

How to Cheat at Designing a Windows Server 2003 Active Directory Infrastructure

Chapter 8 The Physical Design

This chapter is excerpted from *How to Cheat at Designing a Windows Server 2003 Active Directory Infrastructure*, by Melissa Craft, et al [ISBN: 159749058X]. Copyright (c) 2006, Syngress. Reprinted with permission. All rights reserved.

To learn more about Syngress, please visit www.syngress.com.



The Physical Design

Solutions in this chapter:

- Design Internet connectivity for a company.
- Design a network and routing topology for a company.
- Design a TCP/IP addressing scheme through the use of IP subsets.
- Specify the placement of routers.
- Design IP address assignment by using DHCP.
- Design a perimeter network.
- Design the remote access infrastructure.
- Plan capacity
- Ascertain network settings required to access resources.
- Design for availability, redundancy, and survivability.

Introduction

Continuing the theme of physical design, this chapter focuses on the aspect of Active Directory Network Services infrastructure and the considerations relating to physical components. As the previous chapter discussed topics relating to the physical design and placement of the Active Directory Services infrastructure, this chapter focuses more on the practical aspects of networking in support of Active Directory services.

The foundation of networking in any organization is the network topology, which includes routers, switches, network protocols, intranets, extranets, and connectivity to the Internet. If the underlying network is poorly designed or not able to scale to meet business demands, then any infrastructure layered on top will more quickly highlight these limitations and shortcomings.

This chapter begins with a discussion of some of the best practices for designing a network topology, including routing, router placement, Internet connectivity, addressing and subnetting, and firewall considerations. These components all have a bearing on the infrastructure used “on top” of them and therefore require careful consideration when deploying an infrastructure.

Throughout the book we discussed the topology and behavior of a network within the perimeter of the network; however, access to the network from outside that perimeter requires specific attention. This involves the design and implementation of a remote access solution, which when scaled appropriately is robust, secure, flexible, and highly available, yet offers the external users access to all resources they need. We, therefore, discuss the considerations that need to be kept in mind when designing a network topology, such that a remote access solution can easily be integrated into whatever designs are chosen.

Networking and Routing

The first principle of building any reliable and scalable network is assessing and designing a network that can support your current requirements and scale to any future requirements. You need to ensure that you have a supported private internal IP addressing scheme and a properly registered external IP addressing scheme for your network.

Another factor that needs to be considered is how to properly segment the internal and external network. Careful consideration will need to be taken for router placement and security. This topic is discussed in the section *The Network Perimeter*.

Internet Connectivity

Let's first take a look at the external design of your network. In today's networks, Internet connectivity is essential. It provides a means of communication that is both cost effective and expedient. This Internet connectivity is used to support the business in many different ways. An internal network needs to connect to the Internet for such applications as research, e-mail, and e-commerce. It is very important when designing a network that steps are taken to ensure that your organization can connect to the outside and do business, and other organizations and customers can connect to your organization to conduct business.

Configuring & Implementing...

Connecting Your Business to the Internet

This is not a practical exercise where you will be able to walk through the steps in a hands-on fashion at the PC that hosts your newly installed version of Windows Server 2003. Nevertheless, it is an important exercise from the perspective that the knowledge of the steps is important for application in the real world. Certain things must be in place for an organization to be properly connected to the outside world and to be recognized and securely accessible by the public.

1. Assess your organization's business requirements for Internet connectivity. For example, investigate your organization's ability to or desire for hosting its own Web site or electronic mail. Perhaps third-party hosting providers will take care of all hosting, and basic connectivity is all that is required.
2. Procure a link to the Internet from your selected telecommunications carrier and apply for a static IP address (or range of addresses), if required. This will most likely require a visit from an installer who will leave you with some networking equipment, such as a router, firewall, remote access device, or some combination of the three, and an active link to the Internet.
3. Ensure that the IP address for the public network is properly configured on the router's public network interface. Make sure that the router's private interface is configured with an IP address on one of your subnets that the entire organization can access. This will be the gateway address for the clients on your internal network.
4. Choose an available domain name that represents your organization and register it with a domain registrar. When choosing a top-level domain (TLD), you should aim to use .com or .biz if your organization is a company, or .org if your organization is a nonprofit group. TLDs are explained in Chapter 5, "Name Resolution."
5. Arrange for DNS hosting so that your newly registered domain name resolves to the IP address that will be used for your mail and Web site hosts (either on your site or with a hosting provider). Your domain registrar usually takes care of this for you.
6. With a way out for your organization, and perhaps a way in for the public if you went that route, you will need a way to secure the gateway between the private and public networks. Entire shelves of

Continued

books have been written about firewalls and network security—read a few of them. The importance of securing this connection should never be underestimated. If you are using third-party hosting providers, then the concern over what traffic should be allowed disappears since only outgoing traffic is required and no public incoming traffic is permitted.

7. Configure the clients on your private network to use the new gateway to the Internet. This can be done by configuring the IP address of the router's private interface as the default gateway in DHCP, or if you only have a handful of clients, it can be manually entered on every workstation.

Additional activities, such as firewall configuration, setting up a demilitarized zone (DMZ), port forwarding, installing a proxy server, traffic logging and auditing, and Web caching have not been covered in these seven basic steps, and would better covered elsewhere. These steps will get you to the point where your organization can make full use of the Internet, and, if applicable, where the general public can make full use of the services you provide.

Designing the topology requires both a physical and logical approach. The “physical” approach to designing a network topology involves network design that matches the physical, or geographic locations, of your organization. You would use routers to segment the network according to where the clients on your network sit. The alternative to that would be to take a “logical” approach, where your topology would match an organization chart. Using technology embedded in networking hardware, such as virtual LANs (VLANs), you can aggregate individual client connections from many LANs who work together into VLANs. This provides the performance on the same physical LAN to individuals who could be scattered over many. Consequently, you will need to consider best practices in hardware *and* software to ensure that your organization can communicate efficiently and securely with the Internet. Connecting private networks to the Internet creates a whole new set of variables that will need to be considered.

WARNING

Be sure you have an understanding of the types of equipment and services that you can use to connect an organization to the Internet. For a small organization that does not host its own Web services and e-mail, the solution might be to use cable or digital subscriber line (DSL) to implement Internet connectivity, and use the Web hosting and messaging services offered by its Internet service provider (ISP). A larger organization might have to use a telecommunications carrier to provision it with a higher bandwidth solution, so that it can reliably host its own Internet- and Web-based services at a satisfactory performance level for its employees and for customers. This solution will require higher-end routers and firewalls to support and secure the traffic.

Domain Name Registration

Any organization that wants to conduct business over the Internet needs a domain name. To acquire an appropriate domain name, you need to deal with companies that specialize in registering these for you.

The first thing you need to do is choose a domain name. This will not be easy, because most organizations want a “.com” and most of these are taken. You will also need to research the chosen domain name to ensure that there are not any trademark conflicts. Once you have chosen the name, it is time to register it. There are many different registries available to help you with this. Please check www.internic.net/alpha.html for a list of approved registries.


Obtaining a registered domain name is useful for your organization to conduct business on the Internet. It is also useful to have a registered domain name for internal use with Active Directory. Maintaining a registered name internally helps to resolve any conflicts in the future. For example, you create your Active Directory domain with the name `company.com`, but `company.com` is already a registered domain name for a business on the Internet. Users on your internal network will probably not be able to access that site because their Web browsers will think it is internal. A good solution is to select an internal domain name with a suffix that is not a TLD (.com, .org, .net, .edu, .mil, and .gov) or any of the country-specific domains, such as .ca for Canada. An appropriate alternative might be a domain that ends in .dpt (for department) or .internal, or any other name that seems appropriate, but differs from a TLD.

Segmenting the Intranet from the Internet

Most organizations use two different yet similar methods of separating the internal network (intranet) from the Internet. Routers are used as both a stand-alone method and in conjunction with a firewall. Some routers have built-in firewall features to help alleviate having multiple pieces of equipment. Depending on how much work will be required of the router, it might make sense to have a separate firewall to offload the work from the router.

An *intranet* is an internal Web environment that serves an organization's personnel, and is generally not accessible to the public. An *extranet* is means of selectively extending an organization's intranet to individuals and organizations through the Internet who are not physically connected to the organization's network.

Routers will help to route IP traffic in and out of the intranet and Internet. Firewalls are mostly used to filter what IP traffic can pass from the Internet to the intranet. Proxy servers and authentication servers are used for filtering and monitoring what IP traffic flows from the intranet to the Internet. See Figure 8.1 for an example of how this segmentation is achieved with the use of routers with a firewall.

Figure 8.1 Private and Public Segmentation

Proxy servers are very beneficial in separating the intranet from the Internet. Proxy servers provide for a means for your organization's users to access the Internet quickly and securely. Proxy servers can use caching to speed Internet access. For example, if User1 accesses www.microsoft.com, much of the content from that Web page is cached on the proxy server. When User2 accesses www.microsoft.com, User2 will pull some of the content from the local proxy server. This will allow User2 to receive the data at LAN speed. Since proxy servers are an intermediate step in the Internet process, the user's identity is masked behind the proxy server to allow for a more secure experience. Proxy servers also allow you to control what Web sites you don't want your users to access.

To allow internal and external clients to securely access the same resources from the LAN and across the Internet, respectively, place the desired Web infrastructure in a DMZ. A DMZ—also called a perimeter network—is a network segment between an organization's trusted internal network and an untrusted, external network such as the Internet. This segment is protected on both sides by firewalls. The term *firewall* comes from the field of automotive manufacturing. In a car or truck, a firewall is the barrier that separates and protects the passenger compartment from the heat and noxious gases of the engine compartment. In networking, a firewall is a barrier between one network and another, typically an internal network and all manners of destruction that lurk on the Internet.

An organization would be wise to locate its Web, mail, and proxy servers that will be accessed by the public in a DMZ. Microsoft Internet Security and Acceleration (ISA) Server is an extensible enterprise firewall and Web cache server that would be suitable for use as both a firewall and a proxy server. A proxy server, like the name suggests, is a server that acts on behalf of other servers. Specifically, a proxy server operates as a relay between the client and server. For example, a proxy server conceals an organization's internal addressing scheme through Network Address Translation (NAT), and accelerates the retrieval of frequently accessed Web pages that are cached in memory on the proxy server. NAT converts the private IP addresses of an internal addressing scheme to one or more public IP addresses for the Internet by altering the packet headers to the new address and keeps track of each session. When packets return to the organization's network, NAT performs the reverse conversion to the IP address.

**WARNING**

It's important to know the types of devices and services that can segment the Internet from an organization's intranet. It is important to know these devices and how they work. Firewalls, routers, proxy servers, RRAS, and NAT are all examples of these devices and services.


Network Topology Definitions

Physical topology is comprised of geometric components that make up the local area network (LAN) or wide area network (WAN). There are three basic physical topologies: bus, ring, and star. Whatever physical topology is used in the environment, the components of that topology are the same:

- **Subnets** A “division of a network into an interconnected, but independent, segment, or domain, in order to improve performance and security. Because traffic is often the heaviest within a department, and Ethernet is the common network technology, the subnet limits the number of nodes (clients, servers) that have to compete for available bandwidth to a confined geographic area.”
(www.techweb.com/encyclopedia)
- **Routers** Devices that interconnect different physical networks to connect these subnets into one or several contiguous networks, depending on how your organization's network architecture is designed.
- **Switches and hubs** Devices used to aggregate network connections from workstations and to connect different network segments within the same physical network. Switches create dedicated connections between network nodes to take advantage of all available bandwidth. Hubs merely connect all nodes together and the available bandwidth is shared.
- **Perimeter defenses** Consist of devices and software that sit at the edge of your network, most commonly between your internal network and the Internet, and protect the integrity of your network by controlling what traffic is allowed to enter and leave. These devices include firewalls, anti-virus scanners and gateways, and virtual private networking (VPN) devices, among others.

Bus Topology

Bus topology uses an open-ended cable in which all network devices are connected. Both ends of this cable must be terminated. Generally, this topology is not used much anymore and is best suited for small networks because it does not require the use of a switch or hub. Support for this topology has become limited in recent times. See Figure 8.2 for an example of bus topology.

Figure 8.2 Bus Topology

Ring Topology

Ring topology uses a cable that is connected to all network devices, but in a ring formation. In this topology, there is no termination because there are no open ends. In the early days, there was a physical ring (see Figure 8.3). This became a problem because if the ring was broken in any place, your network communication failed. Once IBM introduced Token Ring, a concentrator was used to create a logical ring (see Figure 8.4). In this case, if one of the connections was broken, the ring continued around without losing network communication. There are still network in existence that use this topology; however, generally, new network are not designed in this fashion.


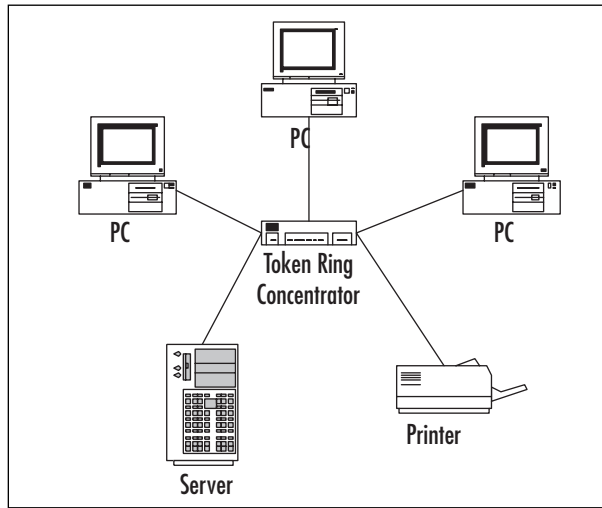
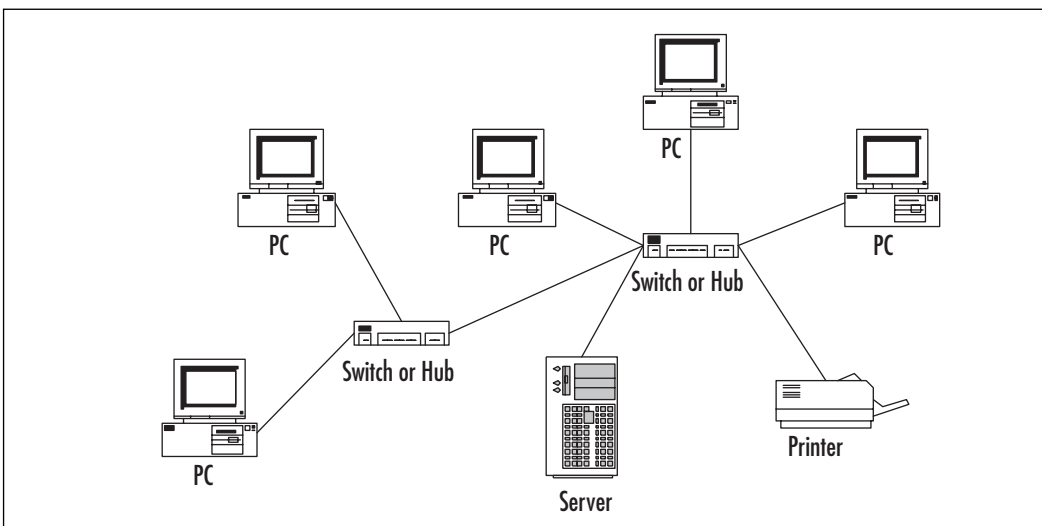
Figure 8.3 Physical Ring Topology

Figure 8.4 Token Ring

Star Topology

Star topology is the most common physical topology used today. In a star topology, each device is connected centrally to a concentrator (switch or hub). The star topology looks similar to how Token Ring is physically set up. However, the star topology is physically and logically the same. Each device is independently connected to the media and does not have to concern itself on how the other devices are connected. See Figure 8.5 for an example of the star topology.

Figure 8.5 Star Topology

Segmenting the Organization into Subnets

A subnet is just a way of taking a complete network and reducing it to manageable and optimized chunks. When designing a network, you want to create a network that will be both fast and secure. Creating subnets will help you to achieve this goal by reducing the size of the network and thus help to control network traffic.

In some instances, you will create subnets to separate groups of devices from one another. For example, you might want to have all of your servers on the same subnet. You might also want to have each floor of your building on a different subnet. Both of these are good practices for creating subnets.

You will also create subnets out of necessity; for example, if your network consists of a WAN. You might have an office in New York City and an office in Philadelphia. These offices will be separated by expensive slow links. This will require the use of separate subnets. In addition, consider the amount of devices you have at each location. For example, if you have 2000 Windows XP Professional desktops at the Philadelphia office, you might have to segment these users into multiple subnets to better manage them.

When segmenting an organization into subnets, you will use routers or devices that can provide routing. You will also require an IP addressing scheme that will make managing the subnets easy. In the *Addressing* section, we discuss the addressing of the subnets further. In the *Router Placement* section, we discuss the placement of routers and routing devices in more detail.

Addressing and DHCP

The Dynamic Host Control Protocol (DHCP) service in Windows Server 2003 is used to provide automatic TCP/IP addressing and management of these addresses. In this section, we discuss what information a designer needs to gather to create a strong DHCP design.

These designs will consist of the three management features supported by DHCP:

- Scopes
- Superscopes
- TCP/IP options

These terms and the concepts behind them are described in depth in Chapter 3, “Developing the Network Services Design.”

We will also look at the DHCP server and the DHCP client. It is important to understand the differences and uses of both. DHCP can distribute IP addresses from a pool (scope) of addresses, or it can always give a device the same IP address.

Why Use DHCP?

As networks increase in size and complexity, the need to ease the management of IP addressing becomes increasingly important. DHCP is a service to assign and manage the IP addresses. DHCP is a client/server process. For DHCP to be successful, there needs to be a DHCP server and a DHCP client. Windows Server 2003 can host the DHCP Server service to facilitate the assigning and managing of IP addresses. For each device that you want to

automatically manage, there needs to be either a DHCP client or Bootstrap Protocol (BOOTP) Client.

DHCP is a message-based service. The client sends out a request for an address and the server responds to the request with an address. The DHCP server keeps track of what address it gives out. This is done to ensure that duplicate IP addresses are not distributed to the clients.

DHCP Design Requirements

One of the very first things you must consider with designing DHCP is how many hosts are in the environment and how many of them will be using DHCP. Keep in mind that not only your users will be using DHCP. Servers, printers, network devices, and other devices can benefit from the use of DHCP—or should we say that administration would benefit if all devices could use DHCP.

In addition to determining how many hosts are in your environment, you also need to determine how many subnets are required in your network design. These factors will help in determining how many scopes or superscopes you will need to include. Collecting this type of information will help you in the design of your Windows Server 2003 network, and with the design of your Active Directory environment.

The first management feature to discuss further is *scopes*. A scope is a range of IP addresses that will be used by a subnet to assign needed IP addresses. These addresses are the first things that are set up and should be the first thing you consider when designing for DHCP.

The second management feature is *superscopes*. Superscopes are a grouping of scopes to support a particular subnet. Superscopes allow for allocating more IP addresses for a subnet without actually extending the scope. An important advantage of superscopes is their ability to use noncontiguous IP address ranges.

Scope options or *TCP/IP options* make up the third management feature of DHCP. These options allow you to create default TCP/IP settings to be delivered to the DHCP client when they receive the IP address assignment. Some of these options include domain name, gateway (router), and DNS servers.

There are four levels for defining these options on a DHCP server.

- **Default global options** These options are applied to all scopes on that DHCP server. By default, there are no global options—the admin must configure them manually.
- **Scope options** These options are applied to a particular defined scope.
- **Class options** If a client has a particular DHCP Class ID, it will receive the specified options. A DHCP Class ID can be set on a client, and then options can be issued from the DHCP server. This becomes useful if, for example, you wanted a particular option for the HR department.
- **Reserved client options** If you create a reservation for an IP address to go to a particular MAC address of a network adapter, then you can also define a particular option to associate with that reservation.

Cheat Sheet...

DHCP Superscopes

Let's discuss DHCP superscopes in more detail. Most importantly, why would you design superscopes into your network? Superscopes are used in environments with more than one DHCP server. When designed properly, they will eliminate DHCPNak (negative acknowledgment) that can occur when a scope for a subnet is split between two DHCP servers.

Here is the scenario: A DHCP client will renegotiate its IP address at 50 percent of the lease time with the server it originally received the IP address. If it cannot renew the IP address at that time because the DHCP server is down or unavailable for some other reason, it will try again at 87.5 percent. This second attempt is called the *Rebuilding* state. This time, it sends a broadcast for the IP address instead of going straight to the DHCP server it originally received it from. If the other DHCP server receives the request, it will check its scope configurations and see that it does not service that IP address in any of its scopes. At that point, it will send the DHCPNak.

To avoid this problem, you can design your DHCP environment to use superscopes. Superscopes allow DHCP servers to have an understanding of the entire DHCP environment—both the scopes they manage, and the scopes that other DHCP servers manage. The superscope will contain member scopes. These member scopes will be scopes on that DHCP server that have addresses it can give out. Additionally, the superscope will contain scopes with exclusions, which are basically scopes that are live on other DHCP servers and are not controlled by that server.

Going back to our earlier example, when the nonmanaging DHCP server receives the request that was broadcast at the 87.5 percent time frame, the DHCP server will check the superscope for the IP address. When it sees that it is excluded, it will simply ignore the request. This will allow the IP address to finish out its lease and request a new IP address instead of trying to renew the current IP address.

Microsoft DNS Integration

An important feature of Windows Server 2003 DHCP services is that it can integrate with Microsoft DNS. Windows 2000 and later clients can register their IP addresses in DNS for name resolution. However, previous versions of these clients do not have this capability. Microsoft DHCP Server can register their IP addresses for them in Microsoft DNS. Furthermore, Dynamic Update is enabled on the DHCP server so that it will register an IP

address that it gives to the DHCP client with a Microsoft DNS server. By default, this option is not enabled and would need to be considered in your design if you are still supporting down-level clients.

TIP

Due to the features that are available in Active Directory Integrated zones, employing Windows Server 2003 DNS is always the best solution when designing DNS for a homogeneous Microsoft infrastructure. . Refer to Chapter 5 for a more detailed description of DNS zones and zone transfer.

DHCP Design Best Practices

As with anything, some general best practices and guidelines should be considered when designing a DHCP environment. These are not rules that must be followed, but are some methods of designing and deploying a strong DHCP environment. DHCP server requirements and placement are discussed in depth in Chapter 6, “Remote Access and Address Management.”

To provide fault tolerance with DHCP servers, use at least two DHCP servers. You should use the 80/20 rule when creating your scopes. That means that on one DHCP server, one scope should contain 80 percent of the address in the scope, and the other should contain 20 percent of the address. You will need to ensure that all TCP/IP options are the same on both servers. In addition, if you are using reservations for any of your clients, be sure to create that reservation on both DHCP servers to ensure that it can obtain its address from either DHCP server.

If you are deploying DHCP in a routed environment and there will not be a DHCP server in each of the physical subnets, use routers that support DHCP relaying or BOOTP relaying. If your routers do not support either, you can design a DHCP relay in that subnet. This relay will forward DHCP requests to a DHCP server for assignments.

When deploying DHCP servers, you should use Windows Server 2003. DHCP services under Windows Server 2003 allow for secure updates, authorization of DHCP servers, and Dynamic DNS (DDNS) services. With secure updates, the computer needs to be authenticated in Active Directory before an IP address is issued. With DHCP authorization, someone cannot bring up another Windows 2000 or 2003 DHCP server without authorizing it in Active Directory—to do so requires Domain Admin rights. With DDNS integration, the DHCP server can be set up to update DNS with the current IP address for a workstation that does not have the capability to perform the update.

Addressing

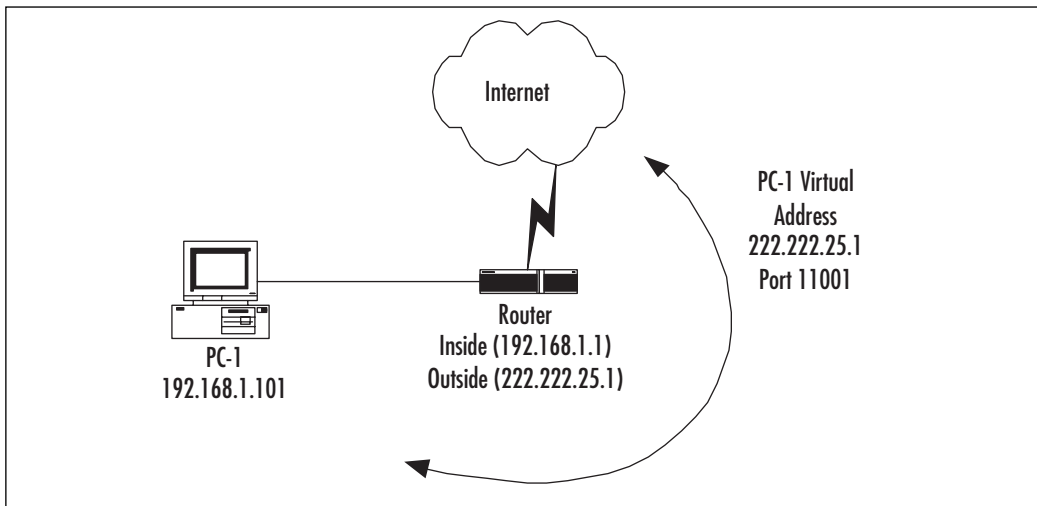
DHCP is useful in managing an IP network, but what addressing do you use? This brings us to the topic of using private and public IP addresses in your network. You need to make sure an appropriate addressing model is used, so as to ensure your network is designed properly.

Should you use a public IP address with your network? Well, this determination comes from examining the needs of your clients. If there is a need for your clients to have public IP addresses, then you will have to ensure that you have enough to satisfy them. In most cases, this is not necessary. It is very difficult to get a large amount of public IP addresses because of the limitation of how many are available. It is better to approach your internal address with the use of private addresses.

RFC 1918 contain ranges of IP addresses that can be used on an internal network. These IP addresses will not be routed on the public Internet, thus eliminating potential conflicts with private IP addresses getting on the public Internet. To use private IP addresses to access resources on the public Internet, you will need to include in your design a means to translate the private IP addresses to public IP addresses. This is done by using Network Address Translation (NAT). Many routers and firewalls support NAT. It is your job to ensure that you choose a router or firewall that supports NAT if you will be using private IP addresses.

Figure 8.6 shows how NAT works. First, PC-1 has an internal IP address of 192.168.1.101. The Internet router has an internal IP address of 192.168.1.1. Both of these IP addresses will not route on the public Internet because they are defined in RFC 1918. The router's public IP address of 222.222.25.1 is a routable IP address. What NAT will do is use the public address for PC-1. It will track on an internal table on the router the translation to ensure that communication between PC-1 and the Internet will occur. It does this by assigning a TCP/IP port number to the 192.168.1.101 address. For example, 192.168.1.101 equals 222.222.25.1 port 11001. Clients outside the router reply to the internal client using 222.222.25.1 port 11001, and that the router then translates this to 192.168.1.101 and forwards packets to the internal client.

Figure 8.6 Network Address Translation (NAT)



Developing & Deploying...

IP Subnetting

In the real world, not every organization is given a public IP address range that will accommodate all devices that require an IP address—there simply are not enough IP addresses to go around. To address this shortage, an organization can do one of two things: use internal addresses from special ranges of IP addresses that are designated for internal, or private use only, or use subnetting of a public IP address to add further scalability within the internal subnets. Let's look at a simple IP subnetting scheme that follows the first approach, which is used in many organizations. For our example, we are talking about Net10.

Net10 comes directly from RFC 1918 (*Address Allocation for Private Internets*). This RFC references current best practice in address an organization's intranet without affecting the public Internet. There are three blocks of addresses from which you can choose:

- Class A (10.0.0.0–10.255.255.255)
- Class B (172.16.0.0–172.31.255.255)
- Class C (192.168.0.0–192.168.255.255)

We will now discuss how you can use the Class A range (Net10) in defining your organization's subnets. First, the default mask for Net10 is an 8-bit mask (255.0.0.0). Using the default mask would give us one subnet that could support 16,581,375 hosts. This would not meet the needs of most, if any, organizations. What we would require is subnetting, for the purpose of creating more subnets that can contain a small amount of hosts.

The first thing you would need to do is find out how many subnets your organization has or will have. Do not forget to plan for scalability in gathering the amount of subnets needed. Another consideration you will need to gather is the amount of hosts in a subnet. Once we have this information, we can plan our subnet addressing accordingly.

To keep this example simple, we will assume that there are no more than 254 users in a subnet. If this is the case, we can use a 24-bit mask (255.255.255.0). A 24-bit mask will allow for 254 users per subnet and allow for 16,581,375 subnets. This will give you plenty of room for scalability. The following list will give you an idea of the ranges you will use for deployment:

- 10.0.0.1–10.0.0.254 (mask 255.255.255.0)
- 10.0.0.2–10.0.0.254 (mask 255.255.255.0)

Continued

- 10.10.0.1–10.10.0.254 (mask 255.255.255.0)
- 10.10.0.2–10.10.0.254 (mask 255.255.255.0)

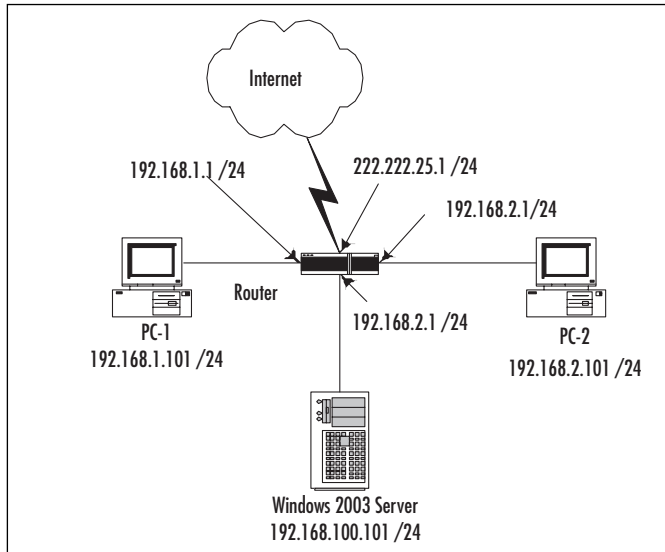
Router Placement

Router placement in your network is important in controlling access and bandwidth. Let's do a quick review of the need for routers in a network. Routers are used as a border for broadcast domains (subnets). Every device on each side of a router can broadcast information to other devices on the same subnet. These broadcasts can be requests for IP addresses (DHCP requests) so that clients can communicate on the network. They can also be requests for IP addresses for another client via the NetBIOS protocol. By default, routers do not let this traffic pass between the subnets; hence the need for devices such as DHCP relays.

Routers also serve as devices that can route data from one host to another. They use different routing protocols to learn the easiest path from one subnet to another, even across multiple routers. Figure 8.8 is an example of how routers segment a network into different subnets. The router in Figure 8.8 is dividing four subnets. The first subnet is a client subnet in which PC-1 is located—192.168.1.x /24. The /24 represents a 24-bit subnet mask, which is 255.255.255.0. The second subnet is a client subnet in which PC-2 is located—192.168.2.x /24. The third subnet is a server or data center subnet in which the Windows 2003 Server is located—192.168.100.x /24. The last subnet is the Internet—222.222.25.x /24. As you can see in Figure 8.7, the router serves to route traffic between the subnets and to contain broadcasts from traveling across to the subnets.

TIP

Knowledge of IP addressing and calculating subnets is essential. These subnets could be within the same building or on different continents, but the principles are the same. Know where to place the routers (for example, with the router's interfaces pointing to different subnets), and how to calculate a subnet with enough available hosts to accommodate the number of nodes in a particular location.

Figure 8.7 Router Placement

It is important when designing a network that you assess the current router placement or design a new router placement that will provide a fast and stable network. The following are considerations when designing router placement.

- **Performance** You want your clients to communicate with your servers with minimal impact. If you place too much complexity in your design, you can impact the performance of the network. Start simple and grow the design from that point.
- **Redundancy** Examine how downtime can affect the business. If the business cannot tolerate any downtime, you will need to design redundancy into your router placement. This means that you will need to provide multiple paths to the business-critical servers and services. Redundancy can also help with performance.
- **Scalability** It is important that the router placement accommodate growth in the network. You want to install routers that can handle more users and subnets than are currently in the environment.
- **Manageability** Any router you place in the environment must be easy to manage and monitor. It is important that you can proactively gauge performance and potential issues before they happen.
- **Security** One of the main advantages to a router is its ability to provide a security boundary. With a router, you can control what type of traffic can flow through it. When designing the placement of the routers, take into consideration what traffic must flow through the router and what traffic should be blocked. Certain subnets might have special requirements in controlling what flows through. For example, a Human Resources department might require that only the IP address assigned to their workstations be allowed through a router to their servers.

- **Cost** Cost should always be a consideration. Because of budgetary constraints, you need to be realistic in placement of your routers. One router with multiple ports might cost less than multiple routers.

The Network Perimeter

Security is one of the most important aspects of a network design. Protecting your network from the outside is difficult. You need to ensure that you design your network with this protection in mind. We are not just talking about the Internet. There are other ways in which threats can get into your network that need to be considered. Your network perimeter will consist of a combination of firewalls, routers, and, perhaps, remote access equipment.

Your router is your first line of defense against the Internet. You can use IP filtering to control what data gets through to your network. You should also consider a firewall in your design. The firewall will serve as an important device in controlling access to your network. A firewall is designed to better handle network perimeter security than a router can, and should always be used in a network design. By design, a firewall inspects incoming and outgoing packets and compares them to a configured set of rules to determine if they should be denied access, dropped, or permitted to pass through to the connected network. A router merely reassigns packets based on the address and port without inspecting the type of packet. Routers can filter by address, but firewalls are much more granular and can filter by MAC address, IP address, TCP and UDP port, and protocol, among others. Therefore, the ability to configure rules for controlling the type of traffic that is permitted to pass through a firewall provides the organization with a much greater degree of control over the integrity of the network and the activity of the clients that connect to it.

Another access into your network can come from dial-up remote access. If there is a need to have dial-in access to your network, you will need to secure this access. You will need to ensure that you only allow dial-in to your network to occur through a device that provides authentication and auditing.

As mentioned earlier, Microsoft ISA Server is a product to consider in this design for securing the network perimeter. This product can provide firewall and dial-up protection for your network. There are also solutions available from other vendors you can consider if ISA does not fit your needs.

TIP

These days, security is job one. That goes for Microsoft solutions as well. Always steer toward solutions that provide a security advantage.

Configuring & Implementing...

Protecting Active Directory

Network perimeter security planning plays a key role when it comes to Active Directory. Active Directory is the database of many objects in your environment. It houses usernames, computer names, and passwords. It is vital that you create perimeter security to protect this investment. Dangerous elements on the Internet and telephone lines pose a real threat.

If you allow these elements to compromise an Active Directory domain controller (DC), you are giving them the keys to your organization. They can access any information that used Active Directory as its authentication mechanism. Another thing to keep in mind is that if they cannot access the Active Directory DC to gain administrative power, they might be able to crash the DC with denial-of-service (DoS) attacks, resulting in downtime and financial consequences.

A best practice to protecting your Active Directory is not to have it exposed to the Internet. Keep it behind a firewall and safe from those dangerous elements. Use a tier approach if you need to authenticate to Active Directory from the Internet. One example of this would be the RADIUS server discussed in the, *Design the Remote Access Infrastructure* section, and, *Ascertain Network Settings Required to Access Resources* section.

Designing Requirements for Remote Access Infrastructures

Let's face it, today, almost every organization needs a way for its clients to connect to the internal, private network when they are not in the office. These connections are necessary for accessing files and checking e-mail, among other things. When you design these solutions, you need to determine what type of access will be required. The following are questions you will need to ask when beginning to design your remote access infrastructure:

- Will there the clients require VPN or will they dial in directly?
- Will partners need to access your network?
- Will the WAN design require the use of the Internet to piece it together using VPNs?
- Will there be a need for demand dialing to connect to remote offices?

Remote access solutions have broad requirements for design and implementation. We will dive in a little deeper to determine the requirements for a remote access design. We will discuss why you need a remote access solution and what information you need to collect to design the solution. We will discuss the perimeter requirements to ensure that the designs can securely accommodate connections from outside the network. We will discuss what is required to allow partners to connect from an extranet perspective.

Finally, we will discuss what is required to establish the remote access solution within the company. This will incorporate how the design interfaces with the current or proposed environment.

Design Requirements

Before designing what hardware and/or software you would need for the remote access solution, you need to determine how your remote access solution will be used. There is certain data you need to collect to ensure you are designing a remote access solution that will fit the needs of the environment and meet the needs of the future.

The first question you need to answer is, how many clients will need to access the network remotely? A good rule of thumb is determining how many clients have notebook computers. These users are mobile for a reason. It will also help to find out if there are any work-at-home users. The organization might have supplied these users with home workstations that will connect back to the environment. You will need this information so you can scale the server(s) to meet the demand.

The next question you need to answer is, are there any partners who will require access to your network environment? If so, what do they need access to? With this information, you can determine how to properly design the VPN and/or dial-up access to allow partners to get to the necessary information. This can also work in reverse. You might have to design the solution to allow a connection to come from your network and access information on the partner's network. In that case, you will need to ensure that the partners can't come back into your network from theirs while the connection is established.

Next, will there be any sites that will require a VPN or dial-in solution to connect them to the network? This is very important because you might have business needs that require these remote sites to deliver or obtain data from the main network. This data might be e-mail or even Active Directory authentication. The solution you might have to establish could be a dial-on-demand using VPN, ISDN, or analog telephone lines.

As you can see, there are many pieces of information you need to collect in determining the design requirements for the remote access solution. Next, we will discuss what hardware and software requirements will be needed to make the design work.

Perimeter Requirements

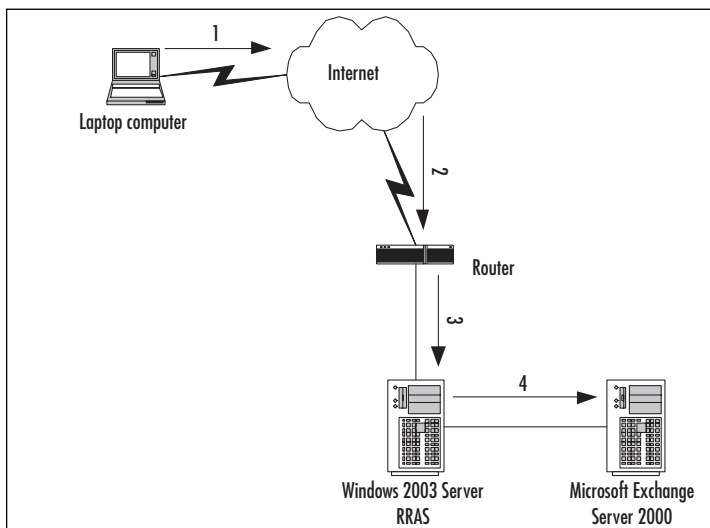
The perimeter is the point at which all remote access will flow into your network environment. Whether access is coming in from your clients or partners for dial-in or VPN, they will access your network through the perimeter. Let's look at what is needed at the perimeter to establish a remote access solution.

Windows Server 2003 is a good solution for implementing on the perimeter to support the remote access solution and provide security for this solution. Windows Server 2003 right out of the box can support dial-in access and VPN access by using Routing and Remote Access Server (RRAS). It can also provide TCP/IP filtering to help protect it from intruders. This is important since it will be located at the perimeter of the network.

Figure 8.8 shows how Windows Server 2003 can be used at the perimeter of the network. Each step in the following list follows the numbering in Figure 8.8:

1. The laptop user connects to the Internet.
2. The user initiates a VPN connection that is routed to your network via the Internet.
3. The perimeter router allows the VPN session through to terminate on the Windows 2003 RRAS server.
4. The user is a virtual node on the same local network as the Microsoft Exchange 2000 Server, and can retrieve his or her e-mail.

Figure 8.8 Windows 2003 Server RRAS




Extranet Requirements

To support an extranet, you and your selected partners need to ensure that you are using a secure remote access solution and that they are using methods for connecting to your network that are compatible with your remote access solution. This could be a Web browser, but you might find that the best solution is typically a site-to-site VPN. Windows Server 2003 can provide this solution with the use of RRAS and dial-on-demand.

Figure 8.9 shows how the site-to-site VPN works. In step 1, when traffic that is destined for your network from your partner's network occurs, using the existing Internet connec-

tion, a VPN connection is initiated from your partner's Windows Server 2003 RRAS. In step 2, the VPN connection is established with your Windows Server 2003 RRAS. This is done with the assistance of dial-on-demand and can occur in either direction.

Figure 8.9 Site-to-Site VPN



Intranet Authentication Requirements

To support a secure remote access solution, you need to establish authentication. To support authentication, you will have requirements on your intranet that will be accessed from the perimeter remote access solutions.

You have two choices for authentication:

- Windows Authentication
- Remote Authentication Dial-In User Service (RADIUS)

Windows Authentication

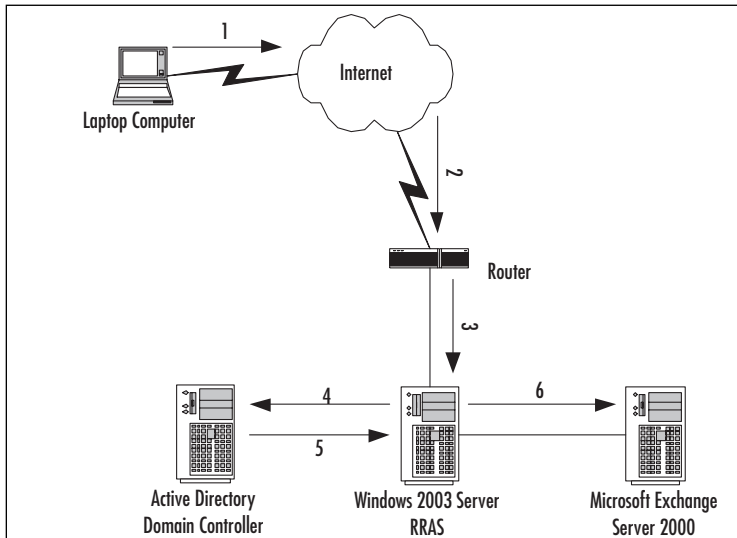
If you only are planning on one RRAS server, then using Windows Authentication will suffice. Your Windows Server 2003 with RRAS will use Active Directory for authentication if it is a member server. If it is a stand-alone server, it will use its internal user database.

Figure 8.10 shows an example of Windows Server 2003 RRAS as a member server and the steps a user would take to make a secure connection from outside the network.

1. The laptop user connects to the Internet.
2. The user initiates a VPN connection that is routed to your network via the Internet.

3. The perimeter router allows the VPN session through to terminate on the Windows Server 2003 RRAS server.
4. The Windows Server 2003 RRAS server used Windows Authentication to authenticate the user with the Active Directory DC.
5. Authentication is approved and the user's VPN is established.
6. The user is virtually a node on the network and can connect to the Microsoft Exchange 2000 Server to retrieve e-mail.

Figure 8.10 Windows Authentication



RADIUS

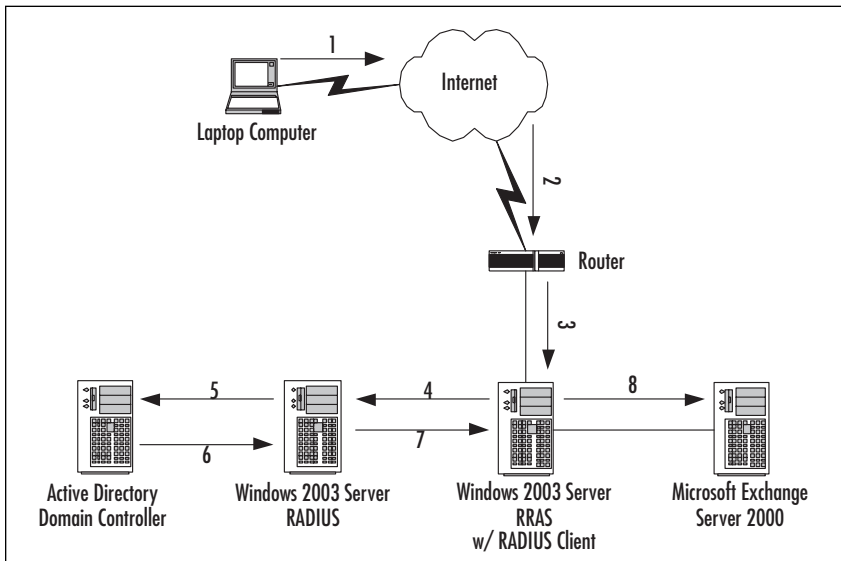
If you are planning to incorporate more than one RRAS server, then Windows Server 2003 should be configured to use RADIUS for authentication purposes. RADIUS is an access control protocol that uses a challenge/response method for authentication. Each Windows Server 2003 RRAS server acts as a RADIUS client. Each of these RADIUS clients authenticates via a top-level RADIUS server, which itself can then authenticate to Active Directory.

Figure 8.11 shows an example of Windows Server 2003 in a RADIUS setup.

1. The laptop user connects to the Internet.
2. The user initiates a VPN connection that is routed to your network via the Internet.
3. The perimeter router allows the VPN session through to terminate on the Windows Server 2003 RRAS server.

4. The Windows Server 2003 RRAS server uses the RADIUS client to authenticate the user with the RADIUS server.
5. The RADIUS server authenticates the user with the Active Directory DC.
6. The authentication approval is sent back to the RADIUS server.
7. Authentication is approved at the RADIUS client and the user's VPN is established.
8. The user is a virtual node on the network and can connect to the Microsoft Exchange 2000 Server to retrieve e-mail.

Figure 8.11 RADIUS Authentication



RADIUS Policies

Another intranet requirement is policies. RRAS policies allow you to control connection times, user and group access, connection security, and others. Using these policies is beneficial for creating a secure RRAS environment. There might be a time of day in which you do not want users to connect because of maintenance. You might also want to force users to use L2TP instead of PPTP for security reasons. Policies allow you to control how you want clients to connect to your organizations network. This is discussed in greater detail in Chapter 6.

Determining Sizing and Availability of Remote Access Infrastructure

Now that we know what we need to design a remote access solution, we have to determine how much of it we require. As with designing any network, you need to know how it is going to be used. You also need to know how many hosts will be using the network. The same goes for remote access.

Let's now look at some of the best practices and principles for sizing the remote access solution. We are going to determine what and where we should place these solutions. We are also going to examine the level of scalability and availability we need to design into the solution.

Sizing Remote Access Components

How many and how powerful should your remote access components be to support your environment? Well, the first thing you need to determine is how many users will need to connect remotely via VPN and/or dial-in. You will also need to determine any other remote access clients, such as site-to-site. This is the starting point for sizing.

Many network designs today tend not to use dial-in because of cost and speed. A better choice is VPN because it does not require the provisioning of additional analog or ISDN lines within the company. With the advent of cable modems and DSL, many companies and users prefer VPN for accessing the organization's network.

Sizing for remote access is a difficult task. A general rule-of-thumb is the 8:1 rule; that is, for every eight users, you should provide one port (dial-in or VPN). Given that rule, you must after implementation monitor these ports to determine that they require being increased. Windows Server 2003 can support 1000 concurrent VPN connections, which equates to 8000 users per server. Keep in mind that this would require robust server hardware in a configuration that is carefully managed and monitored.

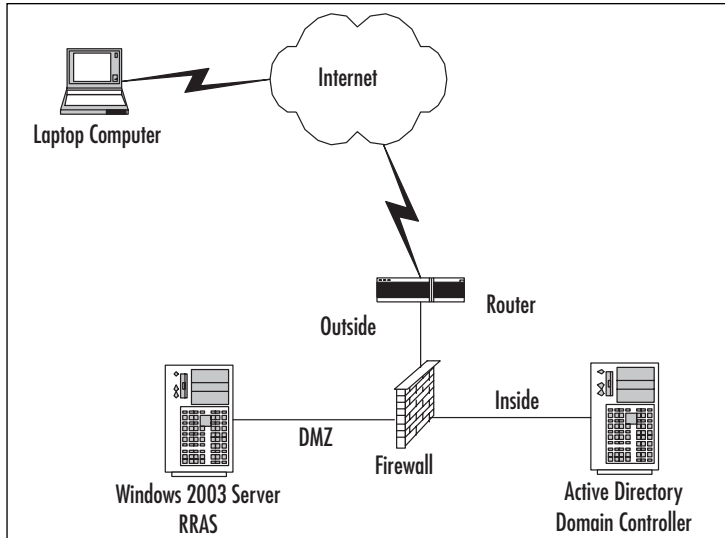
Placing Remote Access Components

We covered the requirements for placing remote access components in earlier sections. Now let's look at where we should place these devices and why. It is important that we place these devices where they can function efficiently and securely. Functionality and security is always a constant trade-off. For example, a network can be locked down to the point where it is barely usable. Alternatively, security measures can be ignored to give clients the greatest degree of latitude in using the network; however, the network could be easily crippled by an attack, rendering it nonoperational. The art of designing any system that has a security aspect associated with it is to get the right balance between security and operation.

If we are dealing with a Windows 2003 Server that is only providing dial-in access to the network, then it makes sense to place this server inside the network perimeter. Since users are only dialing into the server, we do not have to worry about many of the attacks that could be waged from the Internet. For someone to penetrate the network via this type of server they would need to first be authenticated.

If we are deploying a Windows Server 2003 server that is providing VPN access to the network, it should be placed in a DMZ behind a firewall. This will help to protect the server from most attacks, and the DMZ will isolate the inside network from that server in the event it is compromised. See Figure 8.12 for an example.

Figure 8.12 Securing a VPN Server



To reduce cost for connecting remote offices, you can design dial-on-demand routing into your infrastructure. Windows 2003 RRAS server can support dial-on-demand routing. To use dial-on-demand, you need to install RRAS and set it to detect requests beyond the local subnets. This is done by using the RRAS server as a gateway. A gateway is merely a point of access from one network to another. In this context, the RRAS server will initiate and receive outgoing and incoming dial-up connections to another similarly configured server on the perimeter of another network. When a connection needs to be established, it will connect to the destination according to the routing table. Encryption can be used with dial-on-demand. You have a choice for the connection: “no encryption,” “optional encryption,” or “require encryption.” The encryption will only work if it is connecting to another Windows 2000 or Windows Server 2003 RRAS server.

Providing Scalability, Availability, and Failover

It is important that you provide a remote access solution that can scale for the future. A good start is the use of Windows Server 2003. Since each server is capable of providing up to 1000 concurrent VPN connections, you have a solution that is scalable. What will become important is that you provide the scalability in the hardware to ensure that the server can maintain more connections than are required. Monitoring the server’s system resources is the key to maintaining this availability.

When installing RRAS on a server, you will be given the choice of creating a pool of IP address to give to clients or to use DHCP for IP addressing. A safe bet is to use DHCP for IP addressing. Using DHCP will allow you to better manage your organization's IP addressing. The RRAS server will reserve 10 IP addresses from the DHCP server when the service starts. Once these are used up, it will reserve another 10 IP addresses. If your RRAS server is located where it does not have access to the DHCP server, then you might want to provide a pool of addresses for the RRAS server to use for clients.

Availability and failover go hand and hand. To ensure availability, you need to provide the means for failover. The easiest way to do so is to provide multiple remote access servers. You can then either provide users with multiple remote access entries or with a dial-in solution and a VPN solution. Giving them both choices will help to lessen the impact of a failure.

Another consideration for remote access availability and failover is providing dial-on-demand for backing up routers. For example, you can use RRAS to provide a backup if your Internet connection goes down. With dial-on-demand, you can use asynchronous telephone lines, cable, DSL, X.25, or ISDN to provide the dialing access. These connections will need to be set up, and in the case of cable or DSL, you might need to establish a VPN if the connection will be to another site via a site-to-site VPN.

Summary

This chapter covered one of the most important parts of designing a network and Active Directory. Without a good physical assessment and design, you cannot be successful in designing the logical network and Active Directory. These things need to be in place before moving forward.

One of the first sections we discussed was *Networking and Routing*. This section encompassed the physical aspect of how we are going to communicate within our network and beyond. This included Internet connectivity, where we talked about the importance of connecting to the Internet; for example, e-mail and e-commerce. We also discussed how we segment the Internet and the intranet. We covered the importance of using routers and firewalls. This included why it is important to protect your network and, most importantly, Active Directory. We covered the need to register a legal domain name for your organization and how to get started. We then took a look at the different types of topologies that you might find when assessing an organization's existing network. This included the bus and ring topologies, which are not in much use today. We looked at the star topology, which is the most popular and the one you will most likely run into or design.

Next, we jumped into segmenting this physical network into subnets. Subnets let you take a large network and break it into more manageable pieces. We covered how to properly address these subnets using TCP/IP subnetting. We looked at proper router placement and why this is necessary. We also covered how DHCP can be used to help with the assignment and management of TCP/IP addresses.

Digging deeper into DHCP we looked at how Microsoft's DHCP can be integrated with DNS. We looked at some best practices to support redundancy and availability. This included using multiple DHCP servers and superscopes. It also included the use of DHCP relays where the protocol could not pass a router to have communication between the server and the client.

Finally, we covered remote access and the importance of it. We looked at the design of it, which included assessing how much is needed to accommodate your organization. We covered the requirements for designing and implementing such a solution. This included RRAS solutions on the perimeter, and protecting and segmenting them from the intranet and Internet. We looked at what is required for the clients and partners who might connect to your organization. We also looked at solutions like RADIUS to provide support for larger deployments of remote access.

All of the preceding pieces go into the physical design of your organizations network. It is important that you provide a strong assessment of any current environment so you can provide a strong foundation for a new network environment.

Solutions Fast Track

Networking and Routing

- ☑ An organization needs to connect to the Internet to support research, e-mail, and e-commerce.
- ☑ You need to have a valid registered domain name if your organization wants to perform business on the Internet.
- ☑ You should have a valid registered domain name to be used for Active Directory.
- ☑ Routers route traffic between the intranet and the Internet. They also route traffic between segments or subnets in your intranet.
- ☑ Firewalls should be used to help protect your intranet from elements on the Internet.
- ☑ The main components of a physical topology are subnets, routers, switches and/or hubs, and perimeter defenses (firewalls).
- ☑ Networks should be segmented into subnets to reduce the management size of the network.
- ☑ The three management of DHCP are scopes, superscopes, and TCP/IP options.
- ☑ A DHCP scope is a range of sequential IP addresses to be assigned from a DHCP server to a DHCP client.
- ☑ DHCP superscopes provide a mechanism for managing more than one DHCP server in an organization.
- ☑ DHCP TCP/IP options are additional information assigned to the DHCP client. These are set at the server.
- ☑ Using DHCP reduces the size and management of assigning IP addresses to hosts.
- ☑ To use DHCP, you must have a DHCP server and DHCP clients.
- ☑ DHCP relay can be used if there is not a local DHCP server and the router(s) will not allow DHCP requests to cross them.
- ☑ Microsoft DHCP can integrate with Microsoft DNS to provide proper name resolution in a dynamic environment.
- ☑ Private IP addresses should be used in an organization as referenced in RFC 1918.

- ☑ NAT can be used to provide private-to-public address translation so that intranet users can communicate on the Internet.

Designing Requirements for a Remote Access Infrastructure

- ☑ Remote access is driven by the demand of the organization and partners.
- ☑ Microsoft RRAS can provide VPN and dial-in solutions for an organization.
- ☑ The RRAS server should be placed on the perimeter network to allow access and to block bad elements from the Internet.
- ☑ RADIUS can be used as a solution for integrating multiple RRAS servers with one authentication point.

Determining Sizing and Availability of Remote Access Infrastructure

- ☑ Windows Server 2003 can support 1000 concurrent VPN connections.
- ☑ Windows Server 2003 RRAS server can be configured to support dial-on-demand routing. To deploy dial-on-demand, you need to install RRAS on a perimeter server and configure it to listen for requests that originate beyond local subnets.
- ☑ Deploying redundant hardware and software in support of RRAS and VPN will dramatically increase the reliability and availability of the secure remote access solution.

Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to www.syngress.com/solutions and click on the “Ask the Author” form.

Q: Why do I need to register a domain name to perform business on the Internet?

A: To properly account for all businesses on the Internet, there needs to be a legal registration of the domain name. To enable a customer to find your organization, business legal registration is required.

Q: Why are firewalls important to an organization using the Internet?

A: Firewalls provide a means of filtering incoming traffic to your organization’s intranet. Some of the incoming traffic might be of a malicious intent. Routers can provide some of this filtering also, but is it best to offload this work to the firewall.

Q: How and why would a partner to my organization connect using remote access?

A: If a partner needs to access an application or transfer data that business has deemed necessary, then a secure connection might be necessary. This can be done simply by allowing the partner to VPN from one client or as complex as creating a site-to-site VPN to virtually connect the organizations. It will all depend on the necessity.

Q: How do I know how many VPN connections to create to support my organization?

A: A good rule-of-thumb is the 8:1 rule. Create one connection for every eight users who will require VPN access.

Q: I have three RRAS servers in my design; what is the best way to implement this solution?

A: You can include RADIUS in your design. This will create a tiered support for the RRAS solution by allowing all of the RRAS servers to use RADIUS for authentication. RADIUS will handle communication back to Active Directory. This will allow you to centralize the solution.