoft Office applications. Because it's based on database-storage, the database server can be configured as a failover cluster with multiple active nodes of the database distributed geographically to provide local LAN-based access to the contents of the document libraries. In this scenario, SQL Server database replication technologies are used instead of file-based replication technologies found in DFS. WSS can also be deployed as a single-server solution for use in workgroups, departments, or teams that require a dedicated resource for collaboration.

In addition to document repositories which supplant the need for distributed file services, WSS also provides conventional web-based services, such as team collaboration, event coordination, wikis, blogs, and personal management services, such as calendaring and to-do lists which can be integrated with Microsoft Outlook and Exchange Server. WSS is a free product and ships as an installable role on Windows Server 2008.

Microsoft Office Sharepoint Server (MOSS)

Built on top of Windows Sharepoint Services, the Microsoft Office Sharepoint Server provides enterprise resources, portal services, and entire community-based sites. MOSS is sometimes deployed to help consolidate and integrate multiple WSS services that have been independently deployed throughout an enterprise. MOSS includes advanced services for content management, business intelligence, enterprise search, project management, personal sites, portals, and several enterprise product vendors have implemented MOSS-enabled interfaces to their products, such as SAP and Dynamics AX.

Server Virtualization

Server virtualization is rapidly growing in enterprise usage and addresses a number of concerns in the realm of data available and business continuity. Virtualization of legacy operating systems running legacy line of business applications that cannot be upgraded or replaced allows those services to continue without dependence on failing hardware (or replacing legacy hardware with severely underutilized new hardware). The use of server virtualization has made accessibility to entire training labs and testing labs now within reach of many organizations who previously could not afford the cost of the several physical servers needed to implement such labs, never mind the physical space that would also be required; today an enterprise administrator can have an entire testing network sitting under the desk.

Virtualization allows for the implementation of network load balancing services and failover clustering services at significantly lower costs than if those services were implemented on individual physical machines. Two virtualization servers can be configured in a failover cluster and facilitate cluster-backed services on dozens of virtual machines running on such a cluster. Virtualization allows for the rapid redeployment of a server or service on an alternative physical machine, in an alternative physical location, providing exceptional capabilities for business continuity. Many organizations now find it affordable to maintain an entire offsite virtual environment which allows for a limited-load business critical environment to be turned up within hours, if not minutes, of the loss of services at the primary location.

One of the primary uses for virtualization in today's enterprise is server consolidation. An artifact of the OneService-OneServer philosophy is that many servers have significant amounts of underutilized resources: CPUs run at 5%, Disk I/O sits at 90% idle. Server consolidation allows for the maintenance of the OneService-OneServer philosophy but sharing of unused hardware resources. Rather than having four servers (DC1, DC2, File, Web) each with 5% usage of a 2GHz CPU and 10% usage of a single RAID mirror array, a single host can be configured with four virtual machines, which now has 25% CPU usage and 50% Disk I/O idle time (allowing another 5% CPU and 10% disk for the host operating system).

The key to planning a server consolidation project is properly identifying candidates for server consolidation through virtualization and identifying appropriate combinations of virtual machines to exist on a common host. One of the simplest ways to properly identify candidates for virtualization is to know how to identify servers that are *not* candidates for virtualization. Virtualization will not reduce the amount of resource consumption already being used by a server, so if a server is already consuming significant amounts of CPU cycles, resides on a busy RAID-based disk array, or requires x64 hardware with many gigabytes of dedicated RAM, it is not a candidate for virtualization. Everything that's left is a candidate!

When planning virtualization, resource availability and allocation is the primary concern. With today's hardware, CPU resources and RAM availability is rarely a constrained resource, but perhaps the most commonly under planned resource is disk requirements for the virtual host. On a virtual host, the host operating system should be installed on a dedicated physical volume, preferably on a dedicated RAID1 mirror. VHDs for the virtual machines should have their own dedicated spindle resources or a RAID1+0 array consisting of one spindle for every planned VHD would be a good starting point for a production server. Development and testing servers may be able to survive with less, and particularly so, if the virtual machines are not constantly powered on.

The process of creating a virtual machine is fairly simple and essentially the same regardless of the virtual hosting service being used.

1. Identify a name for the virtual machine. Virtual machine names are different from the machine's network hostname; the latter needs to be fairly concise. VM names, however, can be quite descriptive, as their sole purpose is to identify the VM in the VM management console.
2. Identify the memory required for the machine. Memory is a resource that cannot be over shared on a virtual host. Memory allocated to a virtual machine is not available for use by the host or other virtual machines. Thus the sum total of memory required for each virtual machine plus the memory required for the host must be less than the actual physically installed memory in the host machine.
3. Identify the network requirements. Virtual machines can be connected to an external network, to an internal (VM-only) network, or to a network that includes the host machine but no other physical machines. The latter two configurations are especially useful for testing environments and other scenarios that can benefit from an isolated environment.
4. Identify a virtual hard disk. The virtual hard disk can be a pre-existing disk, created as a separate task prior to creating the virtual machine, or created at the time the virtual machine is created. In the latter scenario, the VHD is usually empty and will need to have an operating system installed. The more common scenario is that a virtual machine's VHD is derived from a master installed image on another VHD.

Virtual hard disks can be specified with a fixed size, which will permanently allocate that physical space from the installed physical disk resources or can be specified as an auto-growing virtual size, which will use only the actual space required (or available). In addition to being fixed or dynamically sized, a VHD can also be defined as a child of another VHD. This is known as a differencing disk. A differencing disk contains only the changed bits from a standard parent disk.

A common practice is to create a single parent disk with a standard non-customized operating system installation, and then create multiple differencing disks to establish several virtual machines all deriving from the same operating system but to be used for different purposes. In this scenario, the common (parent) VHD can be placed on a high-speed read-optimized array. The differencing (child) VHDs, which will be write intensive, should be placed on one or more write-optimized arrays. While extensive discussion can be had about which RAID format is best for these scenarios, the most productive configuration is probably using separate controllers with appropriately configured caching. The controller for parent VHDs should be weighted 100% in favor of read caching; the controller for child VHDs should be significantly weighted in favor of write caching, but note that child VHDs also have read requirements when the changed bits are required to be reread.

Windows Server 2008 Hyper-V

There are two current methodologies for implementing server virtualization services available today. You need to be familiar with the requirements and limitations of both so that an appropriate implementation of virtualization services can be achieved. The current, premier, virtualization technology is Windows Server 2008 Hyper-V. Hyper-V makes use of hardware-assisted virtualization and 64-bit technology to allow implementation of several virtual machine instances on a single physical server. Hyper-V is available in a number of different implementations, as a dedicated Hyper-V only server, as a Server Core implementation, or as a full Windows Server implementation. Hyper-V is available on Standard, Enterprise, and Datacenter Editions of Windows Server 2008 x64; Enterprise and Datacenter Editions provide enhanced licensing arrangements for the deployment of virtual machines—Enterprise Edition provides four additional licenses for deployment of virtual machines, and Datacenter Edition provides unlimited licenses for virtual machines.

The Hyper-V server is managed from the Hyper-V console, which is installed on the Hyper-V Server but can also be installed as a remote console on Vista or Windows 7 clients. A significant feature advantage of Hyper-V is the availability of snapshots. Snapshots are a point-in-time image of a virtual machine. Hyper-V snapshots can be made on a live machine or with the machine powered off. Snapshots are an excellent tool to establish “recovery points,” particularly in testing scenarios. They can also be used to establish branching points of a system installation, such as with various service pack levels for a particular operating system or with various application installations for evaluating co-existence scenarios.

Snapshots can also be used as a performance enhancing tool. Because snapshots are implemented by the use of differencing disks, a virtual machine with a baseline installation actually has at least three VHDs involved:

- A parent (read-only) VHD which likely contains the core operating system installation.
- The primary (child) VHD which contains the baseline installation for the particular virtual machine, such as current updates, applications, etc., which becomes a volume for read activity only once the snapshot is created.
- The active (snapshot) VHD which is the target of all active writes.

So, the environment can be further optimized by placing parent VHDs on one array, primary (machine-specific) VHDs on a second array, and active snapshots on a third array.

Virtual Server 2005 R2

The alternative virtualization technology is Virtual Server 2005 R2. Virtual Server is a free application available for both x64 and x86 architectures that runs on server and desktop operating systems.

Virtual Server’s primary advantages over Hyper-V are:

- Installable on x86 hardware
- Installable on Windows XP and Windows Server 2003

Virtual Server’s primary disadvantages over Hyper-V are:

- Cannot run x64 virtual machines
- Virtual machines do not have access to symmetric multiprocessing

The primary scenario in which you might deploy Virtual Server is on an x86 Windows Server 2003 machine where you need to only support a couple of x86-based virtual machines within the 4GB confines of a 32-bit hardware environment, typically for legacy system survivability, but it is conceivable to implement Virtual Server on an x64 host with up to 256GB RAM which will support up to 512 virtual machines. An interesting comparison note is that Virtual Server on x64 actually supports a higher number of VMs than Hyper-V does (512 vs. 192 for Hyper-V) and a higher number of CPU cores on the host system (32 CPUs vs. 24 cores on Hyper-V); thus, despite the premier status of Hyper-V, there may be some scenarios in which Virtual Server may be a better solution than Hyper-V.

In addition, Virtual Server has a significant advantage for developers and testers in that it can be installed on Windows XP and Windows Vista desktop systems to host server operating systems for local use. It’s also useful for personal users to host legacy operating systems such as Windows 2000 or Windows 98 to facilitate running legacy applications on the current desktop operating system, although a desktop oriented product, Virtual PC, is the recommended solution for desktop virtual machines on a desktop host operating system.

System Center Virtual Machine Manager

In addition to the core virtualization technologies, there are two virtualization management tools of interest. System Center Virtual Machine Manager (SCVMM) 2007 is a virtual machine management environment designed to consolidate management of Hyper-V, Virtual Server, and VMWare environments in a domain-based enterprise. SCVMM can manage up to 400 virtual servers containing up to 8,000 virtual machines. It supports the creation and management of libraries that contain machine templates, hardware profiles, VHDs, ISOs, and scripts; the migration of virtual machines from physical to virtual, between libraries of virtual machines, and between virtual servers when implemented on a Fiber Channel SAN; capacity planning tools; and load balancing and distribution tools that ensure new VMs are deployed to the servers with the most available resources.

An SCVMM implementation consists of the SCVMM server, agents deployed on virtual servers, a SQL Server database, the administrative console, a self-service portal website, and a library server which contains the catalog of resources (ISOs, scripts, hardware profiles, machine templates, and VHDs) used to create new virtual machines.

Virtual Server Migration Toolkit

The second tool used in virtual environments is the Virtual Server Migration Toolkit (VSMT), which was primarily designed to facilitate the migration of physical machines to virtual machines for the Virtual Server application. It is a command-line based utility that builds and uses XML configuration files to assist in the migration process. Because Virtual Server machines can be imported into Hyper-V environments, it's also possible to use VSMT to migrate physical machines to virtual machines for subsequent import into a Hyper-V environment by first converting them to a Virtual Server 2005 virtual machine and then importing that virtual machine into Hyper-V. VSMT can also be used to convert VMWare virtual machines into Virtual Server 2005 virtual machines.

Patch Management

The last component of interest to us in this domain is implementation of patch management or update management technologies to assist in ensuring data availability and business continuity by protecting against server failures or service outages caused by security vulnerabilities that have been exploited or operating system defects that have been remediated through available OS updates.

Windows Update (WU) / Microsoft Update (MU)

Windows Update is a web based service, originally launched for Windows 98, to provide add-on packages, device driver updates, security updates, and critical updates. The "Year 2000" updates for Windows 98 were distributed via Windows Update. Today it continues to provide security updates and critical updates for Windows operating systems.

The functionality is implemented through a browsable website, an ActiveX control, and the Windows Update Agent. Required updates are identified by the WUAgent via the ActiveX control, and the user is offered the opportunity to choose which updates will be installed. The update content is then downloaded and installed immediately under the supervision of the ActiveX control and the Windows Update Agent. Activity from these sessions is logged to the WindowsUpdate.log, the ReportingEvents.log, and the Application Event Log.

The primary limitation of WU is that it only provides access to critical and security updates for the operating system and operating system components such as Internet Explorer, Outlook Express, Windows Mail, and Windows Media Player.

In 2005 an enhanced version of WU was made available, known as Microsoft Update. MU is an opt-in service and provides updates for additional products such as Microsoft Office, SQL Server, Exchange, Virtual Server, Visual Studio, and Windows Defender. It is the stated objective of Microsoft that all Microsoft products will be updatable using the MU engine.

With the release of Windows Vista, the user interface for this on-demand service was shifted from a web-based property to a control panel application in the operating system. If you attempt to browse to WU/MU from a Windows Vista or newer system, the control panel applet will be automatically launched.

Automatic Updates (AU)

The primary disadvantage of the WU/MU technologies is that they require user interaction to detect and install updates. Furthermore, this requires a user to actually remember such updating is needed. In 2000, shortly after the launch of Windows Update and as a replacement for an ill-designed Critical Update Notification Tool, the Automatic Updates service was launched.

AU is an automated update methodology that delivers security updates, critical updates, and service packs to individual computers once enabled following system installation. The update content is still limited to the operating system and key OS components. AU is implemented as a combination of functionality provided by the Windows Update servers and the Windows Update Agent on the local computer. Every 22 hours (minus a random offset of 0.2 to 4.4 hours (1%-20%)) the Windows Update Agent initiates a session with the Microsoft Automatic Updates servers and scans for available updates for the local machine. If updates are found, they are downloaded to the local computer using the Background Intelligent Transfer Service and scheduled for installation at 3am the next morning. If the computer is powered off at 3am, the updates are installed when the computer is next powered up.

With Microsoft's security push in 2002, and particularly with the launch of XP Service Pack 2, AU became a prominent figure in the installation and configuration of an operation system. Previously AU was simply an available service, but unless it was known the service existed, little was done to promote its use. Beginning with XP SP2, the user is now prompted at installation, as well as at other key installation points, such as upgrading to Internet Explorer v6, to configure the computer to use Automatic Updates. By 2005, almost 75% of all systems were being updated through Microsoft's services were being done via Automatic Updates.

Windows Server Update Services (WSUS)

For enterprises and even small and midsize businesses, however, the functionality of AU, WU, and MU does not entirely meet the needs of maintaining business computers. As noted, WU/MU requires user interaction, and AU, while automated, is an all or nothing scenario, particularly when combined with automated scheduled installations. Add in questions of update conflicts with line of business or other customized applications, and the prospect of having updates automatically installed onto a business computer becomes somewhat intimidating, and as some organizations discovered, particularly troublesome.

In 2003, Microsoft released a server-based product known as Software Update Services (SUS) designed to address these needs of businesses and organizations. SUS was a free downloadable application that was installed on a Windows 2000 or Windows Server 2003 system. It provided an alternative server interface for the AU client component and a web-based administration tool to allow administrators to selectively approve critical and security updates for installation. The behavior of the client system was still functionally the same, except now the AU client can only download and install updates that have been explicitly approved by the system administrator for installation. The following year, SUS Service Pack 1 was released and added the ability to manage the deployment of service packs.

Concurrent with the release of Microsoft Update, the next generation of SUS was released, named Windows Server Update Services. Like its predecessor, it installed on a Windows 2000 or Windows Server 2003 system and provided a web-based administrative interface to allow the approval of updates for installation. Like its new cousin, MU, it provided expanded access to updates for server and application products. In 2007, the current generation of WSUS products was released, v3, which migrated the administrative console from a web-based engine to an MMC-based snap-in console, which could be integrated into other custom MMC consoles for enterprise administrators. It also removed support for installation on Windows 2000 Server and shifted the underlying database engine from the MSDE (or SQL Server 2000) engine to the SQL Server 2005 engine (shipping the Windows Internal Database as a built-in component). WSUS v3 added support for installation on Windows Server 2008 with the release of Service Pack 1 in 2007; however, it uses IIS7 running in IIS6 compatibility mode. In January, 2009, an update to the Server Manager was released allowing WSUS to be installed to Windows Server 2008 as a role and integrating the WSUS MMC administrative console into Server Manager. This update also introduced the first real use of the new "Dynamic Installer" technology, which allows the installation of complete applications, on demand, via the WSUS/MU infrastructure.

The current version, WSUS v3 Service Pack 2, was released in August, 2009, and provides support for installation on Windows Server 2008 R2 editions. The primary feature enhancement of WSUS v3 SP2 was to support installation on Win2008 R2, and to recognize the Windows 7 client. In addition, other bugfixes and minor enhancements were added, the most notable being a new set of reports. The Windows Server 2008 certification exams are based on the WSUS v3 Service Pack 1 product, and there are no significant differences between SP1 and SP2 that would likely affect specific questions in any event. (It's likely that any subsequent Windows Server 2008 R2 certification exams covering this topic will be based on WSUS v3 Service Pack 2, given that is the only version that will run on Windows Server 2008 R2.)

Windows Update Agent (WUAgent)

To fully understand and appreciate the capabilities of AU and WSUS, one must understand the capabilities and flexibility of the WUAgent. With the release of WSUS many administrators became aware of the configuration capabilities of the WUAgent, but in reality, these capabilities have existed for many years.

The WUAgent determines its behavior through a collection of registry configuration values and default behavior. The registry configuration values can be used to override the default behavior of the WUAgent, and the registry can be configured directly (not recommended), with Local Policy, or with Group Policy. Comprehensive documentation for the registry values is available in the WSUS Deployment Guide, "Configure Clients in a Non-Active Directory Environment," and documentation for the policy settings is available in "Configure Clients Using Group Policy."

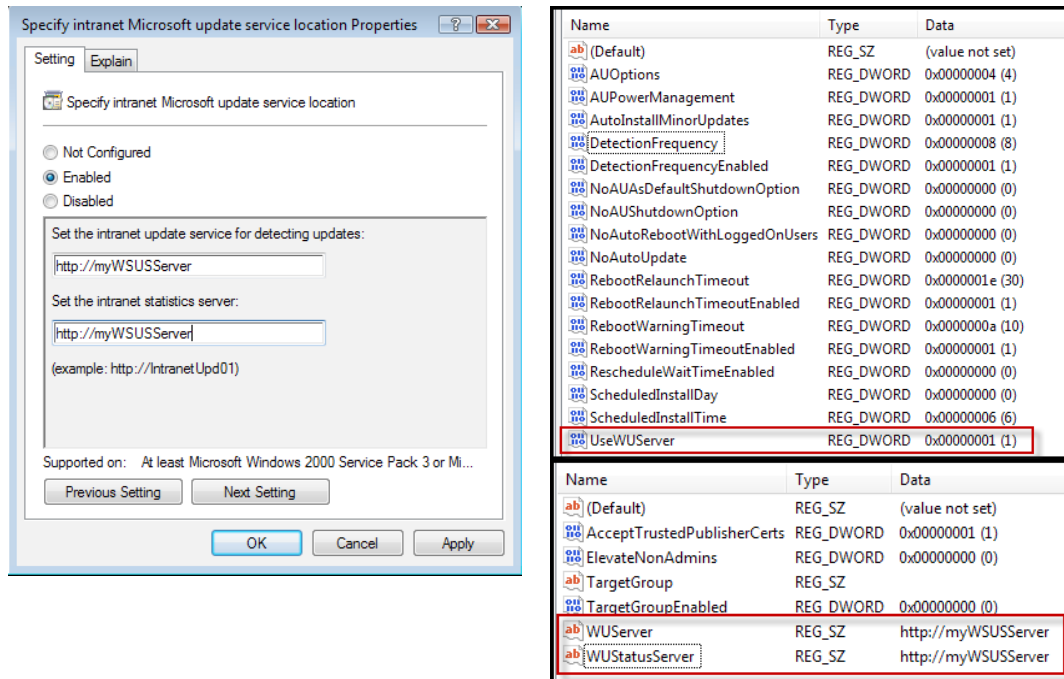
The registry values are contained in the registry key: HKLM\Software\Policies\Microsoft\Windows\WindowsUpdate and the subkey ~\AU, and the policy settings are contained in the path Computer Configuration | Administrative Templates | Windows Components | Windows Update. Here we will discuss the most significant of these several options.

Update Server

By default, the WUAgent will attempt to initiate a session with the Automatic Updates service at Microsoft. The WUAgent is configured to use a WSUS server with three registry values or a single policy setting.

Policy: *Specify intranet Microsoft update service location*

Enable the policy and enter the URL of the WSUS Server in both "Set the intranet update service for detecting updates" and "Set the intranet statistics server" text boxes. If this policy is disabled, the WUAgent will revert to using the AU service.



Registry: Set the *UseWUServer* value in the AU key to dword:0x1; set the *WUServer* and *WUStatusServer* values in the WindowsUpdate key to the URL of the WSUS Server. If *UseWUServer* is set to dword:0x0 or deleted, the WUAgent will revert to using the AU service. If *WUServer* and *WUStatusServer* are not identical, the WUAgent will revert to using the AU service.

Installation Policy

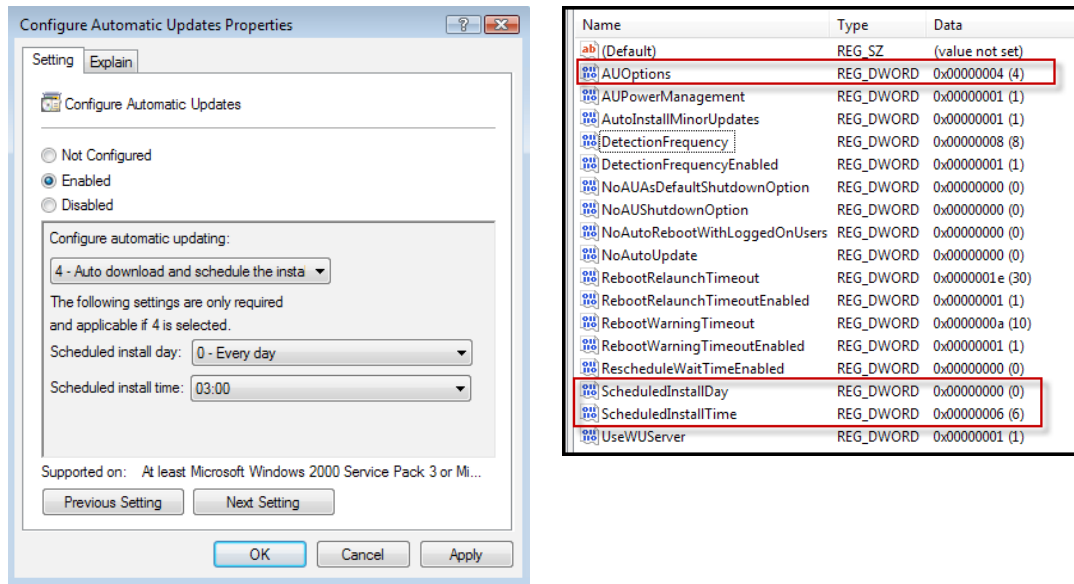
By default, the WUAgent will download all approved updates immediately after performing the scan, stage the content on the local drive, and install the updates at 3am the next morning following completion of the download of the update. The automatic installation behavior can be changed, as can the time of the scheduled event.

Policy: *Configure Automatic Updates*

Enable the policy. Set the "Configure automatic updating" value to the desired option:

- 4 – Install at the scheduled time
- 3 – Prompt for installation
- 2 – Prompt for download

If option #4 is selected, set the “Scheduled install day” and “Scheduled install time” to the desired values. If this policy is disabled, the WUAgent will revert to the default behavior of installation any day at 3am.



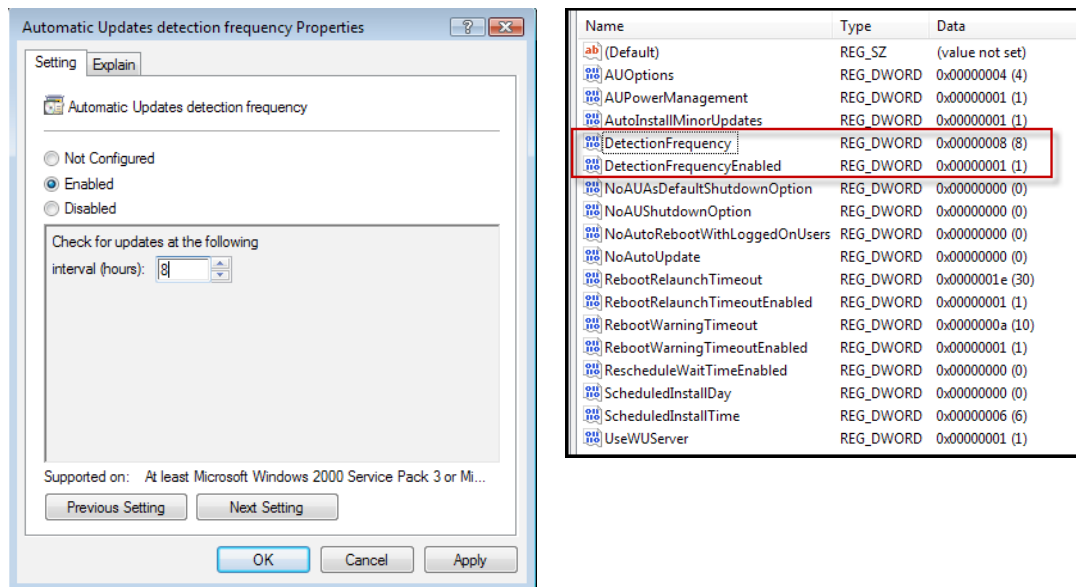
Registry: Set the *AUOptions* value in the AU key to the dword value based on the option desired (e.g. dword:0x4, 0x3, or 0x2); Set the *ScheduledInstallDay* value in the AU key to the dword value based on the option desired (e.g. dword:0x0 for every day, or dword:0x1 – 0x7 for the specific day of the week, Sunday thru Saturday); Set the *ScheduledInstallTime* to the dword value that represents the start time of the installation event based on a 24 hour clock (e.g. dword:0x3 for 3am, dword:0x10 for 4pm; dword:0x17 for 11pm). If the *AUOptions* value is not dword:0x4, then *ScheduledInstallDay* and *ScheduledInstallTime* are ignored. If *AUOptions* is dword:0x4 and *ScheduledInstallDay* or *ScheduledInstallTime* are missing or misconfigured, or *AUOptions* is missing, the WUAgent reverts to the default behavior of installation every day at 3am.

Detection Frequency

By default, the WUAgent initiates a detection session with the target service every 22 hours, minus a random offset of 1%-20% of the configured interval. This offset is recalculated at the completion of each detection event. The detection interval can be configured from 1 hour to 22 hours. One hour is useful for diagnostics and testing, but generally speaking anything less than 4 hours provides no useful benefit.

Policy: *Automatic Updates detection frequency*

Enable the policy and set “Check for updates at the following interval (hours)” to the desired value. If the policy is disabled, the WUAgent will reset the detection frequency to 22 hours.



Registry: Set *DetectionFrequencyEnabled* in the AU key to dword:0x1 (true); set *DetectionFrequency* in the AU key to the dword value representing the desired hourly interval (e.g. dword:0x16 for 22 hours, dword:0x4 for 4 hours). If the *DetectionFrequencyEnabled* value is missing or not set to dword:0x1, the *DetectionFrequency* value is ignored and the default value of 22 hours is used. If the *DetectionFrequency* value is invalid (e.g. dword:0x0 or greater than dword:0x16) the WUAgent will revert to the default detection frequency of 22 hours.

System Center Configuration Manager (SCCM) 2007

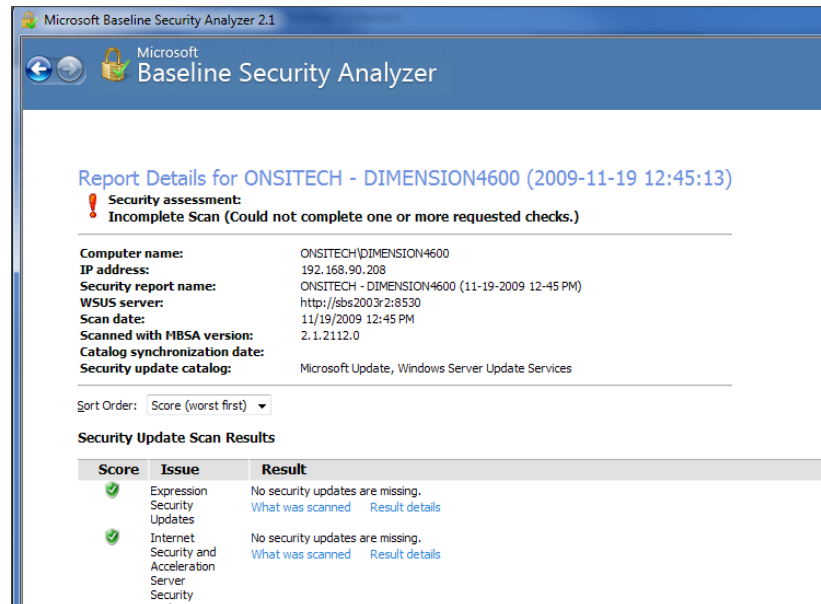
SCCM is an enterprise-grade tool designed to support end-to-end configuration management services, including operating system deployments, application installations, and update management. It requires a separately licensed instance of SQL Server 2005 or SQL Server 2008. For update management, SCCM uses a Software Update Point, which is a dedicated WSUS server customized to provide scanning functionality only but not package delivery. The package delivery is handled by SCCM. This allows the WSUS infrastructure to provide scanning services for substantially more client systems than can be provided when package delivery services must also be provided.

System Center Essentials (SCE) 2007

SCE is a product specifically designed for small and midsize organizations with less than 30 servers and 500 clients. It combines the features of WSUS with some of the features of System Center Configuration Manager and System Center Operations Manager. It provides inventory management capabilities and deployment of third party updates. Active Directory Domain Services is required. A SCE server can manage intra-forest systems in other domains, and multiple SCE servers can be deployed in a multi-domain forest but not more than one SCE server per domain.

Microsoft Baseline Security Advisor (MBSA)

The last component of a patch management process should be the use of the Microsoft Baseline Security Advisor. MBSA is designed for small and midsize organizations to evaluate the security status of their environment. It scans individual machines for missing security updates and other security-related configuration data, such as password policies, weak or missing passwords, and provides specific remediation guidance for such issues.



Microsoft Baseline Security Analyzer 2.1

Microsoft
Baseline Security Analyzer

Report Details for ONSITECH - DIMENSION4600 (2009-11-19 12:45:13)

Security assessment:
Incomplete Scan (Could not complete one or more requested checks.)

Computer name: ONSITECH\DIMENSION4600
 IP address: 192.168.90.208
 Security report name: ONSITECH - DIMENSION4600 (11-19-2009 12:45 PM)
 WSUS server: http://sbs2003r2:8530
 Scan date: 11/19/2009 12:45 PM
 Scanned with MBSA version: 2.1.2112.0
 Catalog synchronization date:
 Security update catalog: Microsoft Update, Windows Server Update Services

Sort Order: Score (worst first) ▼

Security Update Scan Results

Score	Issue	Result
✓	Expression Security Updates	No security updates are missing. What was scanned Result details
✓	Internet Security and Acceleration Server Security	No security updates are missing. What was scanned Result details

Practice Questions

Chapter 1

- You work as an enterprise engineer for a company that manufactures cables and wiring for the garage door industry. The company has a website and is implementing Active Directory for the first time. The website is located at www.CablesWiresAndMore.com and it is hosted by a third-party. The network administrator has asked you about the naming plan for the Active Directory implementation. You currently use a network operating system that does not depend on DNS. He wants to know if you can use the DNS service provided by the third-party hosting provider. Can you and why or why not? Select the best answer.

 - A. No. Because the hosting provider will not allow for dynamic DNS entries and will not want the extra load on the DNS servers.
 - B. No. Because no Internet DNS servers can support dynamic DNS.
 - C. Yes. Because ISPs and web hosters always support dynamic DNS.
 - D. Yes. Because a single DNS host name is all that is needed by Active Directory.
- You are implementing IPv4, but need to support IPv6 for the future. The dual-stack mode implemented in Windows Server 2008 will assist you with this. You must use WINS to support older Windows clients. In what modes can the Windows Server 2008 dual-stack operate? Choose all that apply.

 - A. IPv6-only enabled
 - B. IPv4 and IPv6 enabled
 - C. IPv4-only enabled
 - D. Neither IPv4 nor IPv6 enabled
- You are deploying applications in a Windows Server 2008 environment. All clients run Vista and all servers run Windows Server 2008. You are considering the use of SoftGrid for application virtualization. Terminal Services is already in place within the organization. What is a key difference between Terminal Service applications and SoftGrid applications? Select the best answer.

 - A. Terminal Services applications run on the client rather than on the server
 - B. SoftGrid application run on the client rather than on the server
 - C. SoftGrid is compatible with more applications
 - D. Terminal Services is compatible with more applications
- You manage the Terminal Servers for your organization. During a planning session, the IT Director asks you how you will handle a scenario where Terminal Server licenses are needed for devices that must actively use the server, but the licenses are all granted to other devices even though those devices are not currently active. You use TS CALs that are Per Device. The Director wants a solution that can be implemented in five minutes or less. What could you do in this scenario? Select the best answer.

 - A. Revoke 20% of the licenses
 - B. Configure the license server to handle 125% of the valid pool
 - C. Revoke all licenses
 - D. Purchase new licenses

5. You are implementing a Terminal Server. You begin by installing Windows Server 2008 Enterprise Edition. Next, you install thirteen applications that will be used by the users of the Terminal Server. Finally, you install the Terminal Server role and configure it to use a pre-existing Terminal Server License Server. After installation is complete, several of the thirteen applications do not work. These applications have been tested in Terminal Server environments in the past and they have always worked without modification. What should you change about your initial installation procedure to resolve this issue? Select the best answer.
- A. Install the Standard Edition of Windows Server 2008
 - B. Configure the Terminal Server to act as its own licensing server
 - C. Install the TS Application Compatibility toolkit
 - D. Install Terminal Server before installing the applications
6. You are deploying Terminal Server with a TS Gateway. The server running TS Gateway is a Windows Server 2008 Enterprise Edition server. It has 2 GB of RAM and two 3.4 GHz processors. In order to function as the TS Gateway, what two primary features must be installed and enabled on the server? Select the best answer.
- A. IIS and HTTP
 - B. IIS and RPC over HTTP
 - C. DHCP and DNS
 - D. IPSec and 802.1X

Chapter 2

1. You are attempting to install Active Directory on a Windows Server 2003 server. The server will participate in a Windows Server 2008 domain. The existing forest is running at the Windows Server 2008 functional level. When you attempt to install Active Directory on the Windows Server 2003 server, the installation fails. Why is the installation failing? Select the best answer.
- A. You cannot install Windows Server 2003 DCs in a Windows Server 2008 functional level forest.
 - B. The DNS is configured wrong.
 - C. The Windows Server 2003 server cannot connect to the Windows Server 2008 DC over the network.
 - D. You cannot install Windows Server 2003 DCs in a Windows Server 2008 functional level forest.
2. You are working with intra-forest authentication issues and attempting to plan an optimization strategy. Users in the research.sales.ohio.mycompany.com domain need to access resources in the research.sales.michigan.mycompany.com domain. Both domains exist in the same forest. Currently, the authentication process must traverse up the research.sales.ohio.mycompany.com branch and then down to the research.sales.michigan.mycompany.com child domain. How can you optimize this process? Select the best answer.
- A. Create a forest-level trust
 - B. Make the Domain Admins of research.sales.ohio.mycompany.com members of the research.sales.michigan.mycompany.com Domain Admins group
 - C. Create a shortcut trust
 - D. Create accounts in the research.sales.michigan.mycompany.com domain for each user in the research.sales.ohio.mycompany.com domain

3. Which one of the following administrative tasks is not part of the data management category when planning AD administration delegation? Select the best answer.
- A. Managing computer accounts
 - B. Adding and removing DCs
 - C. Managing security groups
 - D. Managing application-specific attributes for applications that integrate with the AD DS
4. You have installed a Windows Server 2008 server with the default installation options. You have also installed AD on the server. Every configuration setting is still configured as the default. You want the Audit Directory Service Access to be enabled on the server. What should you change to make this happen? Select the best answer.
- A. Enable the policy in the Default Domain Policy GPO
 - B. Enable the policy in the registry
 - C. Enable the policy in the ini file
 - D. Nothing
5. You are planning fine-grained password policies. You have created four PSOs that you need to apply to users in four different OUs. What should you do? Select the best answer.
- A. Assign the PSOs to the appropriate OUs
 - B. Create a distribution group for each OU that contains the users in that OU and then assign the appropriate PSO to that shadow group
 - C. Create a shadow group for each OU that contains the users in that OU and then assign the appropriate PSO to that shadow group
 - D. Install the appropriate PSO on the local computer for each user
6. What features of group policies should be used sparingly in order to remove confusion from GPO troubleshooting and maintenance? Choose all that apply.
- A. Blocking inheritance
 - B. Enforced
 - C. Permissions
 - D. Custom ADM templates

Chapter 3

1. You are planning the upgrade of an existing Windows Server 2003 domain. The current domain includes three DCs: 1 Windows Server 2003 Enterprise DC-2 Windows Server 2003 Standard DCs. You want all three servers to be upgraded to Windows Server 2008 standard edition to act as dedicated DCs for the new Windows Server 2008 AD DS domain. You need to accomplish the desired outcome with the least amount of effort. What actions should you take? Select the best answer.
- A. Perform an in-place upgrade on the two Windows Server 2003 Standard DCs. Perform a clean install on the Windows Server 2003 Enterprise DC and then add the AD DS services.
 - B. Perform an in-place upgrade on all Windows Server 2003 DCs.
 - C. Perform a clean install on all Windows Server 2003 DCs. Install a new domain.
 - D. Perform an in-place upgrade on the one Windows Server 2003 Enterprise DCs. Perform a clean install on the two Windows Server 2003 Standard DCs and then add the AD DS services.

2. You work as an enterprise administrator for a large organization. The organization operates three AD DS forests. All three forests run Windows Server 2008 DCs. The first forest includes three domains: prod.com, res.prod.com, and dev.prod.com. The second forest includes two domains: sales.internal and new.sales.internal. The third forest includes 1 domain: default.internal. The users in dev.prod.com need to access resources in sales.internal. No users in sales.internal need access to resources in dev.prod.com. How should you create the needed trust relationship while maintaining the most secure environment? Select the best answer.
- A. Create a one-way trust from dev.prod.com to sales.internal allowing users of sales.internal access to resources in dev.prod.com
 - B. Create a one-way trust from dev.prod.com to sales.internal allowing users of dev.prod.com access to resources in sales.internal
 - C. Create a forest level trust between the two forests so that all users from each forest can access resources in the other forest
 - D. No trust is needed. All forests automatically trust each other
3. The IT director has asked you to plan the migration of user accounts from one domain to another. The old domain is a Windows Server 2003 domain. The new domain is a Windows Server 2008 domain with all Windows Server 2008 DCs. He wants to ensure that user's passwords do not change during the migration. When you use the ADMT User Account Migration Wizard, what option should you choose on the Password Options page? Select the best answer.
- A. Transfer passwords
 - B. Generate complex passwords
 - C. Migrate passwords
 - D. Duplicate passwords
4. You are planning the steps involved in upgrading to Windows Server 2008. The current AD domain is a Windows Server 2003 R2 domain with all Windows Server 2003 R2 DCs. You are documenting the steps required to add a Windows Server 2008 DC to this domain. What is the first action you must take among the listed choices? Select the best answer.
- A. Adprep /rodcprep
 - B. Adprep /forestprep
 - C. Install the AD DS role on the Windows Server 2008 server
 - D. Perform an in-place upgrade on one of the non Operations Masters DCs
5. Contoso is in the process of upgrading its domains to Windows Server 2008. During this process, they wish to replace 41 Windows 2000 domain controllers that each have 256 MB RAM with 15 more powerful domain controllers. In order to accomplish this, more powerful hardware has been acquired. However, the resulting network will have fewer DCs servicing the same number of clients. The new DC servers have the following specifications: Dual 3.6 GHz physical processors 8 GB RAM 750 GB HD space 100 Mbps NIC. If only one component could be replaced to meet the demands of the clients, which component would it be? Select the best answer.
- A. RAM
 - B. NIC
 - C. Processor
 - D. Hard Drive

6. You must choose between an internal CA and an external third-party CA. The following requirements exist: 21 internal users need user certificates. These internal users must use the certificates with several partner organizations. You have 6 servers that need certificate for authentication to external networks. You choose an external certificate authority. Which of the following is not a disadvantage incurred when using an external CA? Select the best answer.
- A. Higher cost
 - B. Reduced flexibility
 - C. Autoenrollment is unavailable
 - D. Reduced compatibility with other organizations
7. You have installed Services for NFS on a Windows Server 2008 server; however, you are unable to access shares from a Unix machine. You did not perform any prerequisites before installing Services for NFS. After performing research, you determine that you should make changes to the AD schema to support Unix group ID and Unix user ID. How do you make this change? Select the best answer.
- A. By installing the Identity Management for Unix schema extension
 - B. By launching ADSIedit.exe and adding a value to the UnixGroupID and UnixUserID records
 - C. By installing Subsystem for Unix-based applications
 - D. By installing Services for NIS
8. Several users in your organization need to run older Windows applications that will not work on their current machines. The applications were designed for Windows 2000. All users run Vista Business edition. All servers run Windows Server 2008. You have a single dedicated Hyper-V server as well. Seven users need to run the application and 13 users need to access data created by the application. The data is stored in a SQL Server 2008 server. One of the Windows Server 2008 servers is also a Terminal Server. No additional software can be installed on the clients. Which one of the following solutions will work best to resolve the user's problems? Select the best answer.
- A. Run the application from the Windows Server 2008 Terminal Server
 - B. Create a Windows 2000 virtual machine in Virtual PC 2007 on the Vista client
 - C. Run the application in Windows 2000 mode on the client
 - D. Install a Windows 2000 server as a virtual server in Hyper-V and run Terminal Services on that server; allow the users to run the application from the virtualized Windows 2000 Terminal server

Chapter 4

1. You have 11 domain controllers in your environment. All domain controllers run Windows Server 2008 Enterprise edition. All clients run Vista Business. One domain controller has become corrupted and can no longer serve clients. You must restore AD DS on this domain controller. The rest of the operating system seems fine. What kind of restore will you perform? Select the best answer.
- A. Non-authoritative
 - B. Authoritative
 - C. Limited
 - D. Unlimited

2. You are planning business continuity. Specifically, you must design service availability for AD DS. You are considering the Operation Master roles. The network will contain a single Windows Server 2008 domain. All five roles will run on a single server. What must you do if the role server permanently fails? Select the best answer.
- A. Nothing. The domain will work fine without the roles
 - B. Seize only the Infrastructure role on another DC
 - C. Seize each of the roles on one or more other DCs
 - D. Seize only the Schema Master role on another DC
3. You are estimating licensing costs for a virtual infrastructure based on Hyper-V. You can choose among Standard edition full install, Standard edition Server Core install, Enterprise edition and Data Center edition of Windows Server 2008. You need a maximum of 3 guest VMs on each host server and all host servers will have at least 2 VMs. Which edition will serve your needs best when considering the cost? Select the best answer.
- A. Datacenter
 - B. Enterprise
 - C. Standard full install
 - D. Standard Server Core
4. You need to utilize BitLocker for drive encryption on several machines including three Windows Server 2008 servers and 18 Vista Ultimate clients. The specs for the machines are as follows:
- The three servers support the Trusted Platform Module version 1.0
 - The 18 Vista clients support the Trusted Platform Module 1.2
 - You cannot store information on a USB key module, but it must be stored on the TPM.
- What must you do to use BitLocker on these machines?
Select the best answer.
- A. Replace the 18 clients with machines supporting TPM 1.0
 - B. Do nothing; Bitlocker works on either 1.0 or 1.2 TPM systems
 - C. Replace all 21 machines with machines that support TPM 1.5
 - D. Replace the three servers with machines supporting TPM 1.2
5. You have implemented a Windows Server 2008 domain. Currently you manage 11 domain controllers in a single domain. 27 other servers are used for various purposes including SQL Server 2000, SQL Server 2005, Exchange Server 2007, IIS 7 and AD RMS. You want to use AD RMS to provide persistent protection of data even as it leaves the organization through email attachments or theft. You have 1137 clients running Windows 2000 Professional, Internet Explorer 5.5 and Microsoft Office 2000. You want the users that access these clients to be able to access files managed and secured through AD RMS. What action will you need to take in order for this to work? Select the best answer.
- A. Upgrade to Internet Explorer 7.0
 - B. Upgrade to Office XP (2002) and Internet Explorer 6.0
 - C. Upgrade to at least Office 2003 and Internet Explorer 6.0
 - D. Upgrade to Office 2003

Answers & Explanations

Chapter 1

1. ANSWER: A

Explanation A. Correct. Hosting providers do not usually provide dynamic DNS for security reasons. Additionally, the extra load on their DNS servers would be too much. You will need to host your own DNS zones internally.

Explanation B. Incorrect. This statement is incorrect because Internet DNS servers certainly can support dynamic DNS. They simply choose not to in most situations for security reasons.

Explanation C. Incorrect. Dynamic DNS is rarely supported by web hosters.

Explanation D. Incorrect. Active Directory uses server host names, but it also uses dozens of service record (SRV) entries.

2. ANSWERS: A, B, C

Explanation A. Correct. In this mode, only IPv6 communications occur.

Explanation B. Correct. In this mode, both versions of IP are used for communications.

Explanation C. Correct. In this mode, only IPv4 communications occur.

Explanation D. Incorrect. This would be a state of non-IP communications and the dual-stack would not exist.

3. ANSWER: B

Explanation A. Incorrect. SoftGrid application run on the client; Terminal Services applications run on the server.

Explanation B. Correct. Terminal Services provides presentation virtualization only. Applications run on the server, but are presented to the client. SoftGrid allows the applications to run on the client in a virtual environment.

Explanation C. Incorrect. In a like environment, SoftGrid is compatible with the same applications as Terminal Services.

Explanation D. Incorrect. In a like environment, SoftGrid is compatible with the same applications as Terminal Services.

4. ANSWER: A

Explanation A. Correct. You can revoke 20% of the licenses and these would become available to the devices needing immediate access. Licenses expire randomly between 52 and 89 days. This revocation action should only be taken as a temporary solution. More licenses will likely need to be purchased.

Explanation B. Incorrect. No such option is available.

Explanation C. Incorrect. This would cause unnecessary processor and network utilization. Revoking 20% is the Microsoft recommended practice.

Explanation D. Incorrect. Purchasing new licenses may be a long term solution, but it could not likely be done in 5 minutes or less.

5. ANSWER: D

Explanation A. Incorrect. There is no difference between Windows Server 2008 Standard and Enterprise from a Terminal Services application compatibility standpoint.

Explanation B. Incorrect. The new Terminal Server should be able to interoperate with the pre-existing licensing server.

Explanation C. Incorrect. No such toolkit exists.

Explanation D. Correct. When you install the Terminal Server role first, you have the ability to install application properly for use via Terminal Services.

6. ANSWER: B

Explanation A. Incorrect. HTTP is available on all servers by default. IIS is required.

Explanation B. Correct. IIS provides the web page hosting for the launching point to Terminal Servers and RemoteApps. RPC over HTTP allows for the communications between the client and the TS Gateway server.

Explanation C. Incorrect. DHCP and DNS would rarely be installed on the TS Gateway server.

Explanation D. Incorrect. These items are enforcement techniques used with NAP.

Chapter 2**1. ANSWER: D**

Explanation A. Incorrect. As long as the forest is not running Windows Server 2008 functional level, you can install Windows Server 2003 DCs.

Explanation B. Incorrect. Since there is an existing domain, DNS is most likely operational.

Explanation C. Incorrect. Even if it can connect, the installation will fail because the Windows Server 2008 forest is in the Windows Server 2008 functional level.

Explanation D. Correct. Windows Server 2008 functional level forests only support Windows Server 2008 DCs.

2. ANSWER: C

Explanation A. Incorrect. Forest-level trusts are used to establish trusts between forests.

Explanation B. Incorrect. This action will not change the current authentication and authorization processes.

Explanation C. Correct. Shortcut trusts are used to create direct trusts between nested domains in the same forest in order to avoid multi-point authentication requirements.

Explanation D. Incorrect. These accounts would simply be separate accounts in the remote domain.

3. ANSWER: B

Explanation A. Incorrect. This is part of data management. Computer accounts are simply data entries in the directory database.

Explanation B. Correct. Adding and removing DCs is part of the service management category.

Explanation C. Incorrect. This is part of data management. Security groups are simply data entries in the directory database.

Explanation D. Incorrect. This is part of data management. Applications that integrate with the AD DS store data in the directory store. Exchange Server is such an application.

4. ANSWER: D

Explanation A. Incorrect. It is enabled by default.

Explanation B. Incorrect. The policy is properly configured using GPOs, but it is enabled by default.

Explanation C. Incorrect. Very few Microsoft solutions use ini files today. The setting is managed in a GPO and it is enabled by default.

Explanation D. Correct. This policy is enabled by default. Directory service events will be logged in the Event logs.

5. ANSWER: C

Explanation A. Incorrect. PSOs must be assigned directly to users or through security groups.

Explanation B. Incorrect. Distribution groups cannot be used to assign PSOs.

Explanation C. Correct. PSOs cannot be assigned directly to OUs.

Explanation D. Incorrect. PSOs must be deployed through AD.

6. ANSWERS: A, B

Explanation A. Correct. This feature is useful, but if it is heavily used it will make troubleshooting very complicated. Heavy use of this feature usually indicated poor planning.

Explanation B. Correct. This feature is useful, but if it is heavily used it will make troubleshooting very complicated. Heavy use of this feature usually indicated poor planning.

Explanation C. Incorrect. Permissions are used to filter the GPOs and, as long as they are well documented, they will be used commonly.

Explanation D. Incorrect. Custom ADM templates are somewhat complex to create and implement, but once implemented they do not cause troubleshooting problems as they are apparent in group policy analysis reports.

Chapter 3

1. ANSWER: A

Explanation A. Correct. You cannot upgrade from Windows Server 2003 Enterprise to Windows Server 2008 Standard. You cannot downgrade editions while upgrading the operating system.

Explanation B. Incorrect. Since one DC is running Enterprise edition, you must clean install that DC in order to achieve the desired results of running Windows Server 2008 Standard on all DCs.

Explanation C. Incorrect. This would not be an upgrade and all user accounts and AD objects would have to be recreated.

Explanation D. Incorrect. This option is the opposite of what is needed. You should perform an in-place upgrade on the two Windows Server 2003 Standard DCs. Perform a clean install on the Windows Server 2003 Enterprise DC and then add the AD DS services.

2. ANSWER: B

Explanation A. Incorrect. This solution is reversed from that which is needed.

Explanation B. Correct. You can create uni-directional or bi-directional trusts. In this case a uni-directional external trust should be created.

Explanation C. Incorrect. While you can do this, it would not result in the most secure environment.

Explanation D. Incorrect. All domains within a forest automatically trust each other through some trust path; however, no trusts exist between separate forests by default.

3. ANSWER: C

Explanation A. Incorrect. This is not an option.

Explanation B. Incorrect. This would generate new passwords.

Explanation C. Correct. This option will allow the users' passwords to remain the same.

Explanation D. Incorrect. This is not an option.

4. ANSWER: B

Explanation A. Incorrect. This command would be run after adprep /forestprep.

Explanation B. Correct. This command must be run on the Schema Master domaincontroller.

Explanation C. Incorrect. You must prepare the schema first.

Explanation D. Incorrect. While you could take this action, you must prepare the schema first.

5. ANSWER: B

Explanation A. Incorrect. The RAM is likely sufficient.

Explanation B. Correct. Since we are supporting the same number of clients with 15 NICs instead of 41 NICs, we should use a Gigabit NIC in each server.

Explanation C. Incorrect. The processors should be sufficient.

Explanation D. Incorrect. 750 GB is more than enough for a DC.

6. ANSWER: D

Explanation A. Incorrect. External CAs charge monthly or annual fees for using their services and costs are higher.

Explanation B. Incorrect. The external CA will not allow you to directly manage certificates at the same level as an internal CA would.

Explanation C. Incorrect. It is true that external CAs do not provide autoenrollment.

Explanation D. Correct. In fact, this is the primary motivator for selecting an external CA. The external CA can authenticate your certificate for any user in the world assuming it is a well known CA.

7. ANSWER: A

Explanation A. Correct. The Identity Management for Unix schema extension must be installed before Services for NFS will work properly.

Explanation B. Incorrect. No such records exist.

Explanation C. Incorrect. Installing this subsystem does not modify the schema in any way.

Explanation D. Incorrect. Installing Services for NIS allows management of UNIX NIS from AD, but it does not modify the schema for Services for NFS.

8. ANSWER: D

Explanation A. Incorrect. Most applications that do not work on Vista also do not work on Windows Server 2008.

Explanation B. Incorrect. This would require the installation of Virtual PC 2007 on the clients.

Explanation C. Incorrect. Vista does not provide a Windows 2000 mode.

Explanation D. Correct. Since the application was written for Windows 2000 and does not work on Windows Vista, it will not likely work on a native Windows Server 2008 Terminal Server. It will, however, most likely work on a Windows 2000 Terminal Server.

Chapter 4**1. ANSWER: A**

Explanation A. Correct. A non-authoritative restore will receive updates from the other DCs once it is brought online, but it will not force its data onto the other DCs.

Explanation B. Incorrect. An authoritative restore would result in the backup data overriding existing valid data on the other 10 DCs.

Explanation C. Incorrect. This is not a valid restore option when restore AD DS.

Explanation D. Incorrect. This is not a valid restore option when restoring AD DS.

2. ANSWER: C

Explanation A. Incorrect. Eventually you will experience problems. For example, without the PDC role, password changes could become problematic when employing administrative resets.

Explanation B. Incorrect. All five roles are essential.

Explanation C. Correct. The roles are essential to the service. Eventually you will have to seize the roles onto another DC.

Explanation D. Incorrect. All five roles are essential to smooth operations.

3. ANSWER: B

Explanation A. Incorrect. While Datacenter allows an unlimited number of VMs without extra licensing, Enterprise edition is sufficient and costs less.

Explanation B. Correct. Enterprise edition allows for up to four VMs running Windows Server 2008 without extra licensing.

Explanation C. Incorrect. Standard edition does not provide any licenses for guest OS VMs.

Explanation D. Incorrect. Standard edition does not provide any licenses for guest OS VMs.

4. ANSWER: D

Explanation A. Incorrect. Version 1.2 is required. These client machines are fine.

Explanation B. Incorrect. BitLocker only works on version 1.2 TPM systems.

Explanation C. Incorrect. At the time of Windows Server 2008 release, TPM was not on version 1.5. Version 1.2 is the required version.

Explanation D. Correct. Version 1.2 must be supported otherwise USB storage is required. This scenario explicitly forbids USB storage.

5. ANSWER: C

Explanation A. Incorrect. While this would make Internet Explorer compatible with AD RMS, it would not resolve the Office compatibility issue. The clients must be upgraded to a minimum of Office 2003.

Explanation B. Incorrect. While this would make Internet Explorer compatible with AD RMS, it would not resolve the Office compatibility issue. The clients must be upgraded to a minimum of Office 2003.

Explanation C. Correct. Office 2003 is the lowest version of Office that supports AD RMS. Internet Explorer 6.0 supports AD RMS through the RM add-on for Internet Explorer.

Explanation D. Incorrect. While this would make Office compatible, Internet Explorer would still be incompatible. Internet Explorer must be at version 6.0 at a minimum.