Study guide by ExamNotes.net

Exam 70-216

Implementing, Managing and Supporting W2K Network Infrastructure

Abstract

This ExamNotes Study Guide intends to provide you with information to prepare for the Microsoft W2K 70-216 Exam.

ExamNotes Study Guide Topics Covered

- DNS and DDNS
- DHCP
- Remote Access
- Routing
- TCP/IP and Other Protocols
- Authentication
- Security
- WINS

Before you start

This study guide provides you with information on the many different aspects of “W2K Network Infrastructure”. You should not use this information as your first step into Networking, as this exam is targeted towards candidates with real experience and solid background on Windows Networking and TCP/IP. If you are a beginner, I recommend that you first study the material presented in the NT 4.0 track Networking Essential and TCP/IP, and then complete the 210 and 215 exams before working on this one.

DNS

Domain Name System (or Service) is an Internet service that translates domain names into IP addresses. The Internet is really based on IP addresses, but because domain names are alphabetic, they're easier to remember.

Every time a user use a domain name, DNS service must translate the name into the corresponding IP address. For example, the domain name www.myexample.com might translate to 198.105.222.2.

The DNS system is a distributed system. If one DNS server doesn't know how to translate a particular domain name, it asks another one, and so on, until the correct IP address (or an error) is returned.

InterNIC Internet Network Information Center controls top-level domains as follow:

Name	**DEscription**
Com	Commercial organizations
Edu	Educational institutions
Org	Non-profit organizations
Net	Networks or ISP
Gov	Government organizations
Mil	Military government organizations
Num	Phone numbers
Arpa	Reverse DNS

The three components of DNS are: resolvers, name servers, and the domain name space. Resolver is client that sends queries to a name server. The name server returns the requested information, a pointer to another name server, or an error message.

Computers in a W2K network use DNS for EVERYTHING, from name resolution to locating domain controllers for logon. In the past you can use hosts file or lmhosts file, but now you basically won't b able to setup a smooth running W2K network without a properly set up DNS server.

A name server is an Internet host running software capable of processing DNS requests. A popular free software name server is BIND, although the MS implementation is completely different.

Generally, a single name server will be configured as the primary name server for a domain. You make all the changes to the DNS database in the Primary Server. For backup purposes, a number of other name servers may be configured as secondary name servers. Regarding load sharing, the

resolving algorithm simply uses a domain's NS records in the order provided. Typically, the primary name server is listed first, followed by the secondaries, but this is not absolute. A client is free to set his/her own resolving order.

Typically, secondary name servers do not have to be known to the primary. However, we prefer to restrict zone transfers to preconfigured hosts for security purpose - Secondary name servers depend on zone transfers to acquire a complete copy of the primary's database. This action is triggered by a change in the serial number - the SERIAL field in a domain's SOA record must be changed every time a change is made within the domain.

If a name server receives a query it does not serve, it may return a referral to the client citing better name servers, or may recurse by attempting to completely resolve the request through a series of exchanges with other name servers, delaying a reply to the original requester until it is complete.

Most name servers will recurse, as this will permit them to cache the various resource records used to access the foreign domain, in anticipation of further similar requests. Note that every resource record has a Time To Live field which specifies the number of seconds the record may be cached before it must be discarded.

If the DNS server simply caches data rather than to hold a copy of the DNS database, it is a caching only server for performance purpose.

In W2K, 4 types of zones are supported:

- Active Directory-integrated
- Standard Primary
- Standard Secondary
- Caching-only

DHCP

Short for Dynamic Host Configuration Protocol, DHCP is a protocol for assigning dynamic IP addresses to devices on a network. Why dynamic addressing? A device can have a different IP address every time it connects to the network. In a Windows network, the device's IP address can even change while it is still connected, using the IPCONFIG or WINIPCFG command. This simplifies network administration because the software keeps track of IP addresses rather than requiring an administrator to manage the task.

DHCP client support is built into Windows 95/98/NT/2000.

Each time a DHCP client starts to run, it will request IP address information from a DHCP server. Information can include:

- IP address
- Subnet mask
- Default gateway address
- Domain Name System address
- Windows Internet Name Service server address

Administrator should have setup an address pool first. If there is no available IP address in the pool to lease to a client, the client cannot initialize TCP/IP. Fortunately, a Windows 2000-based clients can automatically configure an IP address and subnet mask through Automatic Private IP Addressing (APIPA) by using a selected address from reserved Class B network, 169.254.0.0, with the subnet mask 255.255.0.0. The DHCP client can even test for address conflicts and retry autoconfiguration for up to 10 addresses, avoiding network disruption. This does not end the story. The client will continue to check for a DHCP server every 5 minutes. If finally a DHCP server is found, the client will use the formal DHCP address.

All SERVERS are recommended to use static IP addresses themselves. In W2K DHCP, you must authorize a DHCP Server in the Active Directory service before it can lease addresses out. In fact, if Active Directory is deployed, which is highly likely, all computers operating as DHCP servers must be either domain controllers or domain member servers.

A DHCP scope is a pool of valid IP addresses available for lease to clients. One scope for every DHCP server must be created, and only one scope can be assigned to a specific subnet. Different DHCP servers do not share information, so you must make sure there is no conflicting configuration.

Dynamic DNS - DHCP and DNS Integration

A W2K DHCP server can register with a DNS server and update records on behalf of its DHCP-clients using the Dynamic DNS update protocol. DHCP option code 81 enables the return of a client's fully qualified domain name to the DHCP server.

In DNS Property, you can select “Enable Updates for DNS Clients That Do Not Support Dynamic Update” to involve non-supported client. Also, if your zone type Is Active Directory-Integrated, you should select to accept Only Secure Updates for security purpose, as you don’t want someone to spoof a DNS and update you with incorrect information.

Most DHCP-related problems come from mis-configuration at a client. If you suspect this is a server problem, check the system event log and DHCP server audit logs.

DHCP Relay Agent

A relay agent is a software program that relays DHCP/BOOTP messages between clients and servers on different subnets. You need a relay agent for each IP network segment that contains DHCP clients but NO DHCP server. The reason is that, router block broadcast between subnets, making it impossible for clients in different subnets from asking for leases.

In W2K configuration, DHCP Relay Agent is part of the Routing Protocol.

Routing and Remote Access

Remote Access Policy RAP basically defines who has remote access to the network, and what characteristics will apply. You may accept or reject connections based on day and time, group membership, and type of service. The point is, you need to create policies manually on each server, with the file name IAS.MDB. Also note that you cannot Control Access Through Remote Access Policy on mixed-mode domain controllers.

Every Remote Access Profile has 6 different areas are Dial-in Constraints, IP, Multilink, Authentication, Encryption, and Advanced.

In networking, Routing is the process of moving a packet from source to destination. Normally routing is performed by a dedicated device called a router, but in a W2K server routing is possible via NIC and remote access.

Routing enables messages to pass from one computer to another and eventually reach the target machine. Each intermediary computer performs routing by passing along the packet to the next computer. This involves analyzing a routing table to determine the best path. Please note that even nonrouter host has a routing table that is used to determine the optimal route.

We have three different types of entries in the routing table, namely network route, host route, and default route. Use the command mode command ROUTE to manipulate the routes. Keep in mind that, once you reboot your machine, it lost all the routing entries. Entries will have to be rebuilt again.

To make routers able to self coordinating in terms of managing route information, you use routing protocol. RIP Routing Information Protocol is a protocol defined by RFC 1058 that specifies how routers exchange routing table information. With RIP, routers periodically exchange entire tables. Because RIP is relatively inefficient, you want to restrict its usage in small to medium size network.

RIP is a distance-vector routing protocol provided mainly for backwards-compatibility. RIP uses broadcasts to share information with neighboring routers, and sends periodic RIP broadcast packets containing all routing information known to the router. The broadcasts keep all routers synchronized at the expense of bandwidth consumption.

OSPF Open Shortest Path First is a link-state routing protocol that enables routers to exchange routing information and create a map of the network that calculates the best possible path to each network. Its demand for memory and route computation times is high, but is much more scalable than RIP.

With a Link State protocol, each router sends that portion of the routing table that describes the state of its own links, and it also sends the complete routing structure. The advantage of shortest path first algorithms is that they results in smaller more frequent updates everywhere and quicker converge.

Most of the routers are running all the time. However, there is a special type of routing mechanism which works on a per demand basis. A demand-dial interface is a router interface that will be brought up on demand based on network traffic. The demand-dial link is only initiated if the routing table shows that this interface is needed. The primary advantage of on demand link is cost saving.

VPN, which stands for virtual private network, is a network that is constructed by using public wires to connect nodes. For example, there are a number of systems that enable you to create networks using the Internet as the medium. These systems use encryption and other security mechanisms to ensure that only authorized users can access the network and that data cannot be intercepted by hackers. PPTP Point to Point Tunneling Protocol is an advanced functionality in TCP/IP that allows remote access via “tunnels”, exactly the same as running a VPN.

Routing and Remote Access can use DHCP to lease addresses in blocks of 10, and can store them in the local registry. Client can receive IP address from the Routing and Remote Access server, and may use DHCPINFORM packets to obtain WINS or DNS addresses, domain name, or other options, without getting an IP address.

RADIUS, short for Remote Authentication Dial-In User Service, is an authentication and accounting system used for remote access management. When a user dial in to he/she must enter the username and password. This information is passed to a RADIUS server, which checks that the information is correct, and then authorizes access to the system. With RADIUS, all authentication requests heard by a server are sent to a RADIUS server. By default, the Authentication provider is Windows Authentication. You may, however, change to RADIUS authentication using Internet Authentication Service. IAS is used for centralized administration, and enforcement of access policies. It supports PAP, CHAP, MS-CHAP, and EAP.

IAS create log files in located in the %system-root%\system32\LogFiles folder based on authentication and accounting requests received from NAS. You use these logs to track accounting information for billing purposes. Log files can be started daily, weekly, monthly, or when the log reaches a specific size.

TCP/IP

TCP/IP is installed as the default protocol in W2K. Here are the basic terms to get you started.

- Address is the unique number id assigned to one host or interface in a network
- Subnet is a portion of a network sharing a particular subnet address
- Subnet mask is a 32-bit combination used to describe which portion of an address refers to the subnet and which part refers to the host

To go into the internet, you must have already received your legitimate address(es) from the InterNIC - Internet Network Information Center.

To determine the class of an address, look at the first octet of the dotted-decimal address.

Class A: 1 - 126
Class B: 128-191
Class C: 192-223

In a class A address, the first octet is the network portion. Octets 2, 3, and 4 are for you to divide into subnets and hosts if needed. Class A addresses are used for networks that have up to 16,581,375 hosts.

In a class B address, the first two octets are the network portion. Octets 3 and 4 are for local subnets and hosts. Class B addresses are used for networks that have between 256 and 65,536 hosts.

In a class C address, the first three octets are the network portion. Only Octet 4 is for local subnets and hosts. This is good for networks with less than 256 hosts.

Subnetting allows you to create multiple logical networks that exist within a single Class A, B, or C network. If you don't subnet, you will only be able to use one network from your Class A, B, or C network. Through subnetting,

you can divide your network into smaller broadcast domain and increase network performance. Of course, you need to have routers in between these subnets for them to communicate.

By default, W2K computers will attempt to obtain the TCP/IP configuration from a DHCP server. However, you may opt for a static TCP/IP configuration, especially for servers such as DHCP, DNS, and WINS. If you do not have a DHCP server, you must configure TCP/IP computers manually to use a static IP address machine by machine.

You perform basic TCP/IP configuration and connectivity testing using IPConfig and ping utilities. IPConfig verifies the TCP/IP configuration parameters on a host, including the IP address, subnet mask, and default gateway. It also allows you to release and renew IP addresses on the fly. Ping tests TCP/IP configurations and diagnoses connection failures. Since Ping uses ICMP to determine whether a particular TCP/IP host is available and functional, if the target is set to ignore ICMP, you will still get an error message.

For security purpose, you may use IP packet filtering to invoke security negotiations for communication based on source, destination, and type of IP traffic, plus deciding what to block and what to allow. IP packets can be filtered based on TCP port number, UDP port number, and IP protocol number.

NetBEUI

Pronounced net-booey, Netbeui was originally designed by IBM for their Lan Manager server and later extended by Microsoft and Novell. Short for NetBios Enhanced User Interface, it is an enhanced version of the NetBIOS protocol used by network operating systems such as LAN Manager, LAN Server, Windows for Workgroups, Windows 95 and Windows NT.

NetBEUI is easy to implement, as it does not need any special configuration.

Any computer with a computer name (or NETBIOS name) can communicate. However, it mainly uses broadcast, which results in poor scalability and poor performance.

In the NT 4.0 era, Netbeui is the default protocol. It is no longer the case in W2K.

NWLINK

Netware is a popular network operating system developed by Novell Corporation. It runs on a variety of different types of LANs, from Ethernets to IBM token-ring networks. It provides users and programmers with a consistent interface that is independent of the actual hardware used to transmit messages. IPX, which stands for Internetwork Packet Exchange, is a networking protocol used by the NetWare.

You use Gateway Service for NetWare or Client Services for NetWare to connect to NetWare servers. You may also use Client Services for NetWare or Novell Client for W2K to log on to a NetWare network from a W2K Professional PC. The key is, you must install the NWLINK protocol, which is the MS implementation of Novell IPX. Also, to install Client Services for NetWare, you need Administrator rights to the computer running W2K. Unattended Installation can deal with large deployments of W2K Professional and Client Services for NetWare.

Starting from Netware 4.11 you may use TCP/IP to communicate. For the exam, however, it is better for you to insist on NWLINK.

Security Infrastructure

Authentication is the process of identifying an individual, usually based on a username and password. Please note that authentication is distinct from authorization , which is the process of giving individuals access to system

objects based on their identity. Authentication merely ensures that the individual is who he or she claims to be, but has nothing to deal with the access rights.

Kerberos is an authentication system developed at the Massachusetts Institute of Technology. Kerberos enables two parties to exchange private information across the network through assigning a unique key, called a ticket, to each user that logs on to the network. The ticket is embedded in messages to identify the sender of the message.

W2K uses Kerberos as one of the authentication methods. In fact, Kerberos V5 is the default authentication technology used for any clients running the Kerberos V5 protocol, whether or not they are Windows-based clients, as long as they are members of a trusted domain.

Public Key Encryption is a cryptographic system that uses two keys -- a public key known to everyone and a private or secret key known only to the recipient of the message. An important element to the public key system is that the public and private keys are related in such a way that only the public key can be used to encrypt messages and only the corresponding private key can be used to decrypt them. In addition, it is virtually impossible to deduce the private key if you know the public key. The drawback of this method is speed – it is very CPU power consuming, and is very slow.

Digital Certificate is an attachment to an electronic message used for security purposes. The most common use of a digital certificate is to verify a user's identity and to provide the receiver with the means to encode a reply. An individual who need to send an encrypted message applies for a digital certificate from a Certificate Authority CA. The CA issues an encrypted digital certificate containing the applicant's public key and a variety of other identification information. The CA makes its own public key readily available on the Internet. W2K Server's Certificate Service can act as an internal CA.

Recipient of an encrypted message uses the CA's public key to decode the

digital certificate, verifies it and then obtains the sender's public key and identification information held within the certificate. With this information, the recipient can send an encrypted reply.

W2K can use certificates as an authentication method. This requires that at least one trusted CA has been configured. Certificate version supported by W2K is X.509 V3 certificates.

W2K can also use Preshared Key for authentication. Preshared key is a shared key that is secret and is previously agreed on by two users. It does not require the client to run the Kerberos protocol or have a public key certificate. And most importantly, it is much faster than to use Public key encryption.

To use Preshared key, both parties must manually configure IPSec. We often use this method for authenticating non-Windows-based hosts.

You may use a rule to specify multiple authentication methods, in order to ensure that a common method can be found when negotiating with a peer. A rule actually contains a list of IP tilters and specifies the security actions that will take place when there is a match.

IPSec, which stands for IP Security, is a set of protocols being developed by the IETF to support secure exchange of packets at the IP layer. Once it's completed, IPsec is expected to be deployed widely to implement VPNs. It supports two encryption modes: Transport and Tunnel. Transport mode encrypts only the data portion of each packet and leaves the header untouched. Tunnel mode encrypts both the header and the data, which is more secure.

For IPSec to work, on the receiving side there must be an IPSec-compliant device to decrypt each packet. Also, the sending and receiving devices must share a public key. This is usually accomplished through a protocol known as Internet Security Association and Key Management Protocol/Oakley

(ISAKMP/Oakley), which allows the receiver to obtain a public key and authenticate the sender using digital certificates.

An IPSec policy is a collection of rules and key exchange settings being assigned as a domain security policy or an individual computer's security policy. A domain computer will automatically inherit the IPSec policy assigned to the domain security policy when it logs on as a domain member. One security policy can be created for all users on the same network. If the computer is not connected to a domain, IPSec policies are stored in the local registry.

You can view the default IP Security policies in the Group Policy snap-in. The policies are listed under IP Security Policies on Active Directory: Group Policy Object\Computer Configuration\Windows Settings\Security Settings\IP Security Policies on Active Directory. Note that each IPSec policy is governed by rules that determine when and how the policy is applied. The three predefined policy entries are Client (Respond Only), Secure Server (Require Security), and Server (Request Security). By default, none of them are enabled.

We use IPSECMON.EXE to monitor IP SAs, rekeys, negotiation errors, and other IP Security statistics. Network Monitor, on the other hand, captures all information transferred over a network interface at any given time. Version 2.0 contains parsers for IPSec packets, making the packet itself visible (not the content, of course).

NAT

NAT stands for Network Address Translation. It is an Internet standard that enables a local-area network to use one set of IP addresses for internal traffic and a second set of addresses for external traffic. NAT located where the LAN meets the Internet makes all necessary IP address translations. We mainly use NAT to provide a type of firewall by hiding internal IP addresses, to enable a company to use more internal IP addresses, and to allow a company to combine multiple ISDN connections into a single Internet connection.

WINS

WINS stands for Windows Internet Naming Service, which is a system that determines the IP address associated with a particular network computer.

WINS supports network client and server computers running Windows and can provide name resolution for other computers with special arrangements. WINS uses a distributed database that is automatically updated with the names of computers currently available and the IP address assigned to each one. This automatic maintenance feature is why WINS was so popular in NT4.0. Apart from the "auto" method, an administrator can manually use the WINS console or command-line tools to add or delete statically mapped entries in the server database.

With DDNS, we can achieve auto maintenance with DNS too, so WINS become a second choice mainly for backward compatibility purpose.

Replication enables a WINS server to resolve NetBIOS names of hosts registered with another WINS server. To replicate WINS database entries, each WINS server must be configured as either a pull or a push partner with at least one other WINS server.

WINS console provides backup tools for you to back up and restore WINS database. When backing up the server database, \Wins bak\New folder under the backup folder you have specified as the Default backup path in Server Properties is created. By default, WINS performs complete database backups every three hours using the specified folder. You may also configure WINS to back up the database automatically when the service is stopped or the server is shut down.