Examnotes for Exam 70-226: Designing Highly Available Web Solutions with Microsoft Windows 2000 Server Technologies

Abstract

This exam guide was created to aid you in the focus of your studies for the new Microsoft exam 70-226. You will be expected to know a great deal about the designing of N-tier architecture, Load Balancing, Clustering, IIS, Application Center 2000, Network Infrastructure and security in a DMZ

Audience Profile

Candidates for this exam work in medium to very large Internet or corporate intranet environments that use Microsoft Windows 2000 operating systems. They have a minimum of two years' experience planning and designing highly available Web site infrastructures. They work in multiserver, n-tier application environments that have the following characteristics:

- Concurrent client connections that can exceed 1,000
- Transactional applications
- User databases, such as LDAP server or directory service
- Internet security, such as firewalls, secure protocols, or proxy servers
- High availability services that can include:
    - Network Load Balancing (NLB)
    - Component Load Balancing (CLB)
    - Cluster service
    - Microsoft Application Center 2000

Links

- Exam Home Page and Objectives
- Microsoft Clustering Info
- NT Cluster Info
- Clustering NT and 2000
- Load Balancing Article
- Windows 2000 NLB

Publications

Recommended publications to study for this exam are listed below:

This is the official MS Press training guide but is not yet available

The other MS Press title to help you understand clustering and Load Balancing is the MS Windows 2000 Advanced Server Clustering Services book

Both Titles are good for training and understanding with the help of a lab setup or actual field experience

If you can get your hands on a basic Application server setup manual, or any SAN based network setup books, they will also help you with this exam.

Study Tips

This test is a Microsoft “Design” exam and it revolves its questions around a high level of information based on design and practical application. Make sure you spend considerable amount of time as an admin of clustered and load balanced solutions before you attempt to design them. For exam purposes, you could easily get by with enough study but setting up load balanced and clustered solutions and being able to “design” them for a client will be the true test. Make sure you know the ins and outs of IIS and COM+ components architecture in how they relate to IIS and load balancing. Be familiar with Application Center 2000 as well

As always, do not use this study guide as your sole study source, as it was meant to be a guide to aid and direct your studies in the right direction. Please make sure you thoroughly prepare for this exam properly

Objectives

Skills Being Measured
This certification exam measures your ability to design Web solutions that incorporate Microsoft Windows 2000 Server technologies. Before taking the exam, you should be proficient in the job skills listed below:

Designing Cluster and Server Architectures for Web Solutions
Design NLB solutions to improve availability, scalability, and fault tolerance. Considerations include the number of hosts, number of clusters, placement of servers, multicast versus unicast, failover strategy, priority, affinity, filtering, load weighting, and application types.

Design Cluster service cluster solutions to improve fault tolerance. Considerations include the number of nodes, placement of servers, cluster resource groups, failover and failback strategy, active/active, active/passive, application types, and dependencies.

Design CLB solutions to provide redundancy and load balancing of COM+ components. Considerations include the number of nodes, placement of servers, NLB, and CLB routing.

Design data storage for high availability. Considerations include RAID and storage area networks.

Design a system management and monitoring strategy. Considerations include performance monitoring, event monitoring, services, data analysis, and WMI.

Design a disaster recovery strategy.

Designing a Highly Available Network Infrastructure
Design a TCP/IP network infrastructure. Considerations include subnet addressing, DNS hierarchy and naming, DHCP server environment, and routed and switched environments.

Design a highly available network topology. Considerations include redundant paths, redundant services, and redundant components.

Plan server configurations. Considerations include network adapters, cluster communication, connectivity, and bandwidth.

Analyze and design end-to-end bandwidth requirements throughout an n-tier environment.

Planning Capacity Requirements
Calculate network, server, and cluster capacity. Considerations include memory, CPU, cost, flexibility, manageability, application scalability, and client/server and server/server communications.

Design an upgrade strategy for networks, servers, and clusters. Considerations include scaling up and scaling out.

Calculate storage requirements. Considerations include placement, RAID level, and redundancy.

Design directory services. Considerations include Active Directory, LDAP, availability, authentication, and sizing.

Designing Security Strategies for Web Solutions
Design an authentication strategy. Considerations include certificates, anonymous access, directory services, Kerberos, and public key infrastructure (PKI).

Design an authorization strategy. Considerations include group membership, IP blocking, access control lists, and Web content zones.

Design an encryption strategy. Considerations include IPSec, SSL, certificates, Encrypting File System (EFS), and PPTP.

Design a firewall strategy. Considerations include packet filters, proxy servers, protocol settings, network address translation (NAT), and perimeter networks (also known as DMZs).

Design a security auditing strategy. Considerations include intrusion detection, security, performance, denial of service, logging, and data risk assessments.

Designing Application and Service Infrastructures for Web Solutions
Design a Microsoft Exchange 2000 Server messaging Web integration strategy. Considerations include browser access and Wireless Access Protocol (WAP) gateways.

Design a database Web integration strategy. Considerations include database access and authentication.

Design content and application topology. Considerations include scaling out, load balancing, fault tolerance, deploying and synchronizing Web applications, state management, service placement, and log shipping.

Design an n-tier, component-based topology. Considerations include component placement and CLB.

Design an application management and monitoring strategy. Considerations include detection and notification of application failure.

Exam Notes

Design Fundamentals

For the exam: Designing Highly Available Web Solutions with Microsoft Windows 2000 Server Technologies – you will be responsible for implementing the above objectives into an applicable design for your clients. Remember simple things like Active directory placement, DNS, TCP/IP and Routing / Switching technologies are bound to rear their heads on this exam. (Unavoidable) and they are also in the objectives. (You even see an objective working with Exchange 2000). Make sure you have completed the core exams and have some hands on experience with 2000 / IIS and Application Center before tackling this exam. Good Luck

Notes Compiled to aid in further studies:

The Following Sections are notes that were compiled to aid in your studies for the Design exam. Remember that you are expected to know what each objective is from above, the following terminology and how to apply them to a strong design. This is a list of Items that must be known to help you in your design and exam studies. This is by no means a complete list of all you need to know for this exam:

Design Terminology:

Clustering

- In computers, clustering is the use of multiple computers, typically PCs or UNIX workstations, multiple storage devices, and redundant interconnections, to form what appears to users as a single highly available system
- Clustering can be used for load balancing as well as for high availability
- Advocates of clustering suggest that the approach can help an enterprise achieve 99.999 availability in some cases
- One of the main ideas of clustering is that, to the outside world, the cluster appears to be a single system
- A common use of clustering is to load balance traffic on high-traffic Web sites
- A Web page request is sent to a "manager" server, which then determines which of several identical or very similar Web servers to forward the request to for handling
- Having a Web farm (as such a configuration is sometimes called) allows traffic to be handled more quickly
- Microsoft, Sun Microsystems, and other leading hardware and software companies offer clustering packages that are said to offer scalability as well as availability
- As traffic or availability assurance increases, all or some parts of the cluster can be increased in size or number

Typical Clustered Scenario:

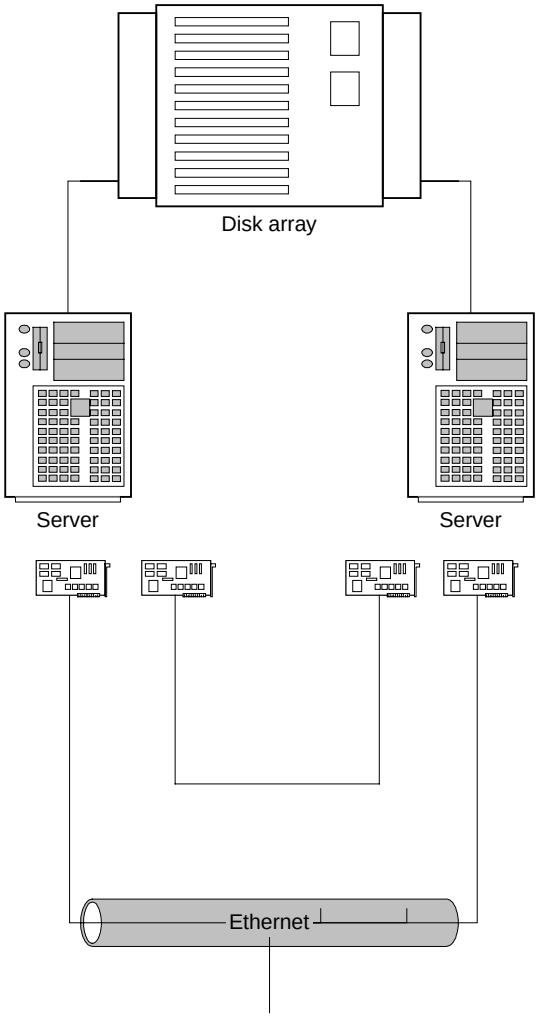

- Remember that a Cluster will be accessed via a Virtual IP and the will have two NIC's to be able to access each other
- The Storage can also be shared as well

Load Balancing

- Load balancing is dividing the amount of work that a computer has to do between two or more computers so that more work gets done in the same amount of time and, in general, all users get served faster
- Load balancing can be implemented with hardware, software, or a combination of both
- Typically, load balancing is the main reason for computer server clustering
- On the Internet, companies whose Web sites get a great deal of traffic usually use load balancing
- For load balancing Web traffic, there are several approaches
- For Web serving, one approach is to route each request in turn to a different server host address in a domain name system (DNS) table, round-robin fashion

- Usually, if two servers are used to balance a workload, a third server is needed to determine which server to assign the work to
- Since load balancing requires multiple servers, it is usually combined with failover and backup services
- In some approaches, the servers are distributed over different geographic locations
- Currently, NT clustering solutions do not support dynamic load balancing (performance clusters), but you can perform manual load balancing
- For example, on one server, you can run Microsoft SQL Server with the accounting department's database, and on the other server, run SQL Server with the order entry department's database
- Each server is fully loaded, with separate databases on one shared SCSI disk array
- Each server is the primary owner for one database, and the secondary owner for the other-both servers are fully utilized but cannot run both databases at the same time

Typical Load Balancing Scenario:

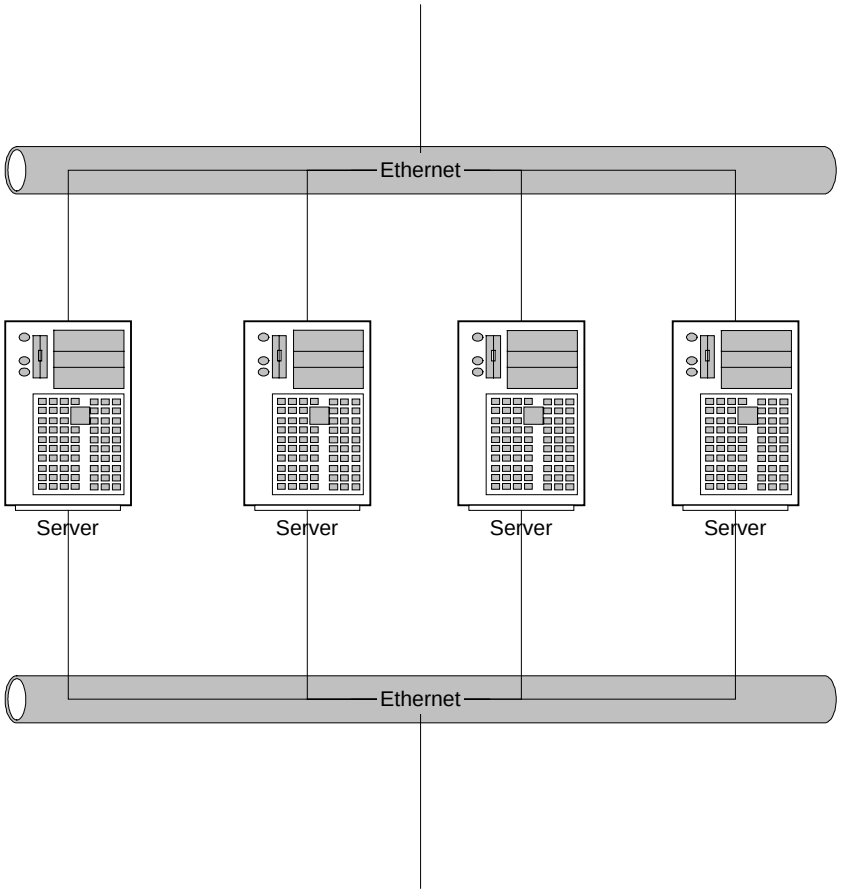

- All servers look like one to the users, and they are all accesses via one IP address (a Virtual IP)

Active/active configuration

- In an active/active cluster configuration, both servers perform meaningful work all the time
- For example, an active/active cluster might have Microsoft SQL Server running on both nodes, and SQL Server on one and Exchange on the other
- Active/active clusters let you fully utilize your server resources and can support load balancing

Active/standby configuration

- In an active/standby cluster configuration, one-server functions as a cold standby (i.e., it performs no meaningful work) to an active server, or one server performs meaningful work but does not run the same application at the same time as the other server
- For example, an active/standby cluster might have Microsoft SQL Server installed on both nodes, but running on only one node at a time; when a failover occurs, the database moves to the other system

Alias

- When you create a cluster, you create an alias that users connect to
- From the cluster's standpoint, the alias refers to the server that owns the service at a particular time
- From the users' standpoint, the alias refers to the service they connect to (e.g., SQL Server), regardless of which server the service resides on
- When a failover occurs, users reconnect (unless the application or client operating system handles reconnection for them) to the alias (not to a server) and continue working
- If users connect directly to the primary server rather than to the alias, they cannot reestablish their connections when the primary server fails

Application recovery kit

- Most clustering solutions have application recovery kits (DLLs, services, etc.) for specific server-side applications, such as Microsoft SQL Server, Microsoft Exchange Server, and Lotus Notes
- The kit enables the clustering software to fail over all resources (files, IP addresses, disk drives) associated with an application from one node to the other

Client agent software

- In the past, some NT clustering solutions (most notably, Digital Clusters for Windows NT 1.0) required agent software to run on the client side
- This software let users access cluster aliases, and even let some client applications understand and handle the service disruption when the primary node failed
- Client agent software provided a way to easily create cluster-aware applications

Cluster objects and groups

- Most clustering solutions employ cluster objects and groups (services)
- An application, a disk drive (for database volumes, files, or directories), an IP address, and so forth are examples of cluster objects
- A cluster group is a collection of related objects that make up a cluster resource such as SQL Server

Clustering APIs

- Most clustering vendors provide APIs so that you can design recovery kits for applications not covered by the basic clustering software
- These APIs let you program to an established (possibly proprietary) standard; several vendors will support the Microsoft Wolfpack API standard

Clustering DLLs

- Clustering solutions have specific DLLs that provide cluster-aware functionality to let programmers create applications that fail over gracefully from one node to the other on system failure

Distributed lock management (DLM)

- Distributed lock management (DLM) enables two servers to access the same physical disk at the same time without corrupting the data
- If a device is updating a particular file or piece of data, the device is locked so that another controller can't seize ownership and overwrite the data
- NT does not currently support DLM, so disks are dedicated to one node or the other

Failback

- Failback switches the service or object from the secondary node back to the primary node after a failover has occurred, and resets ordinary operating conditions
- Failback is usually a manual process, but can be automatic

Failover

- Failover occurs when a node in the cluster or a component of the cluster configuration fails
- Cluster services or resources from the faulty node relocate to the operational node
- You can also trigger failover manually (e.g., by turning off a system to upgrade it), taking one server offline and shifting the work to the secondary server
- System failover capability means that only the entire system can fail over; individual objects (such as applications, disk volumes, etc.) cannot
- The secondary system assumes the identity and workload of the primary system (only on system failure) in addition to its identity and workload
- Application failover capability means that the systems administrator can fail over one application or object at a time, keeping all other services and users intact on the nodes

Fault-tolerant cluster

- A fault-tolerant cluster ties together every action of the two nodes in the cluster, including the instructions running on the CPUs (i.e., the CPUs on each server run in lockstep)
- One server can completely fail without cluster users ever knowing the difference
- The other server takes over instantly because it has been performing the same work as the primary server

Heartbeat

- A heartbeat is the signal that the nodes in a cluster send each other to verify they are alive and functioning
- The nodes transmit the heartbeat over direct (crossover) LAN connections, through a hub or switch, or even via the SCSI bus
- If the heartbeat ceases, one of the nodes has failed and the clustering software instructs the other node to take over
- Employing more than one method to generate a heartbeat eliminates the problem of a minor failure triggering an unwanted failover

Interconnect

- The interconnect provides a communications link between the nodes in a cluster
- Status information, heartbeat, and other intercluster data travels over the interconnect
- This connection can be over your LAN or directly from node to node, using:
    - Ethernet / 100Base-T
    - ServerNet
    - Fibre Channel
    - Serial
    - SCSI
- Fault-tolerant clustering solutions typically use more than one interconnect simultaneously to prevent unwanted failovers

Mirrored-disk cluster

- In a mirrored-disk cluster, the servers contain duplicate drives (no shared disk array)
- Data from the primary server is replicated to the secondary server over a dedicated private network connection (or proprietary high-speed interconnect)
- In the event of a primary server failure, the clustering software shifts object ownership to the secondary server, which uses the duplicate drives to run the applications and data

Node

- A node is one server in a cluster. A node does not include the shared disk array, if one exists

Performance cluster

- Performance clusters support dynamic load balancing
- In a performance cluster, adding nodes to the cluster increases performance by distributing the compute load across multiple systems and provides fault tolerance
- You can improve server throughput in a linear fashion with each computer you add for the same application
- NT clustering solutions do not currently support performance clusters

Primary server and Secondary server

- In a two-node cluster, one server is the primary server and the other is the secondary server
- The primary server usually controls and runs the service
- The secondary server is the failover system
- Each server can be both primary and secondary so that the servers back each other up
- Both servers can perform meaningful work on the network

RAID

- Redundant Array of Inexpensive Disks (RAID) is a strategy that uses technologies such as disk striping, disk mirroring, and disk striping with parity to offer levels of data redundancy and fault tolerance
- All the clustering solutions on NT support fault-tolerant disk subsystems via hardware-based RAID

RAID Types:

- **RAID-0**
    - This technique has striping but no redundancy of data
    - It offers the best performance but no fault-tolerance
- **RAID-1**
    - This type is also known as disk mirroring and consists of at least two drives that duplicate the storage of data
    - There is no striping
    - Read performance is improved since either disk can be read at the same time
    - Write performance is the same as for single disk storage
    - RAID-1 provides the best performance and the best fault-tolerance in a multi-user system
- **RAID-5**
    - This type includes a rotating parity array, thus addressing the write limitation in RAID-4
    - Thus, all read and write operations can be overlapped
    - RAID-5 stores parity information but not redundant data (but parity information can be used to reconstruct data)
    - RAID-5 requires at least three and usually five disks for the array
    - It's best for multi-user systems in which performance is not critical or which do few write operations

Note: There are multiple RAID types but these are the most supported

Shared-disk cluster

- In a shared-disk cluster, the two servers connect to one disk array
- The servers can access the disk array in one of two ways:
    - Simultaneous access
    - Shared SCSI access
- With simultaneous access, each server can read from and write to the same physical disks (not at the same time) on the same disk subsystem bus
- Simultaneous access requires distributed lock management (DLM) to prevent data corruption; NT doesn't support DLM
- With shared SCSI (also known as split-SCSI) access, each server is a termination point on the SCSI bus, each server probably has only one system drive installed, and all failover-enabled applications and data for both servers reside on the disk array
- Only one server at a time owns the disk drives
- When one server fails, the other server takes control of the SCSI bus and all assigned drives

Shared-everything cluster

- Shared-everything clusters follow one of two approaches:
  - Shared memory bus architectures
  - Crossbar systems
- In each approach, all CPUs can access all system resources (memory, disk, network, etc.)
- You will not find the crossbar architecture on the macro-scale in a clustering environment because it was designed for massively parallel systems such as mainframes
- It uses an array configuration of memory blocks and CPU and cache modules (requiring physical locality), resulting in highly scalable systems

Shared-memory cluster

- Think of a shared-memory cluster as an extension of an SMP system's internal design in which all CPUs on one board access one memory system
- A high-speed interconnect (e.g., NUMA) ties together two or more nodes in a performance cluster at the memory bus level to create one large, shared, memory pool that all CPUs can simultaneously access
- Shared-memory clustering provides dynamic scalability, but NT does not support this technology

Shared-nothing cluster

- In shared-nothing clusters, each node in the cluster is a self-contained, fully functional server
- Even if the nodes share a disk array, only one server at a time accesses the disk array, just as if the other system didn't exist
- All the NT clustering solutions the Lab reviewed is a shared-nothing cluster
- The two types of shared-nothing architectures using interconnection networks are network bus architecture and switching fabric architecture
- The network bus is a single connection from system to system, resulting in contention and bandwidth difficulties for heavily loaded systems
- Vinca StandbyServer is an example that uses the network bus for disk mirroring, even though the CPU, memory, and disk systems of the two servers are separate
- The switching fabric architecture uses a high-speed switching technology (such as Tandem ServerNet) to tie the nodes' CPU and memory systems together without the limitations of bus architecture; each node remains a self-contained system

COM+

- COM+ is an extension of Component Object Model (COM), Microsoft's strategic building block approach for developing application programs
- COM+ is both an object-oriented programming architecture and a set of operating system services
- It adds to COM a new set of system services for application components while they are running, such as notifying them of significant events or ensuring they are authorized to run
- COM+ is intended to provide a model that makes it relatively easy to create business applications that work well with the Microsoft Transaction Server (MTS) in a Windows NT or subsequent system

N-Tier Architecture

- An n-tier application program is one that is distributed among three or more separate computers in a distributed network
- The most common form of n-tier (meaning 'some number of tiers') is the 3-tier application, in which user interface programming is in the user's computer, business logic is in a more centralized computer, and needed data is in a computer that manages a database
- N-tier application structure implies the client/server program model
- Where there are more than three distribution levels or tiers involved, the additional tiers in the application are usually associated with the business logic tier
- In addition to the advantages of distributing programming and data throughout a network, n-tier applications have the advantages that any one tier can run on an appropriate processor or operating system platform and can be updated independently of the other tiers
- Communication between the program tiers uses special program interfaces such as those provided by the Common Object Request Broker Architecture (COBRA)
- Here is an Article on N-Tier Development: VB N-Tier

Fibre Channel and Fibre Channel over IP

- Fibre Channel over IP (FC/IP, also known as Fibre Channel tunneling or storage tunneling) is an Internet Protocol (IP)-based storage networking technology developed by the Internet Engineering Task Force (IETF)
- FC/IP mechanisms enable the transmission of Fibre Channel (FC) information by tunneling data between storage area network (SAN) facilities over IP networks; this capacity facilitates data sharing over a geographically distributed enterprise
- One of two main approaches to storage data transmission over IP networks, FC/IP is among the key technologies expected to help bring about rapid development of the storage area network market by increasing the capabilities and performance of storage data transmission

Storage Area Networking

- A storage area network (SAN) is a high-speed special-purpose network (or Subnetwork) that interconnects different kinds of data storage devices with associated data servers on behalf of a larger network of users
- Typically, a storage area network is part of the overall network of computing resources for an enterprise
- A storage area network is usually clustered in close proximity to other computing resources such as IBM S/390 mainframes but may also extend to remote locations for backup and archival storage, using wide area network carrier technologies such as asynchronous transfer mode or Synchronous Optical Networks
- A storage area network can use existing communication technology such as IBM's optical fiber ESCON or it may use the newer Fibre Channel technology
- Some SAN system integrators liken it to the common storage bus (flow of data) in a personal computer that is shared by different kinds of storage devices such as a hard disk or a CD-ROM player
- SANs support disk mirroring, backup and restore, archival and retrieval of archived data, data migration from one storage device to another, and the sharing of data among different servers in a network
- SANs can incorporate subnetworks with network-attached storage (NAS) systems

Network Address Translation

- NAT (Network Address Translation) is the translation of an Internet Protocol address (IP address) used within one network to a different IP address known within another network
- One network is designated the inside network and the other is the outside
- Typically, a company maps its local inside network addresses to one or more global outside IP addresses and un-maps the global IP addresses on incoming packets back into local IP addresses
- This helps ensure security since each outgoing or incoming request must go through a translation process that also offers the opportunity to qualify or authenticate the request or match it to a previous request
- NAT also conserves on the number of global IP addresses that a company needs and it lets the company use a single IP address in its communication with the world
- NAT is included as part of a router and is often part of a corporate firewall
- Network administrators create a NAT table that does the global-to-local and local-to-global IP address mapping
- NAT can also be used in conjunction with policy routing
- NAT can be statically defined or it can be set up to dynamically translate from and to a pool of IP addresses
- Cisco's version of NAT lets an administrator create tables that map:
    - A local IP address to one global IP address statically
    - A local IP address to any of a rotating pool of global IP addresses that a company may have
    - A local IP address plus a particular TCP port to a global IP address or one in a pool of them
    - A global IP address to any of a pool of local IP addresses on a round-robin basis

Note: Other Devices beside Firewalls and Routers can perform NAT, and you can have Windows 2000 Server and ISA do NAT

DMZ

- In computer networks, a DMZ (demilitarized zone) is a computer host or small network inserted as a "neutral zone" between a company's private network and the outside public network
- It prevents outside users from getting direct access to a server that has company data
- A DMZ is an optional and more secure approach to a firewall and effectively acts as a proxy server as well
- In a typical DMZ configuration for a small company, a separate computer (or host in network terms) receives requests from users within the private network for access to Web sites or other companies accessible on the public network
- The DMZ host then initiates sessions for these requests on the public network
- However, the DMZ host is not able to initiate a session back into the private network
- It can only forward packets that have already been requested
- Users of the public network outside the company can access only the DMZ host
- The DMZ may typically also have the company's Web pages so these could be served to the outside world
- However, the DMZ provides access to no other company data

- In the event that an outside user penetrated the DMZ host's security, the Web pages might be corrupted but no other company information would be exposed
- Cisco, the leading maker of routers, is one company that sells products designed for setting up a DMZ and you can use the PIX to set up a DMZ
- Windows 2000 Server ISA server will also allow you to set up a DMZ

IPSEC

- IPsec (Internet Protocol Security) is a developing standard for security at the network or packet-processing layer of network communication
- Earlier security approaches have inserted security at the Application layer of the communications model
- IPsec will be especially useful for implementing virtual private networks and for remote user access through dial-up connection to private networks
- A big advantage of IPsec is that security arrangements can be handled without requiring changes to individual user computers
- Cisco has been a leader in proposing IPsec as a standard (or combination of standards and technologies) and has included support for it in its network routers and VPN and PIX hardware solutions
- Windows 2000 technologies allow for IPSEC as well
- IPsec provides two choices of security service:
    - Authentication Header (AH), which essentially allows authentication of the sender of data
    - Encapsulating Security Payload (ESP), which supports both authentication of the sender and encryption of data as well
- The specific information associated with each of these services is inserted into the packet in a header that follows the IP packet header
- Separate key protocols can be selected, such as the ISAKMP/Oakley protocol

SSL

- The Secure Sockets Layer (SSL) is a commonly used protocol for managing the security of a message transmission on the Internet
- SSL has recently been succeeded by Transport Layer Security (TLS), which is based on SSL
- SSL uses a program layer located between the Internet's Hypertext Transfer Protocol (HTTP) and Transport Control Protocol (TCP) layers
- SSL is included as part of both the Microsoft and Netscape browsers and most Web server products
- Developed by Netscape, SSL also gained the support of Microsoft and other Internet client/server developers as well and became the de facto standard until evolving into Transport Layer Security
- The "sockets" part of the term refers to the sockets method of passing data back and forth between a client and a server program in a network or between program layers in the same computer
- SSL uses the public-and-private key encryption system from RSA, which also includes the use of a digital certificate

Last Tips

This exam includes new types of questions that are very different from other current exam formats. This is an exam that will test all angles of Web Server Design / HA and NLB formats as well as Basic Infrastructure elements. Study hard and prepare well.