Last updated August, 2000.
Click here for updates.

Click here to see additional documents related to this study guide.

Contents

Cramsession™ for Designing a Microsoft Windows 2000 Network Infrastructure

Abstract:

This Cramsession will help you to prepare for Microsoft exam 70-221, Designing a Microsoft Windows 2000 Network Infrastructure. Exam topics include Network Topology, Routing, IP Addressing, Name Resolution, Virtual Private Networks, Remote Access, and Telephony Solutions.

Designing a Microsoft Windows 2000 Network Infrastructure

Deployment cycle:

Design:

Where decisions are made. Requires knowledge of existing network infrastructure and organizational goals. You will need to choose which services to implement and how to combine them to increase performance and simplify management. Also designed at this stage is the management strategy, which specifies how the network will be managed on a day-to-day basis.

Implement:

Takes place after the network design has been properly tested. The network is configured as specified in the design, and monitoring is setup to collect data on its performance.

Manage:

Using the performance data you have collected bottlenecks are identified and removed to enable changes needed to maintain the network within design specifications.

Organizational goals to include in a design:

Functionality	Forms the basic reasoning for implementing a network service. A DHCP server that simplifies network administration by dynamically assigning IP addresses is an example of functional design.
Security	Data is only considered secure when access to confidential data is limited strictly to authorized users. Be aware that implementing security may affect availability and performance.
Availability	Calculated by measuring the percentage of time users have access to a service.
Performance	Based on response times as specified by an organization’s goals.

Important building blocks:

TCP/IP:

This is an open, industry-standard, and routable protocol. It is required for many essential Windows 2000 network services such as DHCP, WINS, DNS, and Active Directory. TCP/IP should be used in heterogeneous environments and whenever Internet connectivity is called for as part of the design.

DNS:

Domain Name Service - resolves fully qualified domain names (FQDN) to IP addresses. Allows network admins to assign "people-friendly" names to network resources. Windows 2000 Active Directory is based entirely on the hierarchical structure of the DNS namespace.

DHCP:

Dynamic Host Configuration Protocol - used to dynamically assign Internet Protocol (IP) addresses to clients and reduce administrative overhead in managing and maintaining a TCP/IP-based network.

WINS:

Windows Internet Name Service - a NetBIOS name service that resolves NetBIOS names to IP addresses in a Windows network. Required for Windows 3.11/95/98/NT4 clients that do not have the Active Directory client installed.

MS Proxy Server 2.0: (KB# Q164084)

This is a combination firewall/proxy server product that provides security by allowing organizations to control the exchange of data between the Internet and their private network. Can also be used to improve the performance of Internet access through its content caching features. It is extremely scalable and suitable for enterprise type deployment within an organization.

NAT:

Network Address Translation – a protocol found in Routing and Remote Access Services (RRAS) in Windows 2000. Used to provide Internet connectivity in simple network environments where all machines are on a single subnet. Provides some security.

IP Routing:

Windows 2000 RRAS supports both static and dynamic routing protocols. Connections over non-persistent links are supported through demand-dial routing.

Remote Access:

Used to allow remote users access to a private network. Can include dial-up connections over the regular telephone system and also Virtual Private Network (VPN) connections over the Internet.

RADIUS:

Remote Authentication Dial-In User Service - provides authorization, authentication, and accounting services for distributed dial-up networks. Used in conjunction with RRAS and IAS (Internet Authentication Service).

TCP/IP:

IP addressing and subnetting: (KB# Q186341 & RFCs: 950, 1518, 1519, 1812, & 1878)

Public addressing schemes: (Tutorial)

All hosts connected to the public Internet require a globally unique IP address. Any network connected to the Internet must have a minimum of one public IP address for connectivity. Used when the organization has a large number of hosts requiring direct Internet access and there is a sufficient pool of registered addresses to work from. Public addressing schemes are expensive and limit network growth as once all available addresses have been exhausted, no new devices can be added to the network unless more IP addresses are purchased.

Address class	Range	Default mask	Used for
A	1 – 126.x.x.x	255.0.0.0	Host/network
B	128 – 191.x.x.x	255.255.0.0	Host/network
C	192 – 223.x.x.x	255.255.255.0	Host/network
D	224 – 239.x.x.x	n/a	Multicast
E	240 – 255.x.x.x	n/a	Experimental

Private addressing schemes: (RFC 1918)

- Used when most hosts do not require direct Internet access and/or when there are insufficient public addresses available. Special ranges of addresses are used for private addressing which are not routable on the public Internet (see table below). This is the most inexpensive route to go and it provides nearly unlimited network growth.
- A NAT device must be installed to pass traffic from the private network to the public network and vice versa. The NAT device must have one valid public IP address and one private IP address assigned to it. (KB# Q243078)

Range	Prefix
10.0.0.0 – 10.255.255.255	10/8 prefix
172.16. 0.0 – 172.131.255.255	172.16/12 prefix
192.168.0.0 – 192.168.255.255	192.168/16 prefix

Subnet limitations:

- Analyze the network bandwidth to ensure it meets design considerations. If the current subnets are congested with traffic consider increasing the number of subnets.
- In an IP-routed network you must consider the number of hosts in each subnet as well as the number of subnets. When the network is IP-switched you need to design for the number of WAN connections only.
- Always allow for future growth when designing subnetting schemes.

Subnets and Active Directory:

- In Windows 2000, domain controllers in the same subnet are automatically made part of the same site. When you move domain controllers between sites you are actually moving them between subnets. When designing your replication topology, you must consider how your subnets will affect Active Directory replication.

Classless Interdomain Routing (CIDR): (RFC 1519 & Tutorial)

- CIDR was created several years ago to help prevent the Internet from running out of IP addresses. The "Class" system of allocating IP addresses was very wasteful; organizations demonstrating a need for more that 254 host addresses were assigned a Class B address block of 65533 host addresses. Even more wasteful were companies and organizations that were allocated

Class A address blocks, containing over 16 Million host addresses! Only a small percentage of the allocated Class A and B address space has ever been assigned to a host computer on the Internet.

- It was determined that IP addresses could be conserved if the original class system was eliminated. By allocating only the amount of address space that was actually needed, the address space crisis could be avoided for many years. This was first proposed in 1992 as a standard called "Supernetting."
- Under supernetting, the classed subnet masks are extended (or made classless), so that a network address and subnet mask could, for example, specify multiple Class C subnets with one address. For example, if 1000 addresses are required, supernetting four Class C networks together will provide the necessary solution.
- When supernetting an address range, treat all the classes of the addresses being combined into a subnet as Class A. Then use whatever method is preferable to determine the appropriate subnet mask.

Other design considerations:

- Automatic Private IP Addressing (APIPA) is used for TCP/IP address configuration for hosts on a single subnet without a DHCP server. Allocated from 169.254.x.x/16 as specified by IANA. (KB# Q220874 & Q255836)
- Some IP traffic such as streamed multimedia is considered "time-sensitive" and requires that bandwidth is reserved for it. The Quality of Service (QoS - bandwidth management) mechanisms built into Windows 2000 allow administrators to prioritize network traffic. (KB# Q233203 & Q233039)

Performance and availability considerations:

- Authentication, logon and encryption traffic are delay and latency sensitive. It may be necessary to place necessary services on both sides of a link exhibiting latency to prevent a disruption in service.
- Increasing the TCP/IP Receive Windows Size through a registry modification may help alleviate problems with network delay. (KB# Q199947)
- If packet loss is high, check your network for router congestion.
- Combine IP ranges by supernetting. Proper use of supernetting reduces routing issues.
- Use variable length subnetting to divide IP ranges. The subnet mask is adjusted in a hierarchical fashion to accommodate a varying number of hosts in each subnet. Keep the number of routers in the hierarchy to a minimum. Routers that support RIP for IP v2, BGP, and OSPF will support variable length subnetting.
- Route cost metrics should be set equally when there is no cost difference between them.
- Higher cost metrics should be assigned to demand-dial links that are backups to less expensive persistent links.

- Place redundant links and routers between locations where high availability is needed. This improves bandwidth performance as well as availability.

Security considerations:

Packet filtering:

- Data and connection security is provided through TCP/IP packet-level filtering (KB# Q259605). TCP/IP filtering allows you to block inbound traffic to any address that does not appear on your exceptions list, limit traffic to dedicated servers, and filter at the application layer. IP packets can be filtered by their protocol type (except for IPSec, ICMP, IGMP, TCP and UDP) and TCP/UDP port number.

IPSec Overview: (KB# Q231585, Q252735, Q253169, & RFC 2401)

- IPSec itself is a protocol, not a service. It consists of two separate protocols: Authentication Headers (AH) and Encapsulated Security Payload (ESP). AH provides *authentication*, *integrity* and anti-*replay.* It does not encrypt data, but is used when a secure connection is needed but the data itself is not sensitive. ESP provides the aforementioned plus *confidentiality* (data encryption). It is used to protect sensitive or proprietary information, but is associated with greater system overhead for encrypting and decrypting data.
- IPSec can be implemented in a Windows 2000 domain using Active Directory or on a Windows 2000 machine using its Local Security settings. It is not available for Windows 95/98 or Windows NT.
- Supported IPSec authentication methods are Kerberos v5 Public Key Certificate Authorities, Microsoft Certificate Server, and Pre-shared Key. (KB# Q240262)
- The IPSec Policy Agent is a Windows 2000 service that runs within the LSASS.EXE process and shows up in the Services snap-in in MMC. It is loaded at system start-up and retrieves an IPSec policy from either Active Directory or the local registry. After the IPSec Policy has been obtained, it will be applied to ***all*** IP traffic sent or received by that system (default behavior - IPSec policy can be modified to allow "soft associations" KB# Q234580).
- Before two computers can communicate they must negotiate a Security Association (SA). The SA defines the details of how the computers will use IPSec: which keys, key lifetimes, which encryption and authentication protocols will be used for example.

IPSec Encryption Algorithms:

- *3DES, 128-bit* – Provides strongest security but affects performance due to overhead associated with longer key length.
- *DES, 56-bit* – Provides performance improvement over 3DES and can be used when a shorter key length is allowed.

- *DES, 40-bit* – Provides greatest performance but the least security. Use mainly when some security is required, but performance is the primary consideration.

IPSec Authentication Protocols:

- *MD5, 128 bit* – (message digest 5). Less secure than SHA. Requires less CPU overhead and increases performance.
- *SHA, 160 bit* – (secure hash algorithm). Provides stronger security but affects performance. Use for U.S. government contracts that require the FIPS (Federal Information Processing Standard). (KB# Q237849)

Diffie-Hellman Groups:

- *Group One* – Low, 768 bits.
- *Group Two* – Medium, 1024 bits.

Protection	Authentication	Encryption	Diffie-Hellman Group
4 (Highest)	SHA-1 (160 bits)	3DES	1024 bits
3	MD5 (128 bits)	3DES	1024 bits
2	SHA-1 (160 bits)	DES	768 bits
1 (Lowest)	MD5 (128 bits)	DES	768 bits

IPSec Key Exchange:

- *Preshared Keys* – Uses a secret key that has been previously agreed upon by two users. They must be manually configured and are used on non-Windows 2000 standalone systems and systems that are not running Kerberos v5.
- *Public Key Certificates* – Computers not running Kerberos v5 use them for authentication. It is preferable to use preshared keys when large numbers of systems are involved.
- *Kerberos v5* – Default in Windows 2000. Used for authentication with any clients in a trusted domain running this protocol.

NetBIOS over TCP/IP: (KB# Q179442)

- Computers in specialized roles, such as proxy servers or firewall bastion servers, should not have NetBIOS over TCP/IP installed. Windows 2000 allows administrators to disable this feature.

DNS:

Planning a namespace:

- In Active Directory, the namespace is based on DNS. You will need to plan your namespace if you choose to use multiple domains.
- There are two types of namespace: Internal (used by Active Directory) and External (registered with Network Solutions for access from the Internet). When implementing AD, you can choose to use the same or different internal and external namespaces.
- Using the same internal and external namespaces has the following two advantages: uses the same logon names both internally and externally (e.g. jdoe@brainbuzz.com could serve as both the logon and e-mail ID) and uses the same tree name (e.g. brainbuzz.com for example is consistent on both the internal network and public Internet).
- Using the same internal and external namespaces results in a more complex proxy configuration and administrators must be careful not to publish internal resources externally. There is duplication of effort in managing resources (e.g., duplicate zone records). As well, users get a different view of internal and external resources even though the namespace is the same.
- Using separate namespaces makes it easier to distinguish between internal and external resources, as there is no overlap or duplication of effort. This makes things easier to manage and proxy configuration much simpler. Disadvantages of using separate namespaces are that multiple names must be registered with an Internet DNS and logon names are different from e-mail IDs.
- MS recommends that you register any domain name you plan to use with AD even if it will only be for internal use. This is to prevent internal clients from being unable to distinguish between the internal name and a name that has been publicly registered by someone else.

Design and interoperability considerations:

- Number of DNS clients per location? The number of clients determines how many DNS servers must be installed per location.
- How many locations in your organization? Typically at least one DNS server will be installed per location.
- Are there any pre-Windows 2000 DNS servers currently in use? Newer features in Windows 2000 DNS may not work with older Windows and UNIX DNS servers.
- Is Active Directory in use or planned in the future? Active Directory integrated zones are only available in Windows 2000 DNS servers (they reduce management overhead by using AD replication to copy the zone databases to all domain controllers).

- Use only RFC compliant (ANSI) characters with NT4 and older BIND DNS servers; they do not support Unicode. (KB# Q255913, Q250488, Q241973, Q241980, Q151416 & RFC 2181)
- In native mode WINS is not necessary. In mixed-mode DNS requests should be forwarded to WINS for NetBIOS name resolution. BIND servers see WINS and WINS-R record types as invalid. If mixing Windows and BIND, specify that WINS records do not replicate to BIND DNS servers. (KB# Q173161 & Q164176)
- For WINS resolution, use a delegated domain as a placeholder for WINS names. When there is a private and public DNS namespace, the WINS sub domain should reside in the private portion. Organizations using the same private and public namespace should place their WINS sub domain under the root of the organization.

Feature	BIND 4.9.6	BIND 8.1.2	BIND 8.2.1	NT4	W2K
DDNS	No	Supported	Supported	No	Supported
IXFR	No	No	Supported	No	Supported
SRV Records	Supported	Supported	Supported	No	Supported
Unicode	No	No	Supported	No	Supported

Working with zones:

Traditional/standard: (KB# Q227844)

- The primary zone is the only type that has a read/write copy of the database (single master model). Only one primary zone is allowed, but there is no limit to the number of secondary zones (read only). If the server hosting the primary zone fails an administrator must intervene immediately to prevent disruption to network services. Traditional zones are completely compatible with BIND-based (UNIX) DNS servers.

Active Directory Integrated: (KB# Q198473)

- Required for secure DDNS. All domain controllers hold a read/write copy of the zone database file (multi-master replication). Since all DNS servers behave as primaries, the failure of a single server will not affect DNS updates (improves availability). Treated as primary zones by BIND-based DNS servers. Data from AD integrated zones can be replicated to other AD integrated zones or traditional secondary zones.
- Reverse lookup zones can be AD integrated, standard primary or standard secondary. The rules listed above apply to reverse lookup zones as well.

Exposing resources to the Internet:

- DNS queries from within your organization can either be forwarded to that organization's ISP or to the Internet's root DNS servers.
  Incoming queries from the Internet can be resolved on an organization's behalf by their ISP (recommended only if resource names aren't changed often) or by a DNS server maintained by the organization in a screened subnet (use when resource names change frequently).
- Place the primary zone inside the organization's firewall and place the secondary zone (read-only database) inside the screened subnet to prevent unauthorized changes to the DNS database. Do not place an AD integrated zone in the screened subnet as it could jeopardize the security of your AD information.
- The public DNS server should contain only those records necessary to do its job. Placing a complete zone database on the machine could expose private information for servers inside the corporate firewall and will also degrade the machine's performance. (KB# Q193837)

Performance and availability considerations:

- With AD-based DNS servers, simply add more DNS servers as needed to handle traffic. With traditional DNS zones, add secondary zones or delegated domains to increase performance. Delegated domains contain a subset of the domain namespace (e.g., cramsession.brainbuzz.com is a subset of brainbuzz.com). (KB# Q164054)
- Incremental zone transfers (IXFR) place less of a burden on the network than full zone transfers (AXFR) – use them whenever possible. Fast zone transfers compress replication data, but are not supported by older versions of BIND. Schedule replication to take place during off-peak hours when possible, to avoid network congestion.
- A caching DNS server simply resolves requests and caches data from resolved requests until its TTL expires. They can be used to reduce traffic across low-speed WAN links where resource information changes infrequently and insufficient bandwidth for zone replication traffic. (KB# Q167234)
- Network Load Balancing redundant DNS zones spread a traffic load across multiple servers. Use when the amount of time it takes to resolve queries has become unacceptable, when DNS traffic exceeds the capacity of a WAN link at a remote location, or when the connection between the two DNS servers supports the extra replication traffic. (KB# Q240997 & Q248654)
- Use MS Cluster Service to increase availability (local servers only: remote servers cannot be clustered). Clustered servers should share a cluster drive so that both nodes have access to the most recent zone database file. Failed servers can be restored more quickly from a cluster drive, as there is no need to resynchronize. (KB# Q259267)

Security considerations:

- Secured updates are only available with AD integrated zones. Use them to prevent impersonation of servers when using DDNS.

Permissions can be assigned to a group, computer or user account. W2K clients can directly update DNS records but this should only be done if:

- It does not create a security risk
- The client station has a static IP address, and
- It does not create unacceptable management overhead in terms of managing permissions.

- Having a DHCP server perform DNS updates is more secure, reduces the headache of managing permissions, and should be used with **non**-Windows 2000 clients (as they cannot automatically update the DNS).
- Encrypt replication data using VPN and IPSec for additional security. Using AD integrated zones provides further protection, as they will not replicate to other AD zones that are not registered with Active Directory.
- Firewalls should be configured to permit only DNS queries from the Internet and zone replication traffic only from the private network.

DHCP: (RFCs 951, 2131 & 2132)

Design considerations:

- Is the network switched, routed, or a combination of both? Consider the location of broadcast domains and the placement of DHCP Relay Agents to forward lease requests through routers that do not accommodate BOOTP/DHCP forwarding. (KB# Q120932 & RFC 1542)
- When using a single DHCP server, place it on the subnet with the highest population of clients – the other subnets will use relay agents or BOOTP/DHCP forwarding on their routers. Use multiple DHCP servers for a geographically dispersed network, low speed WAN links, or dial-up users.
- To what extent have non-Microsoft hosts been deployed through the organization? They may cause problems by not recognizing MS-specific vendor options like *default router metric base*, which provides a base cost for default gateways to the client. Diskless workstations (BOOTP clients) are becoming increasingly popular but are not properly supported by NT4's DHCP server. BOOTP clients should be placed in the same broadcast domain as a W2K DHCP server that has been updated to support RFC 951-compliant devices. (KB# Q174765)

Performance and availability considerations: (KB# Q199160)

- Increase DHCP lease length when network traffic is a concern. The longer the lease, the lower the traffic.
- When working with a small pool of IP addresses, decrease lease length to make greatest use of your addresses. This has the side effect of increasing network traffic. Windows 2000 clients can be configured to give up their lease at shutdown.

- Using distributed scopes with multiple servers in remote locations increases availability in the event of a server failure. Allocate between 50 – 80 percent of an IP address scope to a server on the local subnet and the remainder to a remote server. When the server on the local segment goes down, the remote server can continue allocating addresses.
- Implement vendor classes (KB# Q240247 & Q266675, RFC 2132) when there is a need to provide similar DHCP options to like groups of clients. User classes are used when specific groups of users have different DHCP configuration options than other groups within the company.
- Windows Clustering increases availability by providing automatic failover if the primary node goes down and failback when the downed server comes back online. Clustering is only available to locally placed machines with a persistent high-speed link. (MS whitepaper)
- Network Load Balancing is not an option with DHCP.

Security considerations:

- Placing a DHCP server outside of your firewall or inside a screened subnet poses a security risk since a valid IP address could be allocated to an unauthorized client (allowing access to network resources). Minimize the security risk by extending lease times (this reduces the chance of an IP address being captured), using the smallest possible address range to meet your needs, and manually reserving/mapping addresses to the MAC addresses of specific clients.

WINS: (RFCs 1001 & 1002)

Design considerations:

- Is the network switched, routed, or a combination of both? Consider the location of broadcast domains and the placement of WINS proxy agent to forward broadcast traffic across routers. (KB# Q121004 & Q164765)
- The advent of DDNS in Windows 2000 has obviated the need for WINS, except in networks that are running pre-W2K domain controllers. WINS should be installed when there is a need to provide NetBIOS name resolution services while reducing the amount of related NetBIOS broadcast traffic.
- Non-WINS clients are supported by installing WINS proxy agent (recommended), static WINS entries (next best), or LMHOSTS entries (most work). To avoid changing hundreds (or thousands) of LMHOSTS files whenever a resource is added or removed, use the #INCLUDE statement to reference a centrally managed LMHOSTS file.

Performance and availability considerations:

- Replication across WAN links should be scheduled in off-peak hours. The frequency of replication can also be controlled.
- The best replication convergence times are provided by a hub and spoke model. Aim for persistent high-speed connections between

replication partners whenever possible. Push- or pull-only relationships should be avoided (except for slow WAN links) when planning for WINS replication.

- For remote servers use push/pull WINS replication. Local servers can be clustered for high availability. (KB# Q226796)

Security considerations:

When a WINS server is placed outside a firewall or inside a screened subnet, use pull only replication from its partner. This replication traffic should be encrypted using VPN tunnels or IPSec. (KB# Q179442)

MS Proxy Server 2.0: (Please see related cramsession, KB# Q164084, FAQ & RFC 1918)

Design and interoperability considerations:

- A special install wizard has been released to upgrade a Proxy 2 installation so that it is compatible with Windows 2000. Please see the release notes.
- If there is a need to reduce private network traffic within an organization then consider implementing Proxy 2 with its Web object caching. Its firewall capabilities can also be used to create screened subnets inside a private network to secure data.
- A proxy server at the edge of the private network isolates it from the public network and secures confidential data. It can also reduce traffic on the outbound connection by caching frequently requested Web objects.
- An organization with insufficient public IP addresses can assign one valid public IP to the proxy server and have it service thousands of clients which are using private, non-routable addresses instead(acting as a proxy on their behalf).
- Internet Explorer 5.0 is all that is required for HTTP and FTP traffic. Install the WSP client for any Windows-based Internet application that uses wsock32.dll or NWLink (32-bit only – see FAQ). For UNIX and Macintosh clients, SOCKS4 compatible applications are supported (SOCKS4 supports TCP but not UDP).

Performance and availability considerations:

- When configuring demand-dial connections be sure to specify the data rate and the persistence of the connection, especially if there is a charge for keeping the connection alive. With digital subscriber line (DSL), it is possible to install DSL and use it with a demand-dial interface for creating a VPN tunnel.
- Active content caching makes the most commonly requested objects available in the cache automatically. It will go out and retrieve objects on its own during low traffic periods if needed. Active caching conserves hard drive space but is more CPU intensive. With passive caching, objects are retrieved when requested by a client and stored in the cache until their TTL expires. Passive

caching uses less CPU time but more hard drive space than Active caching. (KB# Q164085)

- Multiple servers can be configured as a proxy array for fault-tolerance. If an array member goes down, the remaining servers pick up the slack. As the Web content cache is spread amongst the array, the cache is lost only on the machine that fails. All servers in the proxy array must share the same array name and belong to the same AD domain and site.
- Setting up multiple proxy servers for Network Load Balancing provides all three machines with a single IP address used by clients making requests. When one of the proxy servers fails, the others will share the work between them.
- You can use round robin DNS resolution to provide fault-tolerance for proxy servers as well. This provides something of a "poor man's load balancing".
- Proxy servers can be "chained" so that requests are forwarded from one proxy server or proxy array to another.
- It is best to setup a machine with multiple interfaces if the resources of the proxy server permit (centralized administration). If resources are an issue, establish multiple proxy servers (decentralized administration).

Security considerations:

- When your proxy server belongs to an Active Directory domain you can assign access permissions to users and/or groups. In a heterogeneous environment install Services for UNIX, CSNW, and/or Services for Macintosh to provide access for non-Windows clients.
- Proxy can also be installed on a stand-alone computer and access granted (or denied) through its local users and groups. The guest account would only be enabled when it is desirable to have anonymous access to resources.
- When designing hierarchical screened subnets, the broadest security belongs at the top of the hierarchy and becomes stronger as you move lower. (e.g. Management has lax security where as the Research division has very strong security). (KB# Q191146)
- Packet and domain filtering provides the ability to completely restrict traffic by protocol, IP address, domain, user, group, and computer.
- Web publishing allows for placement of a single Web server behind a firewall. This increases security, since the proxy server fetches requested pages on behalf of the client and returns them (acting as a Web server). This hides the identity of the real Web server and protects it from attack.

NAT: (KB# Q234815, Q229965 & RFC 1631)

Design considerations:

- NAT is only appropriate for non-routed network environments where all users have the same access privileges but where private addressing for all computers is required.

- A DHCP server is not required, as NAT will automatically assign IP addresses to machines capable of acting as a DHCP client. NAT should not be installed on a machine that is running DHCP, as they both use the same port (or a machine configured for DDNS). (KB# Q250603)
- The following protocols are not supported by NAT: IPX/SPX (NWLink), SNMP, LDAP, Kerberos v5 (DCs cannot replicated AD information through NAT), RPC, and IPSec (header encryption not supported). (KB# Q261203)
- Choose NAT when you want to exchange traffic between two dissimilar network segments (e.g. Ethernet and ISDN), but the expense and complexity of MS Proxy 2 is not desired. NAT can also be used to create screened subnets but lacks the flexibility of MS Proxy 2.
- A DNS proxy is included in NAT to forward name resolution queries to a DNS server belonging to the organization or one belonging to its ISP.

Performance and availability considerations:

- Dedicate system to running NAT. This enhances both performance and availability as there are no other applications running that consume needed resources or can destabilize the system.
- Use multiple Internet connections whenever possible for redundancy. This prevents a resource from being unavailable in the event of a one connection failing and enhances performance by spreading traffic across multiple connections. Also choose persistent connections whenever possible, as demand-dial connections take time to establish (lower performance) and can reduce availability (busy signals).

Security considerations:

- VPN (PPTP) connections can be used whenever remote users need access to resources on a private network or whenever remote resources need to be secured on a user-level basis. Both outbound and inbound connections are supported. (KB# Q255784)
- Use RRAS IP filters on both the Internet and/or private network interfaces to grant or block access by IP address and/or protocol. (KB# Q256644)
- By default, all computers behind NAT are inaccessible from the Internet. If access to the private network is given to a single IP address, you must define its port mappings within RRAS. This is not necessary when using multiple addresses reserved in an address pool, since all IP ports are open unless specifically filtered in RRAS.

Routing:

Protocols:

Protocol	Description
Appletalk	Routable, proprietary protocol developed by Apple and used for integrating Macintosh systems into a Windows network solution. (RFC 1583)
IGMP	Internet Group Management Protocol. Allows Internet hosts to participate in multicasting (RFC 1112)
OSPF	Open Shortest Path First. Dynamic link state routing protocol – more efficient than RIP. Only sends updated information rather than retransmitting entire routing tables. (RFC 1583)
RIP for IP	Routing Information Protocol for IP. Dynamic distance vector routing protocol – uses considerable overhead as routing table is rebroadcast every 60 seconds. (RFC 1058)
RIP for IPX	Routing Information Protocol for IPX. Similar to RIP for IP. (KB# Q203051)
SAP	Service Advertising Protocol. Proprietary broadcast-based protocol developed by Novell and used by IPX/SPX clients to broadcast their resources. (RFC 1634)

- IGMP is used when existing routers are multicast capable, the IGMP clients are directly connected to the same subnet, and multicast traffic needs to pass to and from the public Internet (NetMeeting and Windows Media Player are two apps that can use multicast).
- RRAS has two modes of support for IGMP: Proxy Mode, which simply forwards multicast traffic to a multicast capable router/server; and Router Mode, which can listen for and update the multicast-forwarding table. Router mode cannot propagate group listening information to other multicast capable routers, however.
- RIP is used when it is desirable to reduce the management overhead caused by maintaining static routes. It should be used if frequent changes to routing information occur, demand-dial interfaces are used, the existing routers use RIP, or there are no more than 14 hops between routers. (KB# Q164363)
- Choose RIP version 2 if your network includes variable length subnet masks, CIDR, multicast routing table updates, or password authentication between routers.

Choose OSPF when dynamic routing is necessary, existing routers are OSPF compliant, there are over 50 subnets, or there are redundant paths between your subnets (link state). OSPF design can be subdivided into three hierarchical levels:

- *OSPF Autonomous System* – a collection of networks that share a common administrative authority. Autonomous Systems (AS) are subdivided into OSPF areas.
- *OSPF Area* – a group of routers that connect to contiguous network segments and are all connected by area border routers (ABR) into a backbone area.
- *OSPF Network* – consists of individual segments that are connected by OSPF routers.

Design considerations:

- For router placement, consider the following questions
  - Do you need to logically segment the network (create subnets) to isolate traffic?
  - Are dissimilar network topologies (ATM, ISDN, Token Ring, and Ethernet) being connected?
  - Does the organization require the creation of screened subnets (packet filtering) to secure confidential data?
  - Are connections persistent (higher availability and data rate) or demand-dial (you will have to set persistence for these), which will invariably add to its operating cost?
- Routers placed at the edge of a network (between the Internet and the private network) can provide firewall type security when packet filtering is enabled.
- Static routing is an option when it is desirable to reduce overhead (generated by dynamic routing protocols such as RIP and OSPF) or to increase security (by preventing transmission of routing tables). It should be avoided when it generates unacceptable management overhead because of the number of resources or the frequency of changes. Routes for demand-dial interfaces must be manually added as neither RIP nor OSPF support them. (KB# Q178993 & Q235492)
- Auto-static routing is a combination of static routes and RIP for IP. It allows an administrator to specify a schedule when a demand-dial connection is established and static route entries are automatically updated. It reduces the management overhead associated with static routes, but it can cause availability problems if auto-static updates are not performed frequently enough. Auto-static routing does not support OSPF. (KB# Q241545)

Performance and availability considerations:

- To obtain the best performance use a dedicated computer as a router. If a router is performing more than one role, its performance will be degraded and possibly its stability as well (lowered availability).
- Persistent connections enhance availability and eliminate connection times associated with demand-dial interfaces. Connections should be redundant, maintaining high availability in the event a connection fails.
- Multiple routers should be installed to provide fault-tolerance in the event of equipment failure.

Security considerations:

- In your network design, RIP for IP or OSPF passwords can be implemented to authenticate routers only if a clear-text password exchange is acceptable and all routers use the same protocols.
- IPSec Machine Certificates provide a greater degree of security. It should only be used when all routers support IPSec (servers running IPSec can only communicate with other servers running IPSec), and there is a Certificate Authority that can issue machine-based certificates. IPSec provides authentication and protection against spoofing when using the Authentication Header (ESP) protocol, but does not encrypt the data itself (choose ESP protocol for that).
- VPN (Virtual Private Network) can be used if there is a need to secure routers (which support VPN). Choose PPTP with NT4 routers and L2TP with Windows 2000 routers. Third party routers may be compatible with PPTP and L2TP: check their specifications when planning to use VPN for security.
- MS Point-to-Point Encryption (MPPE) is used by Point-to-Point Tunneling Protocol (PPTP) to secure confidential data. This method is not as secure as IPSec and only provides user-level authentication. It is also used in lieu of certificate-based authentication.
- IPSec tunnels can be used to protect confidential data. Tunnel mode is used strictly for point-to-point communication whereas transport mode can communicate with more than one computer at a time. When used with L2TP, machine-based authentication is possible using certificates.

Remote Users:

Design considerations:

VPNs:

- A Virtual Private Network (VPN) is an extension of the physical network. Rather than restricting the network to local cabling, VPN uses the Internet as a segment backbone. VPNs are used by organizations that have a need for members to access private network resources via the Public Internet. Windows 2000 has two main encryption protocols that are used with a VPN:

  - MPPE (Microsoft Point-to-Point Encryption) is used with PPTP (Point-to-Point Tunneling Protocol). PPTP was developed by Microsoft and others. It has not been widely adopted by most of the Internet community. MPPE uses 40-bit, 56-bit, and 128-bit (North America only) encryption.
  - IPSec (IP Security Protocol) - an open protocol suite that relies on L2TP (Layer 2 Tunneling Protocol) for encrypting user names, passwords, and data. IPSec is used to negotiate the secure connection utilizing DES (Data Encryption Standard/ 56-bit), and 3DES (Triple DES). IPSec is currently supported by Windows 2000 only.

- There are two VPN connection types: compulsory and voluntary. Compulsory tunnels are initiated by the RAS server, do not require client support for tunneling, and require user-based client authentication (RADIUS is optional). Voluntary connections are initiated by the dial-up user and require support on the client end for the tunneling protocols, but the connections do not need intermediate RAS server support for tunneling.

Dial-up Access:

Used when the security risk from allowing access the private network via a VPN tunnel from the public Internet is unacceptable. RRAS support the MS RAS protocol (NetBIOS only) and PPP, but not SLIP.

PPP supports the following protocols:

- TCP/IP
- IPX/SPX (NWLink)
- NetBEUI
- Appletalk

PPP also support the following WAN technologies:

- PSTN (Public Switched Telephone Network)
- ISDN (Integrated Services Digital Network)
- X.25 and X.25 PAD

PPP supports the following security protocols:

- CHAP (Challenge Handshake Authentication Protocol)– is one step above PAP in that it does not use clear-text passwords.
- EAP (Extensible Authentication Protocol)– allows the client and the server to negotiate the protocol that will be used, in much the same way that networking protocols are determined. Possible choices include one-time passwords, username/password combinations, or access tokens (used to encrypt L2TP).

- MS-CHAP (Microsoft Challenge Handshake Authentication Protocol) requires the client to be using a Microsoft operating system (version 2), or a small handful of other compatible operating systems (version 1).
- PAP (Password Authentication Protocol) uses a plain-text password authentication method and should only be used if the clients you support cannot handle encryption.
- SPAP (Shiva Password Authentication Protocol)– is also one step above PAP. It is there for backward-compatibility and is not favored for new installations

RRAS integrates with the following W2K network services (reduces management overhead):

- *RADIUS* – allows centralized administration of remote access policies, distributed client authentication in a heterogeneous network, and authentication/accounting logging from multiple remote access servers.
- *DHCP* – IP addresses can be allocated to remote access clients.
- *DNS* – remote access clients can register their dynamic IP addresses with the DNS server.
- *Active Directory* – remote access policies can be administered through AD in a W2K native-mode network.

Client/server dial-up designs should specify:

- Which users will be granted remote access,
- Remote access policy restrictions by user or group, and
- How many adapters, phone lines, modems, and ports are needed to support client connections.

Demand-dial routing designs should specify:

- What accounts will be used by the RRAS servers when performing authentication,
- Remote access security policy restrictions,
- Routing capabilities of RRAS servers,
- Demand-dial interfaces used by RRAS servers in each location, and
- How many adapters, phone lines, modems, ports are needed to support connections to remote locations.

Dial-up solutions in non-routed environments:
Consider the following first:

- What is the aggregate throughput required by the remote access clients? Make sure the LAN interface in the remote access server can handle the traffic.

- What is the security model in place for remote access users? W2K native-mode domain policies have greater flexibility over mixed-mode policies.
- How many concurrent dial-up sessions must be supported? If PPP Multilink or BAP are being used it may be necessary to provide more than one connection point per client.
- What will the TCP/IP configuration of the clients be (fixed IP address, allocated by RAS server, or allocated by DHCP server)?

Dial-up solutions in routed environments:

Implementing dial-up solutions should be considered when:

- Access from the public Internet (VPNs) is considered an unacceptable security risk to the private network
- Access from the public Internet (VPNs) outweighs the costs associated with providing dial-up access.
- Security policies require use of additional technologies such as Caller ID or callback.
- An Internet connection does not provide a consistent enough sustained data throughput rate due to router congestion during peak traffic periods.
- Client requirements necessitate additional connections to accommodate bandwidth requirements (Multilink or BAP).

Choose VPN as part of a network design when:

- Access to the private network via the Internet is an acceptable security risk.
- The variability of Internet bandwidth is not a concern.
- The organization's Internet connection supports the projected aggregate bandwidth of the maximum number of concurrent remote access client connections.

Performance and availability considerations:

- Position the RAS server on the subnet with the most client-accessible resources in a switched, non-routed LAN to minimize the amount of unicast traffic flowing across all segments and to minimize cross-subnet traffic in routed networks with multiple routers.
- Position the RAS server in a single segment, non-switched LAN when clients are only allowed access to resources on the RAS server.
- Do not try to combine a VPN server using L2TP tunnels and IPSec with a NAT server. The NAT server will be unable to read the encrypted IP headers.
- Use Internet Connection Manager to connect remote dial-up users to your network. Assign each remote-access client a backup phone number in the event of server failure. Connection Manager can also be used to reduce management overhead when distributing updated access numbers to remote clients.

High availability designs must include more than one VPN server. Multiple VPN servers can have their traffic distributed via round robin DNS entries (this uses less resources than Windows Clustering).

Network Load Balancing can make up to 32 W2K VPN servers appear to the client as a single server. It is more resource consumptive but provides immediate failover.

Security considerations:

Restricting access on a private network:

The following client access restrictions can be placed upon remote users:

- Access is confined to RAS server only (set by server, not by user).
- Static routes are defined only to specific subnets where access is granted (can be set by user or server policy).
- Access is permitted to all resources on the routed network (this can only be set by server, not by user).

Place an RAS server in a screened subnet when:

- Security policies specify that all client access must take place through a firewall or filter (this creates a "screened subnet"),
- The majority of resources accessed by remote clients exist in the screened subnet,
- Clients VPN tunnels to connect to the private network, and/or
- The RAS server contains data that is made available to the public Internet.

Place a VPN server outside the firewall when:

- Confidential data is protected behind the firewall and the only access through the firewall is strictly limited to the VPN server,
- Allowing access to the complete range of VPN IP address through the firewall poses an unacceptable security risk, and/or
- It will not compromise the integrity of the network design's security to expose the VPN server directly to the Internet.

RADIUS: (RFC 2138 & 2139)

Overview:

- Internet Authentication Service (IAS) is Microsoft's implementation of the Remote Authentication Dial-in User Service (RADIUS). RADIUS and IAS together perform centralized connection authentication, authorization, and accounting for dial-up and virtual private network (VPN) remote access and

for router-to-router connections. Used in conjunction with RRAS , they enable single- or multiple-vendor network remote access.

Design considerations:

- Place RADIUS clients as near as possible to remote users creating a local point-of-presence (POP - reduce/eliminate dial-up costs), reducing administrative overhead by delegating administration to local network admins in the same region, and reducing the risk of confidential data being exposed.
- RADIUS servers should be placed as close as possible to the server that provides remote user account authentication. This localizes traffic, keeping it within the same private network and helps prevent unauthorized access to the user account database.
- It is possible to outsource dial-up support for remote-access users to an ISP with RADIUS. Local users access the organization's RADIUS server (which performs authentications) through the RADIUS client installed within the ISP's network.
- Dial-up remote access connections are used when organizational security policy dictates additional security measures such as callback, caller ID, or when private network access through the Internet is prohibited. This method entails maintaining a significant number of phone lines, modems, and other expensive hardware.
- VPN connections can be included in a network design when the organization's Internet pipe has enough bandwidth to support the VPN traffic, security policies allow the private network to be accessed via the Internet, and remote-access policy allows for the outsourcing of modems, phone lines, and multi-port communication Adaptors.
- Remote access client protocols:

    - *Appletalk* – used for Apple-based servers, Apple-based file and print resources, and running applications based on the Appletalk protocol.
    - *IPX/SPX* – used for Netware-based servers, Netware-based file and print resources, and running applications based on the IPX/SPX protocol.
    - *TCP/IP* – used for administering W2K-based servers, accessing Web-based applications and FTP servers, and running applications based on TCP/IP.

Performance considerations:

Capacity planning/hardware scaling for an IAS server:

Type of organization	Authentications/second for typical use	Hardware configuration
Small to medium-sized organizations with less than 1000 users	1	Minimum hardware recommended for Windows 2000 Server
Large organizations with 50,000 users	10	Minimum hardware recommended for Windows 2000 Server
ISPs with 2 million users	50	200 MHz Pentium II or higher.
ISPs with u20 million users	300	4-processor Xeon or higher.

Performance Guidelines for a Single IAS Server :

Hardware	Authentication methods	Maximum authentications/second
Minimum hardware recommended for Windows 2000 Server and a remote Active Directory domain controller	CHAP, MS-CHAP v1, MS-CHAP v2	50
200 MHz Pentium II and a remote Active Directory domain controller	CHAP, MS-CHAP v1, MS-CHAP v2	200
4-proccessor Xeon and a remote Active Directory domain controller	CHAP, MS-CHAP v1, MS-CHAP v2	700

- When selecting the data rate and persistence always attempt to specify a persistent connection that exceeds the required data rate.
- When using MPPE encryption, 40-bit provides less security than 128-bit; however it is less CPU intensive because of the shorter key length. If security is paramount, use 128-bit encryption and allocate the necessary resources to accommodate the reduction in server performance.
- The RADIUS server must have a high-speed, persistent connection to the global catalog server. If CPU performance is not an issue, installing IAS on the global catalog server may increase authentication performance.
- RADIUS authentication/accounting performance can be improved by adding additional RADIUS servers as needed, upgrading existing servers, and reducing the level of accounting detail logged.
- To design for RADIUS client availability install redundant RADIUS clients and give remote users phone number for the primary and backups, install sufficient phone lines and modems to handle the user load, and register your redundant RADIUS clients with the RADIUS servers to guarantee proper authentication/accounting.
- To design for RADIUS client availability for VPN connections use Network Load Balancing to provide immediate failover or round robin DNS to distribute the load across multiple RADIUS servers.

Security considerations: (KB# Q246118)

- Authentication can take place from any domain that is accessible to Windows 2000. This includes Windows NT 4.0 domains, Windows 2000 mixed-mode domains, Windows 2000 native-mode domains, as well as any domains that are accessible through trust relationships (e.g. Kerberos 5 authentication realms). RADIUS only supports a single default domain, but users can specify a different authentication realm (domain) if necessary.
- Both the RADIUS client and server use remote-access policies in conjunction with a user account's dial-up properties to grant authorization. While a user is connected, RRAS matches the connection to settings of the user account and remote-access policy profile. As long as they match the connection stays alive (e.g. profile settings allow one hour maximum connection time. When a user goes over an hour, the policy no longer matches and the user is disconnected).
- MS recommends specifying connections between the RADIUS client and the server that encrypts all data and authenticates using VPN or IPSec. RADIUS secrets (KB# Q168667 & RFC 2139) should be used between mutually authenticating RADIUS servers. The RADIUS secrets should be at least 16 characters long and include a mixture of uppercase and lowercase letters and punctuation.

Authentication protocols:

- *PAP* (Password Authentication Protocol) uses unencrypted (clear text) authentication. Only use when no other authentication protocol is supported.

- *SPAP* (Shiva Password Authentication Protocol) provides encrypted authentication for Shiva LAN Rover clients.
- *CHAP* (Challenge Handshake Authentication Protocol) provides encrypted authentication for multiple operating systems (including Mac and UNIX).
- *MS-CHAP* (Microsoft Challenge Handshake Authentication Protocol) provides encrypted authentication for Windows 95/98/NT4
- *MS-CHAPv2* (Version 2) provides encrypted authentication for Windows 2000.

Managing Network Services:

Manual testing:

Schedule regular audits of network services security and performance. Manually test Network Load Balancing (e.g. for Web servers use a tool like WebCAT), failover response of clustered servers (switch one off under controlled circumstances and see what happens), stop and restart services as needed, etc.

Monitoring:

Throughout your testing and the regular operation of the network services, analyze how service uptime, performance, and interaction with other services are affected. The Performance Console, Performance Logs and Alerts, Snap-in and regular monitoring of server logs (can be automated through scripting) will help greatly.

Keeping an eye on things:

Management processes must be put in place to readily monitor the current status of the network services, analyze data that is collected, and identify trends to verify that the operation of the services falls within the parameters of the network design. A system for responding to changes (MS recommends SMS) should also be implemented to bring network services back into design specifications.

Data collection tools and strategies:

Data should be collected from multiple points within your network services infrastructure. This information is usually funneled to a central management point by one of two methods:

- *In-band data collection* – status data traverses the same network that provides services. The traffic from collecting this data can impact the network if large amounts are collected, and the data lost in the event of a network services failure. Should be used when the network has redundant paths (fault tolerance).
- *Out-of-band data collection* – status data is gathered via separate physical/logical network connections. Failure of network services/components

being monitored does not affect data collection. Use when the network being monitored is not fault-tolerant.

With centralized data collection, data is collected and analyzed at a central location (usually a host running specialized management tools). This method generates increased traffic and can affect network performance. In the event of a network failure, status data may be lost.

A *distributed data collection* strategy entails accumulating data on multiple nodes within the network infrastructure where it is processed before being forwarded on to a management node. This reduces the burden of the management node and allows localized responses to failures. Use when design planning calls for independent operation of locations.

Event notification is provided by specialized software (Performance Alerts, SNMP) that monitors a service and generates an event when a pre-defined threshold has been exceeded. These software monitors not only generate events in the form of event log entries, they can also notify administrators of problems via e-mail/pager and even restart failed services and/or servers if necessary.

Useful tools:

Utility	Function
NBTSTAT	Displays protocol stats and current TCP/IP connections using NetBIOS over TCP/IP.
NETDIAG	Performs a series of tests that help to isolate network connectivity problems. Can also diagnose state of network client. Found in \support\tools folder on W2K CD.
NETSTAT	Displays TCP/IP protocol statistics and current connections.
Network Monitor	Packet sniffer. Monitors all network traffic sent to and from the computer it is running on. SMS version can capture all data.
NSLOOKUP	Used for troubleshooting DNS problems (host name resolution failure).
PATHPING	Combination of PING and TRACERT. Helps to pinpoint where packet loss is occurring.
PING	Used to troubleshoot IP connectivity.
TRACERT	Used to trace the path taken from the host to the destination router.

Event logs:

Event type	Function
Error	Indicates problems (failure of services) that may lead to a loss of functionality.
Information	Entry made upon the successful operation of an application/driver/service.
Warning	Events that may indicate future problems (e.g. low virtual memory).
Success Audit	Indicates that a successful access to an audited resource has taken place.
Failure Audit	Indicates that an unsuccessful attempt to access an audited resource has taken place.

Performance console:

- System Monitor, found in the Performance Console (**perfmon**), can be used to collect *real-time* data and logs. Be aware that running System Monitor on the system being monitored can affect the integrity of the status data.
- Performance Logs and Alerts is used to log events over a *period of time* (creating reports and establishing performance baselines) and for *event notification*. Choose the appropriate counters for the service you are monitoring (DNS, DHCP, AD, etc.) and establish a management process for analyzing the results. This will help you to determine whether your network services are within design specifications.

SNMP: (RFC 1157)

- Support for Simple Network Management Protocol (SNMP) services may play a large part in your network design. Many of the hubs, routers, and switches in your existing network infrastructure may already be managed by SNMP. It can be used to remotely configure devices/services (using NMS), monitor network performance, and detect faults (when alarms are triggered the events are generated).
- SNMP Agents are the software and hardware that support SNMP. They all have a defined Management Information Base (MIB – RFC 1213) – a configuration database from which they read and write data. Status info can be collected interactively from an SNMP manager or as an SNMP trap (an event generated by the SNMP manager).

Windows Management Instrumentation (WMI):

- You can acquire data on the status of services on local and remote systems through WMI. It provides a central integration point for accessing status data

from multiple sources within a computer. Use it when scripted/programmed access is needed for performance counters and events, but direct intervention with the services is not desired. It is started by default on W2K systems but must be manually started on Windows 95/98 systems.

Using scripts and programs for data collection/analysis:

- Many administrators run scripts or batch files to read accumulated performance data (application logs, event logs, and performance logs) and generate event notifications when certain pre-programmed thresholds are exceeded. MS recommends using Windows Scripting Host in combination with popular languages such as VBScript and JScript to monitor and network services.
- Data collected in the form of log files, event logs, and so on can be imported into Microsoft Excel to provide visual representations in the form of spreadsheets, imported into Microsoft Access or SQL Server databases for analysis, or analyzed using a custom program or third-party solution.
- Custom programmed applications can be created to manage network services as well. They can take the form of stand-alone executables, ActiveX dlls, MMC snap-ins, and COM components. These programs can automate data collection and analysis, maintenance, and event notification as needed.

Reactive and proactive response strategies:

- Reactive responses occur *after* an event notification and should only be used in a design if there is fault-tolerance built into the network services (e.g. clustering) and some downtime can be tolerated. Reactive responses are usually triggered by events such as a help-desk call, e-mail notifications, performance-counter related events, and warnings from management and monitoring systems.
- Proactive responses happen *before* the problem really becomes a problem and are based on implementing management processes that to future resource usage limits and failures. Proactive responses are reliant on the collection and analysis of status information on performance, services, network traffic load, data from manual testing, etc. Include proactive responses in your design strategy when downtime must be minimized and prior warning of resource issues/limitations and performance related failures is essential.

Combining Network Services:

Advantages:

Combining services (e.g. DHCP and WINS services on the same server cluster) can reduce the number of computers needed which results in cost savings and reduced management overhead. When services are combined properly performance, availability, and security can also be improved. Services should be combined when:

- The organization's goal is to reduce the number of computers,
- Existing computer hardware resources will support the combined services, and/or
- Combining services enhances performance, availability, and security.

Disadvantages:

The most common obstacle to combining services on a single computer is hardware resources. The trick is to recognize which services use which resources and combine them properly so that all resources on the machine are fully utilized (e.g., combining a CPU intensive service with a RAM intensive services).

Hardware Resources:

Service	RAM	CPU	Network	Disk
DHCP	Low	High	Low	High
DHCP Relay Agent	Low	Med	Med	None
DNS	High	Med	Low	High
IAS	High	Med	Low	None
IPSec	Low	High	Low	None
MS Proxy 2	High	High	High	High
NAT	High	High	High	None
Remote Access Server	High	High	High	None
RRAS Router	High	Med	High	None
VPN	Low	High	Low	None
WINS	Low	Low	Med	High
WINS Proxy	Low	Low	Low	None

- Also, the presence of certain applications running on a system may preclude combining certain services because of resource issues or other conflicts (e.g. NAT and DHCP cannot be combined on the same server).
- With DDNS, the DHCP service performs frequent DNS updates. If the services are on separate machines, network traffic is generated whenever updates are

performed. If there is a large volume of updates, consider combining the services on the same machine to reduce network traffic.

- The layout of physical networks may also prevent the combination of services. Services may be combined when the clients that access them are in the same geographic location as the system providing the services. If the routers and network segments between the clients and the systems running combined services can support the extra traffic load then it is acceptable to have some geographic separation.

Combining with clustering services:

- DHCP and WINS are cluster-aware services and automatically store critical data on cluster-based drives. These services will automatically failover when the primary system in the cluster goes down. Always make sure that cluster-aware services are set up for automatic failover.
- When working with services that are not cluster-aware, make sure that both servers have been configured for automatic failover and that critical data is stored on a *shared* cluster drive.

Security considerations:

- Services that define screened subnets (Proxy Server 2.0 and RRAS) should be isolated. When these computers connect to the public Internet, only those services required to create the screened subnet should be combined.
- Services running inside screened subnets should only be combined when all users accessing the system require the same network resources at the same security level.
- Combining services running inside a private network is usually best. These systems are at low risk as security systems and are in place on other systems dedicated to the task.

This Cramsession was authored by
Sean McCormick, MCSE, MCT, MCP+I, A+, Network+
Chief Technology Officer, Internet-University.net
sean.mccormick@techie.com