Internet Security & Acceleration Server

Deploying the Secure Firewall, Proxy, and Web Cache at Microsoft

Introduction

Information technology (IT) and business are becoming synonymous. Modern information technology is essential for automating a wide array of business processes for purchasing, manufacturing, shipping, selling, and marketing new products and services. More often than not, IT streamlines business processes that support an organization's ability to become more responsive to change. In addition, IT is enabling entirely new ways of doing business.

For example, many businesses today are leveraging the global pervasiveness of the Internet to help streamline business-to-business, business-to-consumer, and all line-of-business processes. Using the Internet, many companies are reinventing business as we know it. Internetworked organizations are creating cost-effective, efficient, automated Web-based applications for such central line-of-business activities as invoicing and procurement. The Internet has allowed many organizations to partner more freely with one another while offering more comprehensive services to customers.

As businesses continue to leverage the Internet, the technologies each uses to keep computing and information assets secure are becoming more refined.

Internet-enabled businesses are facing a new set of tough challenges in today's marketplace. Customers expect computing environments to be secure, fast and easy to interact with. Businesses expect that the deployed computing environment will be able to grow to accommodate new demands in the market and their IT professionals demand that such an environment will become simpler to manage and troubleshoot.

Now, to help meet such needs, there is Microsoft® Internet Security and Acceleration Server 2000.

Overview of ISA Server

Microsoft Internet Security and Acceleration Server 2000 (also known as ISA Server) is part of the Microsoft .NET Enterprise server family, which comprises a comprehensive set of server applications for quickly building, deploying, and managing scalable and integrated Web-based solutions and services. Designed with mission-critical performance and integration in mind, the .NET Enterprise servers are built from the ground up for interoperability using open Web standards such as XML. The .NET Enterprise servers, along with the Microsoft Windows® 2000 operating system, supply the foundation for the .NET platform, which enables the third-generation Internet: where software is delivered as a service; is accessible by any device, at any time and any place; and is fully programmable and customizable. The .NET platform is explicitly designed to enable the rapid development, integration, and orchestration of any group of Web services and applications into a single comprehensive solution.

ISA Server is an extensible enterprise firewall and Web-cache server that integrates with Windows 2000 for policy-based security, acceleration, and management of Internetworking. ISA Server provides two tightly integrated modes: a multilayer firewall and a high-performance Web cache server. The firewall provides filtering at the packet, circuit, and application layers; stateful inspection to examine data crossing the firewall; control of access policy; and routing of traffic. The cache improves network performance and the user experience by storing frequently requested Web content. The firewall and cache can be deployed on dedicated servers separately, or integrated on the same box. Sophisticated management tools simplify policy definition, traffic routing, server publishing, and monitoring. ISA Server builds on Windows 2000 security, directory, virtual private networking (VPN), and bandwidth control. Whether deployed as a set of separate firewall and cache servers or in integrated mode, ISA Server can enhance network security, enforce consistent Internet usage policy, accelerate Internet access, and maximize employee productivity for organizations of all sizes.

ISA Server 2000 Enterprise Edition (the focus of this paper) is Microsoft's scalable enterprise firewall and Web-caching server. ISA Server Enterprise Edition was designed to meet the performance, management, and scalability needs of high-volume Internet traffic environments with centralized server management, multiple levels of access policy, and fault tolerance. ISA Server Enterprise Edition provides fast, secure, and scalable Internet connectivity for mission-critical environments.

The Situation within Microsoft

At Microsoft, the Information Technology Group (ITG) is responsible for running the company's internal networks, telecommunication systems, corporate servers, and all line-of-business applications. This group also is expected to deploy new releases of Microsoft products on those systems while those products are in the beta stage. This practice allows each product-development team to receive real-world feedback on its product before releasing it to manufacturing. Ultimately, ITG and the product development-team must jointly sign off on the release of each new product before it is sent to manufacturing.

Employees at Microsoft refer affectionately to the process of deploying each new beta release internally as "eating your own dog food." The phrase captures the challenges of keeping an internal computing information environment running while introducing a product into that environment that is by definition not yet "done." While the process is often challenging, it also results in a customer-ready product and improved morale among employees who contribute to the development or deployment of the new product.

ISA Server is no exception. Before sending it to manufacturing, ITG began deploying it at Microsoft early in the beta stages. Deploying the product this early was key to finding and fixing implementation defects quickly through real-world enterprise deployment feedback.

Internet access is vital to the day-to-day Web lifestyle of Microsoft employees. On an average day, over 40 thousand client computers located at corporate headquarters access over 40 million Internet-based URLs, with an average processing time of just 1.4 seconds per request. Without Internet access from within corporate walls a good part of Microsoft's business would be paralyzed.

Through the internal testing efforts already mentioned and others shared in this paper, ITG has learned many lessons on how to properly install, configure, and deploy ISA Server. The group also has learned how ISA Server can be used to address various business needs and about the benefits provided by the product's exceptional capabilities.

This document captures many of those lessons learned. Although not intended to serve as a general guide or plan for deploying ISA Server, the document illustrates the approach taken by ITG to deploy ISA Server at Microsoft. By capturing and telling the story of how Microsoft deployed the beta release of ISA Server, its authors hope that customers can learn from the internal experience.

Deployment Planning

As with all beta software deployments at Microsoft, ITG began its work with extensive planning and careful consideration of business requirements and product capabilities. Part of those plans were deployment goals and project scope, since both would be key to ensuring that the deployment of ISA Server would satisfy Microsoft's business requirements.

Business Requirements

All computing information environments are different, and therefore all organizations must develop their own strategies, goals, and plans for deploying ISA Server. The following are some of the most critical business considerations at Microsoft, which ITG took into consideration when formulating its strategy, goals, and plans to deploy ISA Server internally:

Customer needs must be met. Microsoft is committed to developing solutions that satisfy customer needs. One such need is a reliable and scalable solution that will enable businesses to communicate with customers and partners using the Internet. To stand behind this commitment, Microsoft developed ISA Server and then kicked off an internal initiative to ensure that the product was enterprise-ready, secure and scalable. ITG and the product-development team created a tight feedback loop to communicate at every step in planning and deploying the beta release of ISA Server at Microsoft to assure that any problems identified were resolved before

release to manufacturing.

Intellectual properties must be secure. Microsoft's intellectual properties are its greatest asset, and ITG is expected to keep that asset secure. For this reason, ITG is extremely careful to avoid compromising the company's security. Before the beta release of ISA Server was deployed in Microsoft's production environment, a team of security analysts reviewed the planning documents and then deployed a small infrastructure based on those plans to determine if the environment could be compromised with techniques commonly used by hackers. They found that it could not. They also found that it was sufficiently secure to deploy at the edge of the internal network, where it would communicate directly with servers on the Internet.

Employees must have rapid access to information. Information is of value only insofar as it can be used to support the day-to-day decision making of employees and executives. Information must be accessible to employees as quickly as they can process it to ensure that business is carried out at "the speed of thought." Rapid access to information over the Internet and information shared on the Internet are crucial business requirements. The Internet has dramatically changed the way Microsoft employees do their everyday jobs. For example, information within the company is provided almost exclusively in electronic, HTML-based form. Most line-of-business applications used at Microsoft now leverage Internet Information Server (the Web server built into Windows 2000 Server), SQL Server™ 2000, and Internet Explorer 5.5. The widespread use of HTML-based content at Microsoft has made ISA Server an ideal solution for securing information and accelerating access to that information.

Distributed environments must be managed using consistent policy. Although most Microsoft employees work at or near corporate headquarters, others are distributed around the world. Employees in all areas of the company need secure and rapid access to the Web and shared information via the Internet regardless of where they work. Managing a geographically distributed environment must be quick and easy, and it is especially important that ITG be able to apply policy consistently to assure the internal environment is secure.

Internet access points are available at many locations throughout Microsoft, allowing a geographically dispersed workforce to take advantage of them. As of this writing there are twenty-two such access points, all of which must be securely monitored and maintained while allowing employees secure and fast Internet access.

The environment must be based on open standards. The day-to-day management of Microsoft's internal computing information environment is simplified thanks to the continual support of many third parties. ITG relies on the dedication and day-to-day support of many solution providers to reduce support costs, improve security, and make the internal environment easier to manage. As a best practice, the technical skills and support tools that are core competencies of third parties are viewed as cost-effective alternatives to internal development. For this reason it is vital that the environment be based on open standards so that third parties can extend the environment to satisfy changing business conditions.

Product Capabilities

As part of its deployment planning, ITG considered carefully how the capabilities of ISA Server would relate to the business requirements at Microsoft. In this context, the following are some of the more significant product capabilities, which ITG became familiar with prior to deploying ISA Server widely within Microsoft.

Multilayer Firewall Security

A firewall can enhance security through various methods, including packet filtering, circuit-level filtering, and application filtering. Advanced enterprise firewalls, such as that provided with ISA Server, combine all three of these methods to provide protection at multiple network layers.

Circuit-Level Filtering

At the circuit level, the ISA Server Firewall service works with virtually all Internet applications and protocols—such as Telnet, mail, news, Microsoft Windows Media™ technologies, RealAudio, and Internet Relay Chat (IRC)—and other client applications. The Firewall service makes these applications perform as if they were connected directly to the Internet. Circuit-level filtering is offered for both firewall and SecureNAT clients.

Circuit-level filtering enables support for virtually all standard and custom Internet applications on the Windows platform. These applications communicate on the network using Winsock and can be supported, unmodified, on client machines that have the Firewall client software installed.

Circuit-level filtering inspects sessions, rather than connections or packets. A session can include multiple connections, providing a number of important benefits for Windows-based clients running the Firewall client software.

Packet Filtering

The packet-filtering capability of ISA Server enables the administrator to control the flow of Internet Protocol (IP) packets to and from ISA Server. When packet filtering is enabled, all packets on the external interface are dropped unless they are explicitly allowed, statically, by IP packet filters, or dynamically, by access policy or publishing rules.

IP packet filtering intercepts and evaluates packets before they are passed to higher levels in the firewall engine or to an application filter. IP packet filters can be configured so that only specified packets will be passed through the ISA Server. This practice provides a high level of security for the network. IP packet filtering can block packets originating in specific Internet hosts and can reject packets associated with many common attacks. IP packet filtering can also block packets destined to any service on an internal network, including the Web proxy, a Web server, an SMTP server, and others.

IP packets filters are static, communication through a given port is always either allowed or blocked. Allow-filters allow the traffic through, unconditionally, at the specified port. Block-filters always prevent the packets from passing through the ISA Server computer.

ISA Server supports dynamic packet filtering, opening ports automatically only as required for communications, and closing the ports when the communication ends. This approach minimizes the number of exposed ports in either direction and provides a high level of security for a network.

ISA Server supports inbound and outbound IP packet filtering. ISA Server's packet filtering also allows for blocking fragments and detecting packet-level attacks against the firewall.

Application-Level Filtering

The most sophisticated level of traffic inspection provided by the ISA Server firewall is the application-level security. "Smart" application filters can analyze a data stream for a given application and provide application-specific processing including inspecting, screening or blocking, redirecting, or even modifying the data as it passes through the firewall. This mechanism protects against known exploits such as unsafe SMTP commands or attacks against internal Domain Name System (DNS) servers. Third-party tools for content screening, including virus detection, lexical analysis, and site categorization, also use application and Web filters to further extend the firewall.

Stateful Inspection

Stateful inspection examines data crossing the firewall in the context of its protocol and the state of the connection. At the packet level, ISA Server inspects the source and destination of the traffic indicated in the IP header and the port in the TCP or UDP header identifying

the network service or application used.

Dynamic packet filters enable the opening of a port only in response to a user's request and only for the duration required to satisfy that request, reducing the vulnerability associated with open ports. ISA Server can determine dynamically which packets can be passed through to the internal network's circuit- and application-layer services. Administrators can configure access-policy rules that open ports automatically only as allowed and then close the ports when the communication ends. This process, known as dynamic packet filtering, minimizes the number of exposed ports in either direction and provides a high level of problem-free security for the network.

Integrated Intrusion Detection

With the help of technology provided by a firm known as Internet Security Systems, ISA Server can help administrators identify and respond to common network attacks such as port scanning, WinNuke and Ping of Death. This technology provides ISA Server with an integrated intrusion-detection mechanism that identifies that kind of attack. The alert also specifies what action ISA Server should take when the attack is recognized, action that may include sending an e-mail message or a page to the system administrator, stopping the Firewall service, writing to the system Event Log, or running any program or script. With additional help from 3rd parties ISA Server can help administrators identify and respond to other common network attacks.

ISA Server implements intrusion detection at both the packet-filter level and the application-filter level. ITG's deployment plans called for using all available intrusion detection.

High performance Web cache

ISA Server has a completely redesigned Web cache that enables it to place cache into RAM. This high-performance Web cache provides greater scalability on the back end as well as a faster overall Web-client response time. This was especially important with respect to the ITG beta deployment of ISA Server because Microsoft employees need fast access to Web content and ITG requires network-bandwidth savings.

Cache Array Routing Protocol

ISA Server uses the Cache Array Routing Protocol (CARP) to provide seamless scaling and high efficiency in an array of multiple ISA Server computers. CARP uses hash-based routing to provide a deterministic "request resolution path" through an array. The request resolution path, based on a hashing of array member identities and URLs, means that for any given URL request the browser or downstream proxy server can know exactly where the requested information is stored in the array. ITG's deployment plans called for configuring ISA Server in arrays to take advantage of CARP.

Chained-Configuration Cache Placement

In this context, the term "chaining" refers to a hierarchical connection between individual ISA Server computers or arrays of ISA Server computers. With chaining, client requests are sent upstream through the chain of cache servers until the requested object is found. During this process, the object is cached at every server until being returned to the client. Consequently, chaining becomes an effective means of distributing server load and fault tolerance.

The chained configuration can be used to position content closer to users who need it, resulting in faster Web-client response times and reduced Wide Area Network (WAN) traffic. ISA Server makes a distributed Web cache possible in which users can obtain Web pages from ISA Server rather than from individual Web sites. ITG's deployment plans called for utilizing the chained-configuration to reduce client computer WAN distance traversal in geographically remote areas of the company.

Active Caching

With a feature know active caching, whereby ISA Server can be configured to automatically update objects in a cache, ISA Server can optimize bandwidth usage by proactively refreshing content. With active caching, objects that are accessed frequently are updated automatically before they expire, during periods of low network traffic.

Active caching is a way to keep objects fresh in the cache by verifying them with the originating Web server before the objects expire and are accessed by a client. The goal is to expedite those client accesses that would ordinarily require a round-trip to the originating server to revalidate the data. Because there is a cost associated with this (in both proxy processing and network bandwidth) the goal is to refresh only those objects that are likely to be accessed in the future by a client.

In contrast, object "popularity" is not a useful criterion for this because many popular pages never expire. This is due to the fact that clients refresh the pages manually to keep the data fresh. In addition, an object may be popular for only a short time. The active-caching code tries to identify objects that follow precisely the pattern of accessed content that would be helped by active refreshing, that is, objects that expire and are then touched again by a client.

Unified Management

ISA Server takes advantage of Windows 2000–based security, Active Directory™ service, VPN, and the Microsoft management Console (MMC). All of these capabilities, especially MMC, help to make administration easier because operations personnel are familiar with it and can manage both the firewall and Web cache from one console .

Using built-in reporting tools, ISA Server supports the running of scheduled standard reports that detail Web usage, application usage, network-traffic patterns, and security. ISA Server provides extensive support for reporting such matters as frequency of Internet access, what is being accessed, and by whom. By alerting and reporting such matters as out-of-boundary activity to administrators, ISA Server can help them to better understand how employees are using the Web. This is helpful for capacity planning and enforcing corporate policies. ITG's deployment plans called for running such reports on a regular basis.

For example, with modification to the Active Directory schema and creation of policy, ISA Server can be configured to run in an array. ITG's deployment plans called for deploying twenty-two different arrays to geographic regions within Microsoft. Arrays allows administrators to apply policy, such as port configuration, at the enterprise level, thereby allowing a single change to be applied to every ISA Server within each array. Active Directory multimaster replication of ISA Server configuration ensures that each server in an array will receive current and automatically updated configurations.

Administrators must control and enforce security policies while simultaneously supporting employees who need Internet access. Thanks to its integration with Active Directory, ISA Server provides comprehensive support for controlling access by user, group, application, destination, content type, and schedule. ITG's deployment plans call for defining enterprise and array-level control.

Enterprise Policy and Access Control

ISA Server also supports the creation of enterprise-level and local array policies, for centralized or local enforcement. ISA Server can be installed as a standalone server or as an array member. ITG initially used standalone servers for the purpose of learning, but later deployed only array member servers in production. For easier management and administration, array members share the same configuration. With modification to the array configuration, all the ISA Server computers in the array are also modified, including all their access and cache policies.

Centralized administration can also mean greater security. Administrative tasks can be performed on one computer and the resulting

configurations are applied to all. This approach helps to ensure that all the servers have the same access policies and is particularly useful at Microsoft, where arrays include many ISA Server computers.

An enterprise can take this centralized management one step further, allowing administrators to implement one or more enterprise policies, which include site and content rules and protocol rules. An enterprise policy can be applied to any array and can be augmented by the array's own policy. This approach enforces enterprise policies at branch and departmental levels while allowing local administrators to further restrict access. ITG's deployment plans called for using enterprise policy to establish corporate standards, which would augment the array policy established by different subsidiaries.

Overview of Legacy Proxy Access

ITG deployed the first version of Microsoft Proxy Server in 1996, while that product was in the beta stage. The benefits of the first deployment were obvious almost immediately. With this deployment employees had a relatively easy means of accessing the Internet from nearly anywhere within the company.

When ITG deployed Proxy Server 1.0, it placed the affected servers within just two data centers because of security and manageability concerns. The first deployment resulted in internal users having to traverse great network distances to gain access to the Internet, causing an increase in WAN usage. Consequently, although the deployment of proxy servers enabled ITG to physically secure and manage the servers, the first deployment did not result in fast Internet access for all employees.

With this in mind, before deploying Proxy Server 2.0 internal networking engineers redesigned the corporate network while converting the backbone to Asynchronous Transfer Mode (ATM). At the same time, they created twenty-two Internet access points at various locations across the company. As part of the Proxy Server 2.0 deployment planning, the engineers evaluated the placement of their Proxy Server 1.0 deployment and determined that by repositioning Proxy Server 2.0–based servers to locations serving as Internet access points, they could increase the speed of Internet access and reduce the requirements on the WAN.

Deployment Scope and Goals

To formulate the ISA Server deployment strategy, ITG defined a deployment scope and deployment goals, which defined clearly what was to be accomplished and allowed everyone on the team to work toward common objectives.

To simplify both the planning and execution of the deployment, ITG divided the project into six smaller deployments, each of which was characterized by a unique configuration designed to satisfy specific business requirements. These six deployments defined the entire beta deployment project.

1. **Corporate-to-Internet Access.** This deployment required the replacement of an existing Proxy Server 2.0–based installation with ISA Server. In this deployment, ISA Server was configured to run in integrated mode, allowing ISA Server to function in both Web-cache and firewall modes simultaneously.

   This deployment required that seventy computers running Proxy Server 2.0 be transitioned to ISA Server. The existing computers had been configured in twenty-two arrays and deployed throughout the company so that employees working from any region could access the Internet securely while working. The business requirement for this deployment was twofold: (1) to provide employees with secure, fast access to the Web and (2) to establish a tight feedback loop to ensure that ISA Server would be enterprise-ready on release.
2. **Chaining Proxy.** This deployment required ITG to configure Microsoft's internal environment by placing caches closest to users in chained configurations. Chained configurations would better support several of the hierarchically interconnected subsidiaries at Microsoft that do not have existing Internet access points. The primary business requirement of this deployment was to make better use of the Microsoft corporate WAN while providing faster Web access to employees who depended on the performance of the network between corporate headquarters and their subsidiary.
3. **Firewalls in the Extranet.** This deployment, which as of this writing is not yet complete, will require ITG to deploy ISA Server within the Microsoft corporate extranet, a highly secured network infrastructure used for establishing secure network connectivity between Microsoft and its business partners and suppliers.
4. **Firewall client deployment.** This deployment required ITG to deploy the ISA Server Firewall client software to more than twenty thousand desktop computers, which would give those computers Winsock proxy access to the Internet. The primary business requirement of this deployment was to test the firewall client before it was released.
5. **Firewall deployment in a subsidiary.** This deployment (in a Silicon Valley company that was later acquired by Microsoft) required ITG to replace a legacy UNIX-based firewall solution with ISA Server. In this deployment, ITG configured ISA Server to run in firewall mode. This deployment uses SecureNAT and takes advantage of network-routing rules rather then client-side application rules, thus providing proxy service to all IP-based clients including Macintosh and UNIX clients running Socks.

   The primary business requirement of this deployment was to replace a legacy infrastructure (deployed by a third party) with the corporate standard while also providing real-world feedback on the firewall components in ISA Server.
6. **H.323 Gatekeeper.** This deployment required ITG to design, deploy, and configure an ISA Server-based infrastructure that would use the H.323 protocol. Although more planning was required in this deployment, it served as a critical test bed on which to validate and confirm that ISA Server, combined with H.323 Gatekeeper, would allow business to communicate with one other over the Internet using Microsoft NetMeeting® conferencing software. The primary business requirement of this deployment was to provide feedback on the opportunities provided by ISA Server that would allow Microsoft to develop new methods of doing business over the Internet.

Deployment Team Structure

To plan and implement the deployment of ISA Server at Microsoft, ITG selected a team of peers based on their knowledge and skills in the following areas:

- **Project management.** Team members needed experience in providing leadership and a business focus. They needed to have performed all the traditional functions expected of a project manager, such as project scheduling, reporting, and risk analysis. This is because they would be ultimately accountable for the outcome of the deployment.
- **Program management.** Team members needed to develop a "product mindset" based on their detailed understanding of how ISA Server functioned. The program manager would need to communicate problems, estimate the impact of such problems on the project, and then develop plans to resolve them.
- **Infrastructure engineering.** Team members needed expertise in the design and capabilities of the existing infrastructure. They also needed a passion for technology. Infrastructure engineers would provide the knowledge and capability to configure the computing environment; they also would formulate plans on how ISA Server should be configured.
- **Product development.** At least one team member needed expertise in the internal software design of ISA Server along with a zero-defect mindset. Team members would need to interact with a representative of the ISA Server development team through a feedback loop for reporting issues.
- **Operations helpdesk.** One team member needed to be able to provide understanding of the internal user mindset and working habits. This person would act as an advocate for some employees who would rely on ISA Server to carry out their day-to-day job functions and other employees who would be performing any needed monitoring and client-side troubleshooting.

Figure 1 illustrates the team of peers used to deploy ISA Server at Microsoft.

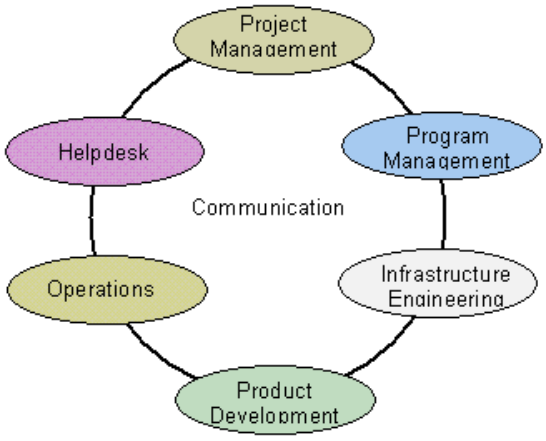

Figure 1 Team of Peers

Each team member brought a slightly different perspective and an entirely different set of skills. Each also had his or her own goals for deploying ISA Server. Figure 2 illustrates the goals specific to each member of the deployment team.

Team Role	Goal
Project Management	Deliver within business constraints
Program Management	Deliver within project constraints
Infrastructure Engineering	Deliver infrastructure design
Product Development	Deliver product-issue resolution
Operations	Deliver installation, maintenance and support
Helpdesk	Deliver internal desktop support

If your browser does not support inline frames, click here to view on a separate page.

Figure 2 Goals specific to each member of the deployment team

Risk Management

The beta deployment strategy for ISA Server required ITG to identify areas of risk that might result in extended server downtime as well as scenarios that might result in compromised internal network security. ITG identified the following risks the team needed to address so as to avoid any adverse impact on deployment:

- **Server downtime.** This risk could not be avoided because at some point ITG would need to perform maintenance on each server. Server downtime was especially an issue since the product was being deployed while in the beta stages. ITG designed its infrastructure so that employees would be unlikely to notice server maintenance, or issues that were being addressed with the beta. As soon as the ITG deployment team encountered a problem on the back end, they removed the server having the problem and distributed the load using Network Load Balancing (NLB). This special infrastructure enabled ITG to keep servers offline for troubleshooting for as long as was needed to discover the cause of each problem with minimal client impact. As soon as it became possible to utilize ISA Server-based distributed caching, ITG removed NLB.
- **Unauthorized network access.** ITG avoided this potential risk by carefully reviewing infrastructure designs and product capabilities to ensure that ISA Server was properly configured. By default, ISA Server denies all requests, requiring administrators to manually allow protocols that are needed by the business. However, ITG still considered it prudent to use a computer lab environment to practice configuring ISA Server and to prove concepts that would later be put to use in the production environment.

Proof-of-Concept

The term “proof-of-concept” refers to activities intended to prove that a documented design or other plan will work in the production environment, before the design is actually put to use in production. The goal of each proof-of-concept is to validate that a proposed design will work in a scaled down test environment that closely matches the production environment. ITG developed each proof-of-concept in order to check its deployment plans before using them to carry out a wide internal deployment.

One proof-of-concept allowed ITG to deploy the beta of ISA Server in a lab located in Redmond, Washington, and then configure that environment so that it would serve as an interim Internet firewall to a Silicon Valley acquisition. In this case, ITG performed the proof-of-concept in the lab long enough to verify and confirm that the proposed changes would not adversely affect any employee who depended on the production environment.

In each case, ITG found proof-of-concept planning helpful in managing risk and in providing an environment conducive to learning. Proof-of-concept was especially important in this case, because ISA Server documentation was not available when the beta deployment began.

The ITG team used other proof-of-concepts to test processes they would later use to remove Proxy Server 2.0–based computers from their respective arrays, install ISA Server, and manage the arrays using the MMC administration tools included with ISA Server. Because each proof-of-concept already had been thoroughly tested in the lab environment, upgrades went smoothly in production.

Capacity Planning

Determining an appropriate server size is key to maximizing server performance and utilization while minimizing cost. As part of its ISA

Server deployment planning, ITG defined “small,” “medium,” and “large” hardware-configuration standards that it would use based on the number of employees that were expected to access each server.

For example, ITG used the large hardware configuration at corporate headquarters in Redmond, where more than 40,000 desktop computers were in regular use. In that deployment, ISA Server was installed on 17 servers configured as a single array, running on an ATM network. Clients accessed them through the network at 200 Mbps, with each server’s network interface card running in full-duplex mode.

In reviewing that deployment, ITG found that on a typical day the array easily handled 284 GB of inbound and 56 GB of outbound HTTP requests, with the HTTP protocol accounting for roughly 97 percent of traffic. The array also easily handled over 40 million Internet-based URL requests each day, with peak usage occurring around lunchtime, when the array handled roughly 25 GB of HTTP requests each hour with an average processing time of just 1.4 seconds for each request. During peak loads, ITG was handling an average of 150 HTTP requests per second per server.

It is important to note that the ISA Server beta deployment at corporate headquarters, along with the selection of a given configuration size for that deployment, was based on long-tested practices at Microsoft. In keeping with these practices, ITG conducts testing on a wide variety of hardware to determine which equipment will be deployed at the Microsoft data centers. ITG considers such criteria as performance, equipment mean time to failure, remote diagnostic capabilities, cost, scalability, and support from third parties. ITG also strives to select a hardware platform on which business units can easily host their line-of-business applications without needing additional components for at least six months following their initial deployments.

Another ITG best practice is to deploy identical hardware to all servers within an array. Based on this, table 1 illustrates the ISA Server hardware standards that were put in place for the beta deployment at Microsoft. (Note that ITG used its Proxy Server 2.0 hardware standards for ISA Server as well, so no additional hardware was required.)

Table 1 Small, Medium and Large ISA Server Hardware Standards

	Xeon CPUs / per server	Memory / per server	Disk Space dedicated for URL cache / per server	Concurrent Users / per array	Number of Servers /per array
Small	2	256 MB	36G	1,500	2
Medium	2	512 MB	54G	5,000	4
Large*	4	512 MB	63G	40,000	9

* Approximately one /each per 5,000 users

Server Placement

In the ISA Server beta deployment at Microsoft, server placement depended on the primary goal at each site, existing network topology, and data-center taxonomy.

Site Goals

Deploying ISA Server primarily for the use of its high-performance Web-caching capabilities required ITG to consider where client-side requests originated and where each Internet access point was located at Microsoft. Also part of this consideration was the question of whether the physical distance between the two locations necessitated a chained configuration for the deployment.

On the other hand, when the deployment was primarily to take advantage of ISA Server’s Internet firewall capabilities, ITG needed to position the ISA Server computers as close as possible to Internet access points so that network traffic from the Internet would pass through ISA Server before entering the corporate network.

Network Topology

ITG also considered the physical design of the network in deciding where to place the ISA Server computers. This consideration was especially important because ISA Server functionality depends on a reliable and scalable network. Microsoft’s computer network is linked using a traditional (ATM) network topology. Where there is a high concentration of employees, there is a hub in the network topology, which will almost always provide Internet access to those employees. Where there are relatively fewer employees, there is a spoke in the network topology.

Reviewing the Microsoft network topology helped ITG to ensure that ISA Server computers would enjoy sufficient network connectivity between them and that Internet access points were available geographically. Figure 3 illustrates the deployment of three arrays to locations serving as Internet access points. Twenty-two such deployments were required so as to provide employees throughout the company with the fastest possible Internet access.

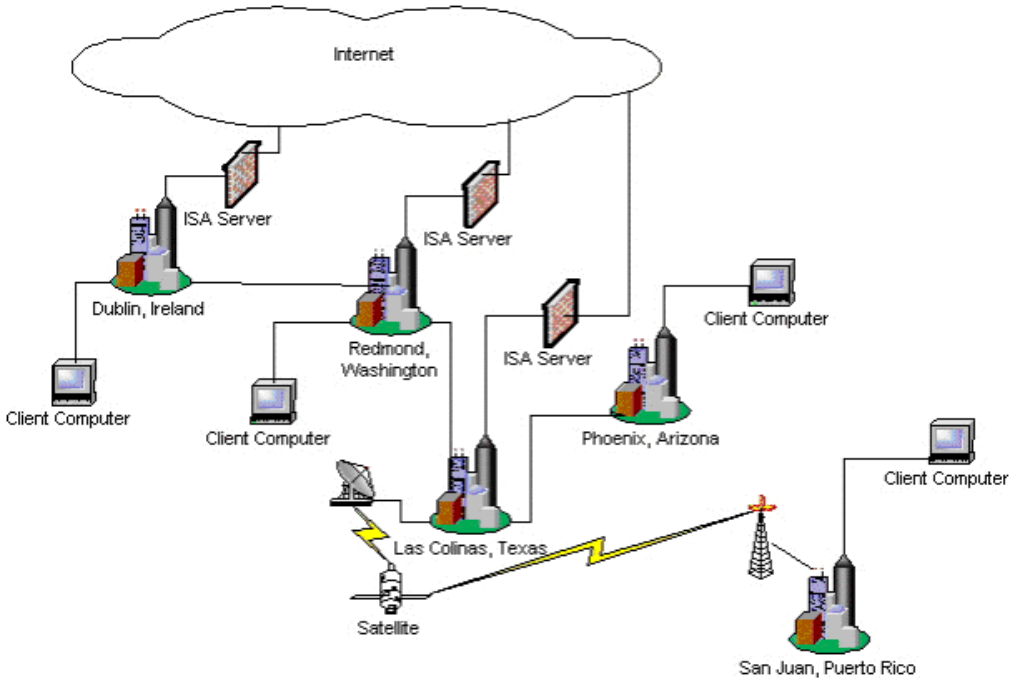

If your browser does not support inline frames, click here to view on a separate page.

Figure 3 Typical ISA Server Placements Providing Corporate to Internet Access

Data Center Classification

According to ITG's taxonomy there are three classes of data centers: enterprise data centers, regional data centers, and site data rooms. The placement of data centers heavily influenced the placement of ISA Servers. The geographic positioning of these data centers corresponds with the corporate ATM spoke-and-hub network topology. Enterprise data centers are located where the majority of employees are located. Regional data centers are geographically dispersed and function primarily to house the networking equipment needed to connect site data rooms with enterprise data centers. Site data rooms are typically located at the subsidiaries. The majority of site data rooms are networked directly to a regional data center using a network speed greater than 256 Kbps.

As for client computers, they are always configured to access the closest ISA Server computer on the internal network where the network path determines the distance between client and server. For example, Windows 2000 Professional client computers located in San Juan, Puerto Rico, access ISA Server located in Las Colinas, Texas, because that is the closest ISA Server deployment. Accordingly, for the ISA Server beta deployment, a team of network engineers determined each physical location that would serve as an Internet access point throughout Microsoft and executed the deployment to take advantage of those existing locations.

The following discusses some of the primary differences of the data-center classifications and some considerations peculiar to the deployment of ISA Server into each classification.

Enterprise Data Centers

Microsoft has enterprise data centers at two locations: Redmond, Washington, and Dublin, Ireland. With staff available on location 24 hours a day, seven days a week, these facilities serve the majority of the company's employees. The Redmond center serves those who work in Washington State, and the Dublin center serves those who work in Europe. Each data center has Internet access and an ISA Server array.

In Redmond, ITG deployed a single array containing 17 computers, each running ISA Server in integrated mode. In this area of the company roughly 40,000 client computers request over 40 million URLs from the Internet each day. Data-center staff monitor the array activity and add computers to the array as required to maximize performance.

ISA Server configurations in Redmond require that each ISA Server computer be configured solely for that purpose and running in integrated mode so as to provide a Web cache and firewall on the same server. That is, servers running ISA Server in this area of the company do not provide additional services such as file, print, or domain control.

Regional Data Centers

At Microsoft, regional data centers have fewer servers than enterprise data centers because they tend to serve far fewer employees than do the company's two enterprise data centers. Generally, the regional data centers run smaller, distributed applications than those deployed in larger data centers and they generally are providing Internet access to each region. Also in contrast to the enterprise data centers, the regional data centers typically do not have dedicated, onsite staff. Instead, they are usually administered remotely with the help of the Terminal Services software included in Windows 2000 Advanced Server.

ISA Server installations located in regional data centers are used less heavily than those located in enterprise data centers, because fewer employees access them. Accordingly, ITG configured ISA Server hardware at the regional data centers to follow the medium or small standard.

Site Data Rooms

Nearly every Microsoft subsidiary office has a site data room capable of securely storing a small Windows 2000 Advanced Server–based infrastructure to serve the needs of employees working there. The infrastructure consists of several computers running the Windows 2000 Advanced Server network operating system, providing Active Directory services; DNS; file, print, and in a few cases even remote access services; proxy service; and e-mail. A typical server is configured to function as a domain controller, global catalog server, DNS server, and Dynamic Host Configuration Protocol (DHCP) server simultaneously. This infrastructure enables employees to (1) log on to a network securely and access shared data within the subsidiary, (2) collaborate with other employees in the subsidiary, and (3) print locally stored documents with or without physical network connectivity to a larger regional facility located elsewhere. The DHCP infrastructure is used to configure client computers to access an ISA Server array that serves the region.

Site data rooms have been instrumental in helping ITG to support a geographically remote workforce. Having a small infrastructure

available within each subsidiary has helped to reduce network latency and network traffic across the WAN, because the infrastructure of each region automatically configures client computers located there to utilize geographically positioned arrays.

Strategy

As part of its ISA Server deployment strategy, ITG used formal checkpoints to measure progress and form consensus on project direction. ITG also established major milestones so as to segment the deployment into distinct phases.

The first phase of the deployment involved understanding the business requirements and capabilities of ISA Server, performing initial proof-of-concept testing, and developing a formal project scope with deployment goals. ITG performed each objective in concert with the others. Project management identified business drivers, infrastructure engineering readied system designs, product development provided information about the product, and operations staff began learning how to configure ISA Server.

The second and subsequent phases involved incrementally deploying a larger number of computers running ISA Server for each of the deployments identified in the project scope. In each phase of the deployment, ITG worked as customer advocates by reporting any issues to the product-development team for resolution before release. This process was repeated continually alongside the many internal beta releases.

Executing the Deployment

To install and configure ISA Server, the ITG deployment team depended on earlier planning in which they had identified how the product was to be used. The team then executed deployment according to the business requirements they established while planning each of the six deployments named earlier. The execution involved configuring ISA Server according to agreed upon standards, the adoption of enterprise and array policy the execution of standard reporting and monitoring, as well as the execution of other more specific configurations.

Standard ISA Server Configuration

At Microsoft, the widespread use of change control and of well-defined corporate standards for server configurations means that when members of the deployment team began their work, they had a predictable and easy-to-manage infrastructure already in place. Also in place were well-defined physical-asset and server-configuration standards that would prove beneficial to the day-to-day support of the ISA Server deployment.

For instance, the team configured all ISA Server computers to use Redundant Array of Inexpensive Disks (RAID), as illustrated in Table 2. (Note that the table represents the standard used by ITG's "large" ISA Server computers, and partition sizes vary depending on the specified amount of cache.) On high-end servers the team configured System, URL, and Log directories so that each of them would use its own physical partition. On smaller servers the team configured System and Log directories to share a single physical partition. In either case, the URL directory always had its own dedicated partition.

Table 2 RAID Configurations

Partition	RAID 1	RAID 5
System	Yes	No
URL Cache	No	Yes
Log Files	No	Yes

Each ISA Server computer ran the Windows 2000 Advanced Server operation system and two additional services, as illustrated in Table 3, provided by Windows 2000 Advanced Server.

Table 3 Windows 2000 Services

Service	Why Enabled?
Simple Network Monitoring Protocol (SNMP)	Required by Compaq Insight Agents
Terminal Services, in administration mode	For remote administration

The team disabled the following Windows 2000 Advanced Server services in the ISA Server deployment because they were not required in the environment:

- Computer browser
- Distributed File System (DFS)
- Distributed link tracking
- Fax Service
- License Logging
- Telephony

As a best practice, ITG disabled these services so as to reduce the complexity of the environment, simplify any needed troubleshooting, and increase system resources available to the required services. Table 4 illustrates the configuration of the external network interface card.

Table 4 Configuration of external Network Interface Card

Network Interface Card	Configuration
01a IP	Default
01b IP	Default
Default Gateway	Default
"Client for Microsoft Networks"	Disable
"File and Printer Sharing"	Disable
"Register this connection in DNS"	Disable
NetBios	Disable NetBios over TCP/IP

The standard ISA Server deployment called also for the following configurations:

Table 5 Additional configurations for the ISA Server deployment

ISA Server	Configuration
Mode	Integrated
Cache	20 GB on each server
IP filters	DHCP disabled. All others left in default state
Reporting	Weekly on every server
Intrusion detection	Enabled

At Microsoft, all servers under the control of ITG use a formal approach toward change control, and the servers installed for the beta deployment of ISA Server were no exception. The change-control process, which results ultimately in the adoption of a new standard configuration, involves three levels. Each level, also known as a "project," is denoted by the color gold, silver, or bronze.

For example, the configuration contained in the gold project represents the current corporate standard for Windows 2000 Advanced Server. For this reason, the majority of production servers at Microsoft use the gold project, which contains Windows 2000 Advanced Server, the current service pack, and hot fixes that resolve issues considered characteristic of the Microsoft environment. The configuration contained in the silver project represents a proposed change to the corporate standard, for example, to support new hardware, tune performance, or include a new services pack or hot fix. All changes to the corporate standard transition through the silver project so that engineers can introduce the changes into the production environment in a very controlled manner. A few dozen production servers at Microsoft use the silver project for testing purposes. If those test servers demonstrate better performance and reliability than servers using the gold project, then the proposed change to corporate standard will be approved. As part of the approval processes the configuration contained in the silver project will replace that contained in the gold project.

The configuration contained in the bronze project represents changes that are proposed to the corporate standard but not ready for production use. The bronze project contains a configuration that will be deployed in a lab environment and for testing purposes only. The lab environment resembles a scaled-down production environment, enabling ITG to determine whether the new configuration will be stable once it is put to limited production use. Each proposed change to the corporate standard enters the change-control process when it is added to the bronze project. The changes then progress to the silver project for limited production testing before being added to the gold project, making the changed configuration the new corporate standard. The duration of testing at each stage is generally at least three days of continuous use. Under ideal circumstances, a much longer duration is used. Figure 4 illustrates the use of these three configurations.

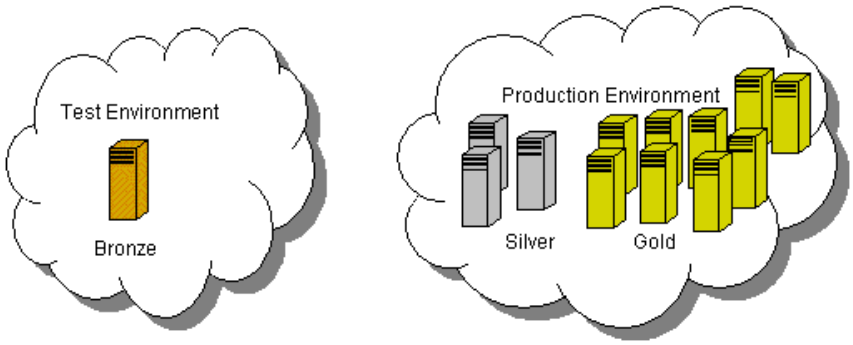

If your browser does not support inline frames, click here to view on a separate page.

Figure 4 Usage of configurations under change control

Migrating Proxy Server 2.0 to ISA Server

When it came time to deploy ISA Server to computers already running Proxy Server 2.0, the deployment team found that additional server preparation was needed before they could proceed. So team members performed the following steps:

1. Removed each server from its Proxy Server 2.0–based array.
2. Uninstalled Proxy Server 2.0 from each server. (Note that at the time no support for upgrades was available.)
3. Uninstalled Internet Information Server (IIS) 5.0 from each server.
4. Formatted the logical drive containing the URL cache.

The reason for Step 3 is that Internet Information Services and ISA Server compete for port 80, the default for Internet Information Services, and also the HTTP standard used by ISA Server. Since ISA Server functionality does not depend in any way on IIS there was no reason to leave it on the servers on which ISA Server would be deployed.

The First Deployment

For the first ISA Server deployment, the team had to initialize the environment, install ISA Server, and configure the environment using the administration tool included with ISA Server 2000 Enterprise Edition. Figure 5 illustrates the steps required to do this:

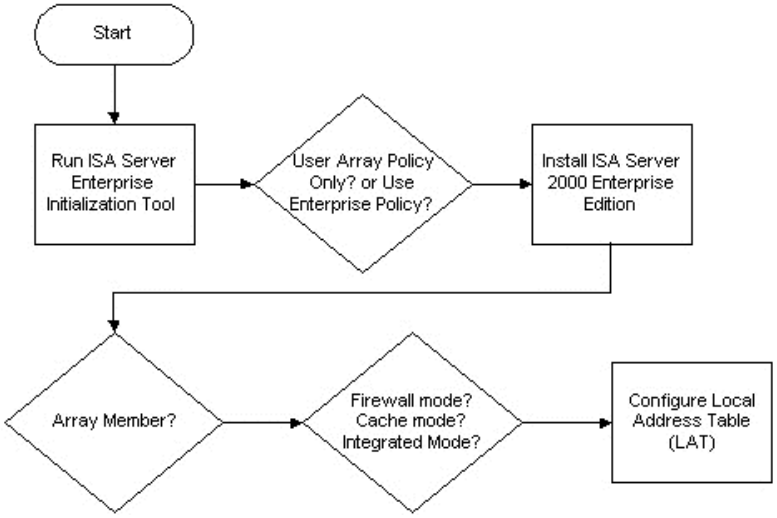

If your browser does not support inline frames, click here to view on a separate page.

Figure 5 Initialization of environment and first deployment of ISA Server.

Note that initializing the environment required the team to run the initialization tool to start the Active Directory schema update. When schema modification is complete, a dialog box will indicate that it is ok to install ISA Server as a domain array member. The initialization tool must be run before either Enterprise Policy or Array Policy (stored in Active Directory) can be used. The utility works by updating the Active Directory schema with a set of base rules that affect how Enterprise Policy applies to Array Policy.

Next, the team initialized its Active Directory schema to allow Enterprise Policy to apply to Array Policy. This would allow Enterprise Policy to be augmented by the array's own policy and allows administrators at regional subsidiaries to adopt governing policies that apply to ISA Server computers deployed in their region while also inheriting the base rules defined by Enterprise Policy.

Without updating the Active Directory schema, ISA Server can only run in standalone mode, which does not permit taking advantage of distributed caching (CARP) and applying multilevel policy to servers in an array. A computer running ISA Server Enterprise Edition in standalone mode can still be promoted to run in the array mode, but the server's configuration will be overwritten by configurations specified in Enterprise Policy and Array Policy. Table 6 illustrates the relationships among these options.

Table 6 Relationships among policy types, distributed caching, and the decision to update the Active Directory schema

Policy	Update Schema?	Distributed Cache	Standalone
Enterprise	Yes	Yes	No
Array	Yes	Yes	No
Standalone	No	No	Yes

As part of their job, team members had to consider how Enterprise and Array Policy should be implemented as well as how such a decision might effect later administration. As part of this consideration the team had to become familiar with how ISA Server administrators create rules governing protocol and packet filters at the array level. Array Policy applies only to the ISA Server computers in the array.

The deployment team decided to configure the environment to permit Enterprise Policy to apply to Array Policy so that the policies could be combined as business requirements necessitated. With the ISA Server initialization tool, the team also provided configuration support for defining a new policy, inheriting a configuration from an existing policy, and using packet filtering at the array level. Figure 6 illustrates this relationship.

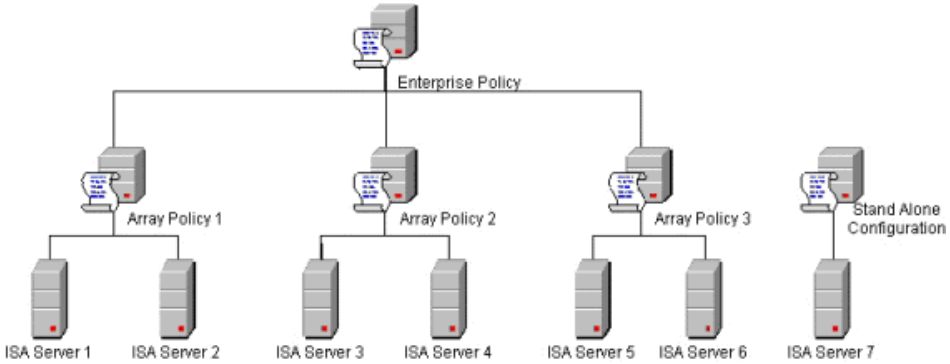

If your browser does not support inline frames, click here to view on a separate page.

Figure 6 The relationship between Enterprise and Array Policy

Once they had finished updating the Active Directory schema with the ISA Server initialization tool, team members were ready to use the ISA Server installation program to set up the first server. When schema modifications were complete, the initialization tool presented a dialog box that indicated it was ok to begin the ISA Server installation. As part of the installation, the team configured the first server as a domain-array member, running in integrated mode.

The last few steps of the installation required that the team define a Local Address Table (LAT). As part of the process, the team selected a range of IP addresses to include in the LAT. ISA Server uses the LAT to distinguish between internal and external networks.

After completing the ISA Server installation, the team ran the product's administration tool to configure the environment so that client computers could begin using ISA Server. Figure 7 illustrates the steps ITG followed while using the administration tool to configure its environment:

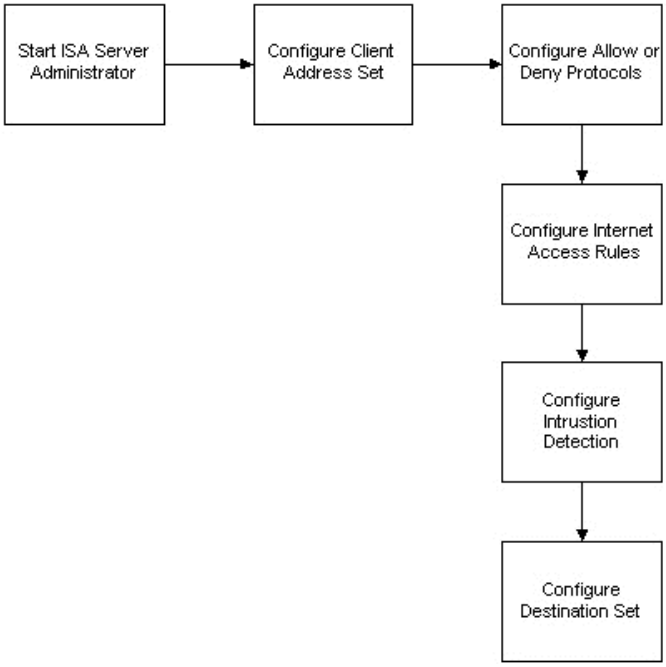

If your browser does not support inline frames, click here to view on a separate page.

Figure 7 Steps taken to configure the initial ISA Server environment

The administration tool supports the configuration of a client address set, which can be based on individual IP addresses or ranges of IP addresses. The client address set specifies the client computers that are authorized to use ISA Server. At Microsoft, almost all employees require regular access to the Internet, so the team configured the client address set using ranges of addresses that would satisfy this requirement.

By default, ISA Server denies all protocols, which means the administration tool must be run specifically to allow the needed protocols. To keep the system as secure as possible, the team took advantage of that restriction by using Enterprise Policy to define only those protocols that were needed at Microsoft and deny all others. The corporate security group used the information to mitigate any risk arising from the chance that the application requiring the protocol might turn out to be a Trojan horse. In a worst-case scenario, such an application could collect sensitive information from internal systems and then use the enabled protocol to transmit the information to an external location. This meant that before the team would “allow” an application, there needed to be some assurance that it had been created by a trusted developer or organization and was needed to meet a legitimate business requirement. To determine this information, the team distributed an employee questionnaire that requested the information, including:

Application requiring the protocol:

- The individual or entity that holds the copyright on the application
- The method by which the Microsoft employee obtained the application
- How other Microsoft employees can obtain the application
- What the application does and how it works
- How the Microsoft employee expects the application to function through the firewall

Connections:

- Initial port number
- UDP, TCP, or both
- Secondary connections
- The systems from which most of the connections will come
- What those systems will most likely be connecting to

Project:

- The number of employees expected to use the application daily
- Estimated number of bytes transferred inbound and outbound daily
- Transaction frequency—all at once or throughout the day

To further secure the ISA Server deployment, the team used Windows 2000 groups and Internet access rules to limit the number of employees who could access allowed protocols. Where applicable, the team applied a schedule to each allowed protocol that would permit it to function only during a specified time period each day. Also where applicable, the team applied a destination set to the allowed protocols, which can block internal clients from accessing specified Internet sites.

Finally, the team configured intrusion detection to support the detection (and notification to administrators) of commonly used intrusion techniques.

Subsequent Deployments of ISA Servers

Once the ISA Server deployment team had implemented the first ISA Server deployment, including the establishment of Enterprise Policy, well-defined goals, and deployment scope, it was time to implement subsequent deployments.

As of this writing, team members are deploying an ISA Server array to each of the twenty-two locations within Microsoft that have Internet access points. Team members are using the Enterprise Policy they established during the first deployment to augment the policy of each of these arrays. In addition, they are using Array Policy to further restrict access by denying protocols that Enterprise Policy had allowed when those protocols are not needed by specific arrays (for example, when employees in certain subsidiaries do not require all the protocols that corporate headquarters needs). Active Directory multimaster replication automated most of the configurations used by subsequent deployments of ISA Server by automatically configuring those servers with settings established by

Enterprise Policy.

Configuring the Internet Explorer and Firewall Client

In both initial and subsequent deployments of ISA Server, team members had to make the necessary configurations available to Internet Explorer and the Firewall client (included with ISA Server) before ISA Server computers could access the Internet. That work, which involved configuring two .DAT files, is detailed in the following.

Configuring the Wpad.dat file

The WPAD.DAT file is used by Internet Explorer to obtain needed information that allows the client browser to use ISA Server's proxy service for Internet access. Clients running Internet Explorer for accessing the Internet may be configured using either DNS or DHCP, but for the ISA Server deployment, team members used DHCP to provide clients with a WPAD.DAT configuration. That's because there are far more DHCP scopes at Microsoft than there are DNS zones. Moreover, with over 130,000 geographically located clients distributed among more than 90 subsidiaries, DHCP would provide the granularity needed to configure clients to use the geographically nearest ISA Server array.

In this case, team members configured DHCP servers in each subsidiary to provide down-level clients with an appropriate WPAD.DAT file using the custom scope option 252. Once the configuration was complete, browsers could detect which DHCP server is being used by a given client computer, which then would request the WPAD.DAT configuration from the appropriate DHCP server.

For the overseas part of the ISA Server beta deployment, the team installed the product on servers in Dublin, Ireland; Munich, Germany; London, England; Les Ulis, France; and Zurich, Switzerland. Each of these servers joined an array responsible for serving each subsidiary. Client computers running Windows 2000 Professional and located in each subsidiary were then configured to use the local array using DHCP scope option 252. With DNS such a configuration would have been impossible because many of these regions constitutes a single domain. The custom option supports specifying the location of the WPAD.DAT file using the following example syntax:

http://ProxyArrayName/wpad.dat/

where `ProxyArrayName` specifies the ISA Server array to be used by clients that obtain information from the DHCP provider.

Team members also registered the array name in DNS, which would be used to perform a round robin with ISA Server computers that join the array for the purpose of load balancing all client requests. Finally, they used a LAT in Internet Explorer to distinguish between internal and external IP addresses and to bypass ISA Server for local intranet access.

Configuring the Wspad.dat file

Many Winsock applications assume that the client computer is connected directly to the Internet and that the computer has an IP address on the Internet. Such applications encounter difficulty when they are used within the Microsoft corporate network, because they cannot obtain a configuration that will allow them to access the Internet by way of ISA Server. This is because the IP packets are not properly addressed to use a proxy server. Some applications will request an Internet address that cannot be obtained unless the application uses a proxy. This means that for Internet access to occur, either the packets must be intercepted and redirected to a proxy or the client must be configured to send the packets directly to the proxy.

Client computers that require Winsock proxy access to the Internet use a Firewall client included with ISA Server. ISA Server supports restricting client-firewall access by domain name, IP address, and subnet mask.

The WSPAD.DAT file provides information needed by the Firewall client to use the ISA Server Firewall service for Internet access.

Employees who require the Firewall client download it from an internal Web site. When the Firewall client is installed on the desktop, there is a default checkbox that enables automatic configuration. This option provides support for querying the WSPAD.DAT file to update the firewall-client configuration.

The Firewall client provides support for such applications as MSN® Messenger, Windows Media Player, Microsoft Outlook® messaging and collaboration client, Hotmail® web-based email service, and many others that require Winsock for Internet access. The deployment team did not need to configure any Winsock applications to use a gateway, because Winsock analyzes local Winsock network traffic and performs lookups against the LAT, both transparently. Winsock determines an appropriate route for network traffic based on addresses within the LAT.

The LAT that the members of deployment team defined when they performed the first installation of ISA Server is stored on each installation of ISA Server and is sent to the Firewall client at regular intervals. The Firewall client uses the configuration contained in the LAT to determine client network traffic that needs to be redirected to the Internet. The Firewall client disables itself automatically when the client computer connects directly to the Internet. This feature is beneficial to roaming computers that periodically establish a direct connection to the Internet.

Managing the Environment

The ISA Server deployment team managed the environment primarily with the help of the administration tool included in ISA Server 2000 Enterprise Edition. The administration tool supports the viewing and changing of configurations of many ISA Server computers at once by treating them as a single logical entity.

Comprehensive Reporting Capabilities

Once team members installed ISA Server, they (along with other administrators and security staff) relied on the product's standard-reporting capabilities for managing the environment. For example, they obtained information on protocol usage through a graphical summary and detailed report, which includes the following information in tabular format:

- Protocol
- Port
- Unique users
- Requests
- Percent of total requests
- Bytes In
- Percent of total bytes in
- Bytes out
- Percent total bytes out
- Total bytes
- Percent of total bytes

The team also relied on a graphical summary and detailed report from ISA Server on the X users generating the largest volume of

network traffic. The ISA Server standard report contains the following information sorted in ascending order of usage:

- User
- Requests
- Percent of total requests
- Bytes in
- Percent of total bytes in
- Bytes out
- Percent of total bytes out
- Total bytes
- Percent of total bytes

Using another report, team members obtained a graphical summary and details on the top X sites accessed through the proxy array. In this report, ISA Server provides the following categories:

- Site
- Unique users
- Requests

The team used other standard reports on array utilization as well, including summary reports in graphical format and a detailed report in tabular format on the following topics:

- Cache performance
- Traffic summary
- Daily traffic summary

Based on their experience, team members found the standard reporting capabilities of ISA Server to be comprehensive and, because they are Web-based, easily accessible. In just a few minutes team members were able to obtain information that in the past might have required hours of searching through log files. For even more extensive reporting capabilities, ISA Server can be extended with third-party products.

The team configured weekly reporting to occur on every ISA Server to determine how employees were using the environment and how hard ISA Server had to work to provide the fast access to information those employees needed. Figure 8 illustrates the Web-based reporting capabilities in ISA Server.

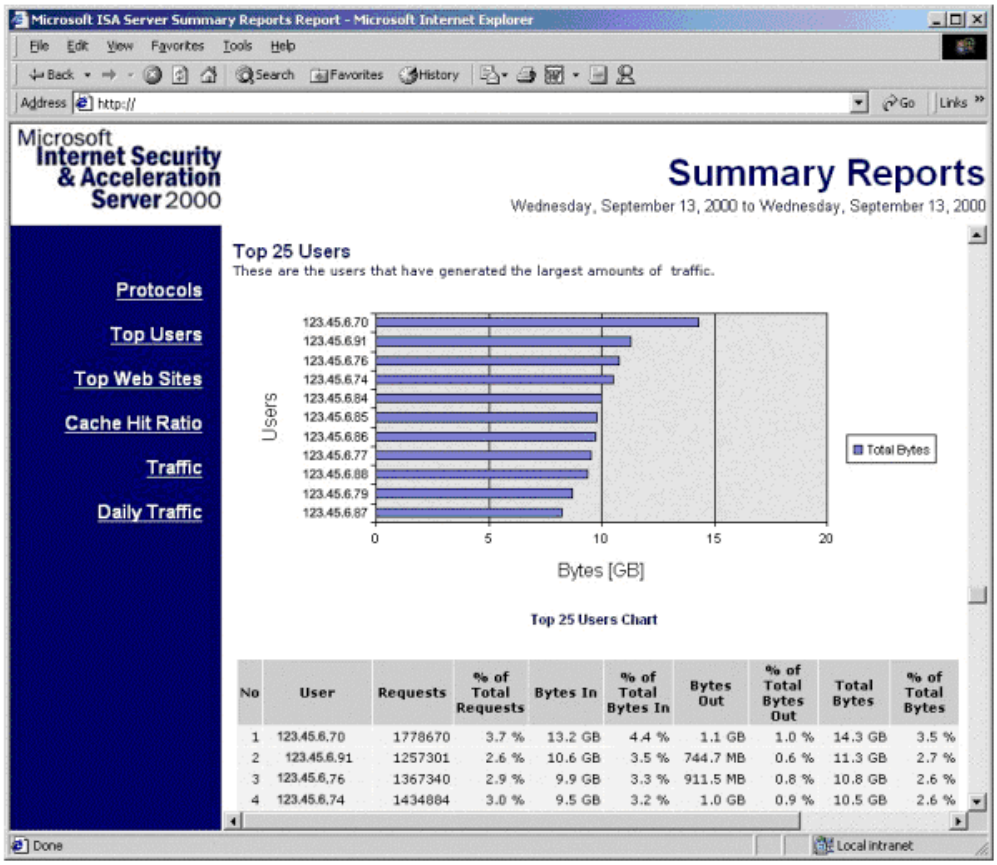

If your browser does not support inline frames, click here to view on a separate page.

Figure 8 Web based report on top users

Client Fail Over and Load Balancing

As part of their work, members of the deployment team took advantage of the fact that in an ISA Server environment, built-in capabilities for smooth client failover permit clients to use an alternative ISA Server computer when maintenance is needed on the back end.

For example, ISA Server uses the Windows Network Load Balancing (NLB) Services of Windows 2000 Advanced Server to provide fault tolerance, high availability, efficiency, and performance through the clustering of multiple ISA Server computers. NLB is especially useful in firewall, reverse-cache (Web publishing), and server publishing deployment configurations.

Team members were able to take advantage of the fact that all servers running ISA Server in an array keep a reference of other

servers that are available. When a server running ISA Server is unavailable others in the array detect it automatically, and an entry in the master mspclnt.ini file is updated to reflect this. While a server is inaccessible, clients will not be directed to it until the server has been returned to an operating state.

Backing Up ISA Server

Third-party products formed two integral components of the team's backup strategy: Veritas Backup Executive Network Storage Executive, for backing up and restoring data, and Arcus Data Security, for storing tapes securely off-site.

Monitoring ISA Server

After deploying ISA Server in the production environment, team members monitored each server closely to determine how well it was performing under real-world conditions. The monitoring provided the team with information necessary for evaluating the environment and determining how well it would scale. This monitoring focused primarily on the following:

- Counters for standard hardware performance (including Processor, Memory, Physical Disk, Paging File, and Network Interface)
- Counters for ISA Server processes
- Counters for ISA Server services

Scenarios

Each identified business requirement required a slightly different configuration. The following section of this paper details the various configurations used at Microsoft.

Corporate to Internet Access

As of this writing, team members have completed their corporate-to-Internet access deployment at corporate headquarters and are well on track to completing deployment of ISA Server in other regions throughout the company. To establish corporate-to-Internet access, the team deployed twenty-two arrays to various corporate locations having Internet access points. These arrays were deployed geographically so as to serve the needs of this geographically diverse workforce.

Corporate-to-Internet access enables more than 55,000 Microsoft employees and contingent staff to access the Internet from any of the more than 130,000 geographically placed desktop computers across the company. Corporate-to-Internet access works as follows: When a client computer makes a request, ISA Server services it from local cache or by passing it to the Internet. When a computer on the Internet responds, ISA Server passes the response back to the client application on the computer that made the request and caches the content on the ISA Server computer (if it is capable of being cached). Figure 9 illustrates a typical deployment of an ISA Server array configured for the purpose of providing corporate-to-Internet access.

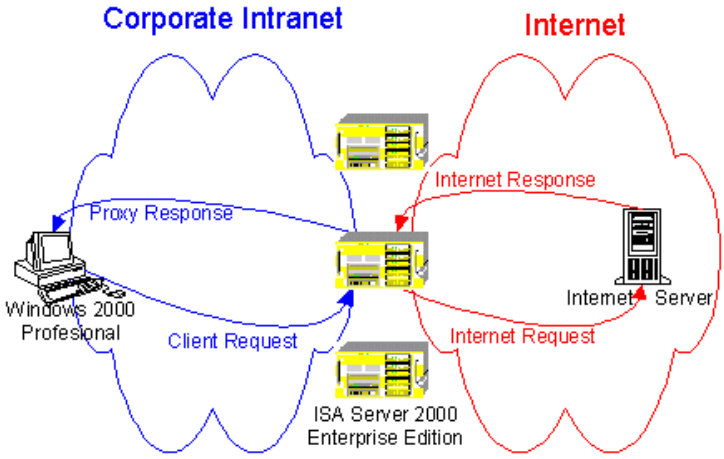

If your browser does not support inline frames, click here to view on a separate page.

Figure 9 A Typical Corporate to Internet Access Configuration

Chaining (or Hierarchical caching)

The deployment team used chaining, which involves linking multiple ISA Server arrays in a prioritized chain that mimics the placement of subsidiary offices, in a subsidiary that had slow (less than 512 Kbps) WAN connectivity to the Internet. The process involved deploying a second ISA Server computer at the subsidiary and configuring it to obtain URL content from an ISA Server computer located across the WAN. The reason for this approach is that LAN speeds within the subsidiary supported 100-Mbps network access within the region.

Team members found that the chained implementation was best suited in situations where the network between the client and Internet access point was constrained or covered a great distance. That's because when a client makes a request, the chained array responds by servicing the request from local cache or passes it to the upstream array where it is either served from its cache or fetched from the Internet (see Figure 11). The chained array stores the content fetched from the upstream array in local cache. Figure 10 illustrates the deployment of an ISA Server array configured to obtain content from another array.

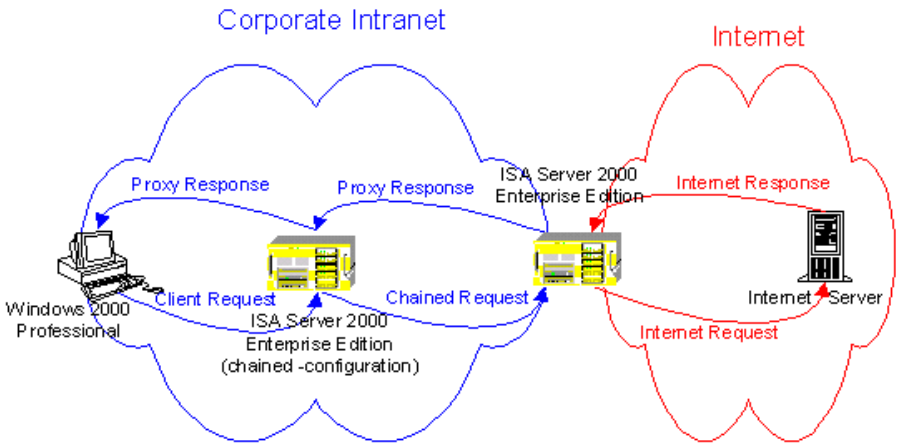

If your browser does not support inline frames, click here to view on a separate page.

Figure 10 Chained deployment of ISA Server

Secure NAT

To support IP-based Macintosh and UNIX clients, the deployment team needed to implement a SecureNAT deployment. This required that they replace a Gauntlet 5.0 legacy firewall (deployed by a third-party vendor that later was acquired by Microsoft) with a corporate standard so the firewall would receive better internal support.

The team's legacy firewall deployment consisted of the following:

- Hardware: Sun Ultra Enterprise 2
- Operating system: SunOS 5.6 (Solaris* 2.6 kernel patch level 105181-16)
- Software: Gauntlet 5.0
- Processors: 2 x SUNW, UltraSPARC-II @ 296 MHz
- RAM: 1GB
- Disk drive: 2 x 4.2G Fast-Wide SCSI disks
- Network: 2 x SUNW, hme 10/100Mbps Ethernet interface

To better understand the legacy environment and formulate a configuration to replace it, team members reviewed inbound and outbound protocols used by the legacy firewall, including the configuration of inbound SMTP and SSH. They also reviewed DNS configuration and reporting capabilities of the UNIX host.

The review was a vital step that helped the deployment team prepare to migrate off of the legacy firewall. The review allowed the team to identify:

- The down-level client requirements (including allowed ports)
- The origin of client requests (including IP addresses)
- What clients were attempting to do through the legacy firewall (including business requirements)
- Expected outcome on the Internet-side of the firewall

Next, the deployment team took the results of their analysis and conducted a proof of concept. The proof of concept consisted of deploying ISA Server in a lab environment so that the team could systematically configure ISA Server to handle network traffic previously handled by the legacy firewall. The lab environment allowed the team to securely develop its deployment plans in a very controlled manner. At the same time, the team was able to document how ISA Server would later need to be configured in production.

Based on its analysis, the team replaced the legacy environment with a Compaq ProLiant 5500R with the following hardware components:

- Quad XEON PIII-550 processor
- 512-KB Cache
- 512-MB RAM
- 7–9.1 GB LVD hard drives

To replace the legacy firewall with ISA Server, team members configured the hardware with a 9.1-GB RAID 1 array and a 45.5-GB RAID 5 array. They installed Windows 2000 Advanced Server on the RAID 1 array and configured it with the addition of SNMP, Terminal Services (in administration mode), and DNS. They also disabled the computer browser, DFS, distributed link-tracking service, and Distributed Transaction Coordinator because these tools were not needed by the configuration. Finally, they configured ISA Server to cache, log, and save reports to the RAID 5 array.

This deployment scenario required ISA Server to use a caching DNS server, which was configured to listen only on the internal corporate IP address. ISA Server was further configured as follows:

Table 9 Configuring ISA Server for replacement of the legacy firewall

ISA Server	Configuration
Mode	Integrated
Cache	20 GB on each server
IP filters	DHCP disabled. All others left in default state
Intrusion detection	Enabled

The team enabled intrusion detection and configured outbound protocols and ports to reflect those of the legacy firewall. Finally, by configuring each Network Interface Card with IP addresses previously used by the legacy firewall, the team transitioned ISA Server into production.

As a result, the same team that manages other ISA Server deployments within Microsoft now manages this deployment without the dependence on a small group of administrators having highly specialized legacy skills.

Lessons Learned

In the early stages of planning, the ITG team charged with executing the beta deployment of ISA Server Enterprise Edition faced a number of obstacles that had to be addressed. ITG is sharing these "lessons learned" in hopes that readers might be able to apply them to their own environments.

Careful planning pays off. Careful planning takes time, but proved to be an essential component of ITG's deployment strategy. Through careful planning ITG was able to thoroughly review business requirements and identify IT employees who had the needed skills to satisfy those requirements while deploying ISA Server. ITG's careful planning also helped team members to maintain the security of the internal Microsoft computing environment because through that planning they gained a better understanding of how to properly position, configure, and deploy ISA Server. Today, business requirements have been met, employees have secure and fast access to information, and Microsoft is in a better position than ever to do business using the Internet.

- **Performance depends on many factors.** Web-client response time depends on such factors as hardware performance on the front and back ends, configuration of the network operating system and ISA Server, network performance, and the back-end server's proximity to the Web client. Proper server placement took several years to refine at Microsoft due to concerns over security, manageability, and the geographical availability of high-speed Internet connectivity. Team members deploying ISA Server had to address all these factors before they could meet the business requirement of providing employees with fast access to Web-based information.
- **Modifying the Active Directory schema must be done with care.** Objects registered in Active Directory are replicated across the Active Directory infrastructure. Objects can be deregistered, but the object container cannot be removed, and the deregistered objects will continue to be replicated throughout the Active Directory infrastructure. Therefore, modifications to the schema should be done correctly the first time they are attempted so as to avoid the replication of deregistered objects. In turn, special care should be taken to avoid creating a need to deregister Active Directory objects in the first place.

  Existing computer hardware was well-suited for ISA Server. ITG used the computer hardware that was previously used by Proxy Server 2.0 because it was equally well-suited for ISA Server. Even though ISA Server provides greater security, better management, and comprehensive reporting capabilities, ISA Server performed so well on the legacy hardware that no additional hardware was required.
- **Limit the number of administrators permitted to change Enterprise Policy.** Unified management of the environment though the application of consistent policy is one of the most powerful management capabilities of ISA Server. A single, quick-and-easy modification to Enterprise Policy may result in a configuration change on dozens of ISA Server computers. Enterprise Policy provides administrators with the ability to lock down or open up the environment rapidly. Because Enterprise Policy is so powerful, ITG limited the number of administrators with permission to modify it. ITG also established rules for requesting protocol modification and for documenting needed changes.

Benefits

The primary benefits of the ISA Server beta deployment at Microsoft involve security, configuration, and troubleshooting.

- **Higher level of security.** Securing intellectual assets is an extremely important job, and information illustrating how the environment is being used is vital to doing that job. The comprehensive reporting capabilities provided by ISA Server are helping ITG to obtain this information much faster than in the past, giving ITG a clearer sense of how the environment is used and an ability to respond more quickly to IT issues that affect the company.
- **Easier configuration.** The ISA Server Enterprise Policy has enabled ITG to deploy a consistent set of protocols across Microsoft with far less administrative overhead than had been required by the legacy environment. Enterprise Policy has made it easy to allow or deny the same set of protocols on every ISA Server computer within the company.

  In contrast, the legacy environment at Microsoft had required manual configuration of every array. This meant that deploying applications such as MSN Messenger required ITG to manually allow the protocol used by the application on every array within the company before employees in all areas could use it. Now, using ISA Server, administrators can apply a single change to Enterprise Policy to allow or deny a given protocol across many arrays.
- **Smarter approach to troubleshooting.** The legacy environment at Microsoft was harder to troubleshoot because ITG often had to manually determine how one array configuration differed from another. Now, using ISA Server and Enterprise Policy, Active Directory multimaster replication virtually ensures that all arrays have a consistently applied configuration, automatically. Unified management of the environment through Enterprise Policy has resulted in a consistent configuration that has helped to reduce the number of required steps needed to diagnose problems.

Conclusion

As more and more businesses move to an Internet model of doing business, the demand for security and performance becomes more crucial than ever. Web-ready enterprises providing goods and services in the evolving global marketplace are seeking better technology to provide those goods and services faster and with greater security than ever before.

The successful deployment of ISA Server at Microsoft (along with other .NET Enterprise servers) has positioned Microsoft to better secure and manage its computing information environment. The secure environment, faster performance and comprehensive reporting capabilities in ISA Server 2000 Enterprise Edition, combined with its ease of deployment at Microsoft, has better positioned ITG to respond more quickly and effectively to security-related issues and other technology requirements throughout the company.

For More Information

ITG will continue to share other deployment stories with customers in hopes that some of what has been learned might better prepare readers from other large organizations to become increasingly Web-ready themselves. More information on this topic and related topics is available from the following sources.

For the latest information on Microsoft Windows 2000 Advanced Server and Windows 2000 Professional go to http://www.microsoft.com/windows/

For the latest information on ISA Server 2000 Enterprise Edition go to http://www.microsoft.com/isaserver/

For the latest information on other .NET Enterprise servers go to http://www.microsoft.com/net/default.asp

For other information that illustrates additional solutions deployed at Microsoft go to http://www.microsoft.com/solutions/HowMicrosoftWorks/default.htm

For support information and self-help tools for Microsoft products on the Microsoft Knowledge Base go to http://search.support.microsoft.com/kb/c.asp?ln=en-us

Here are some of the Knowledge Base articles that assisted ITG in planning the deployment of ISA Server 2000 at Microsoft:

KB Article Q179442 – "How to Configure a Firewall for Windows NT and Trusts"

KB Article Q154596 – "Configuring RPC Dynamic Port Allocation to Work with Firewall"

KB Article Q254949 – "Client-to-Domain Controller and Domain Controller-to-Domain Controller IPSec Support"

To view additional IT Showcase material, go to http://www.microsoft.com/technet/itshowcase/default.htm

For questions, comments, or suggestions related to this document, or to obtain additional information about Microsoft IT Showcase, please send e-mail to showcase@microsoft.com.
