



System i
Security
System i and Internet security

Version 5 Release 4





System i

Security

System i and Internet security

Version 5 Release 4

Note

Before using this information and the product it supports, read the information in "Notices," on page 29.

Seventh Edition (February 2006)

This edition applies to version 5, release 4, modification 0 of IBM i5/OS (product number 5722-SS1) and to all subsequent releases and modifications until otherwise indicated in new editions. This version does not run on all reduced instruction set computer (RISC) models nor does it run on CISC models.

© Copyright International Business Machines Corporation 1999, 2006. All rights reserved.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

System i and Internet security	1
Printable PDFs.	1
System i and Internet security considerations	2
Planning Internet security	3
The layered defense approach to security.	4
Security policy and objectives.	6
Scenario: JKL Toy Company e-business plans	8
Security levels for basic Internet readiness	10
Network security options.	11
Firewalls	11
i5/OS packet rules	13
Intrusion detection	15
Choosing i5/OS network security options	15
Application security options.	16
Web serving security	17
Java Internet security	17

E-mail security	19
FTP security	21
Transmission security options	23
Using digital certificates for SSL	24
Secure Sockets Layer for secure Telnet access	25
Secure Sockets Layer for secure iSeries Access for Windows	25
Virtual private network for secure private communications.	26

Appendix. Notices	29
Programming Interface Information	31
Trademarks	31
Terms and conditions	31

System i and Internet security

Accessing the Internet from your local area network (LAN) requires you to reassess your security requirements.

The integrated software solutions and security architecture of IBM® System i™ product allow you to build a strong defense against potential Internet security pitfalls and intruders. Using these security offerings ensures that your customers, employees, and business partners can obtain the information they need in a secure environment.

This topic collection explains those well-known security threats and how these risks relate to your Internet and e-business goals. This topic collection also discusses how to assess the risks versus the benefits of using the various security options that the system provides for dealing with these risks. You can determine how you can use this information to develop a network security plan that fits your business needs.

Printable PDFs

Use this to view and print a PDF of this information.

To view or download the PDF version of this document, select System i and Internet security (about 456 KB).

You can view or download these related topics:


- Intrusion detection (about 285 KB). You can create an intrusion detection policy that audits suspicious intrusion events that come in through the TCP/IP network, such as incorrectly created IP packets. You also can write an application to analyze the auditing data and report to the security administrator if TCP/IP intrusions are likely to be underway.
- Enterprise Identity Mapping (EIM) (about 1954 KB). Enterprise Identity Mapping (EIM) is a mechanism for mapping a person or entity (such as a service) to the appropriate user identities in various user registries throughout the enterprise.
- Single sign-on (about 1203 KB). The single sign-on solution reduces the number of sign-ons that a user must perform, as well as the number of passwords that a user requires to access multiple applications and systems.
- Planning and setting up system security (about 3992 KB). Planning and setting up system security provides information about how to effectively and systematically plan and configure system-level security.

Saving PDF files

To save a PDF on your workstation for viewing or printing:

1. Right-click the PDF in your browser (right-click the link above).
2. Click the option that saves the PDF locally.
3. Navigate to the directory in which you want to save the PDF.
4. Click **Save**.

Downloading Adobe Reader

- 1 You need Adobe Reader installed on your system to view or print these PDFs. You can download a free copy from the Adobe Web site (www.adobe.com/products/acrobat/readstep.html)  .

Related concepts

Intrusion detection

Enterprise Identity Mapping (EIM)

Single sign on

Plan and set up system security

System i and Internet security considerations

Security issues related to the Internet are significant. This topic provides an overview of i5/OS® security strengths and security offerings.

- | When you connect your System i platform to the Internet, typically one of your first questions is, "What should I know about security and the Internet?" This topic can help you to answer this question.
- | What you need to know depends on how you want to use the Internet. Your first venture into the Internet is to provide your internal network users with access to the Web and Internet e-mail. You might also want the ability to transfer sensitive information from one site to another. Eventually, you can plan to use the Internet for e-commerce or to create an extranet between your company and your business partners and suppliers.

Before you get involved with the Internet, you should think through what you want to do and how you want to do it. Making decisions about both Internet usage and Internet security can be complex.

Note: If you are unfamiliar with security and Internet-related terms, you can review common Security terminology as you work through this material.

After you understand how you want to use the Internet for e-business, as well as the security issues and the available security tools, functions, and offerings, you can develop a security policy and your security objectives. A number of factors affect the choices that you make in developing your security policy. When you extend your organization onto the Internet, your security policy is the critical cornerstone for ensuring that your systems and resources are secure.

i5/OS security characteristics

In addition to a number of specific security offerings for protecting your system on the Internet, the i5/OS operating system has the following security characteristics:

- Integrated security, which is extremely difficult to circumvent compared with add-on security software packages that are offered on other systems.
- Object-based architecture, which makes it technically difficult to create and spread a virus. On an i5/OS operating system, a file cannot pretend to be a program, nor can a program change another program. i5/OS integrity features require you to use system-provided interfaces to access objects. You cannot access an object directly by its address in the system. You cannot take an offset and turn it into, or manufacture, a pointer. Pointer manipulation is a popular technique for hackers on other system architectures.
- Flexibility which lets you set up your system security to meet your specific requirements. You can use the Security Planner to help you determine which security recommendations fit your security needs.

i5/OS advanced security offerings

The i5/OS operating system also offers several specific security offerings that you can choose to enhance your system security when you connect to the Internet. Depending on how you use the Internet, you might want to take advantage of one or more of these offerings:

- Virtual private network (VPN) is an extension of an enterprise's private intranet across a public network, such as the Internet. You can use a VPN to create a secure private connection, essentially by

creating a private tunnel over a public network. VPN is an integrated feature of the i5/OS operating system, available from the iSeries™ Navigator interface.

- Packet rules is an integrated feature of the i5/OS operating system, available from the iSeries Navigator interface. You can configure IP packet filter and network address translation (NAT) rules to control the flow of TCP/IP traffic into and out of your system by using this feature.
- With the Secure Sockets Layer (SSL) protocols, you can configure applications to use SSL to establish secure connections between server applications and their clients. SSL was originally developed for secure Web browser and server applications, but other applications can be enabled to use it. Many applications are now enabled for SSL, including the IBM HTTP Server for i5/OS, iSeries Access for Windows®, File Transfer Protocol (FTP), Telnet, and so on.

Related concepts

“Security policy and objectives” on page 6

Your security policy defines what you want to protect and the security objectives are what to expect of users.

“Scenario: JKL Toy Company e-business plans” on page 8

The typical scenario of JKL Toy Company, which has decided to expand its business objectives by using the Internet, might be helpful for you when you want to set your own e-business plans.

“Virtual private network for secure private communications” on page 26

Virtual private network (VPN), an extension of a company’s intranet over the existing framework of either a public or private network, can help you communicate privately and securely within your organization.

Related information

Connecting to the Internet



AS/400 Internet Security: Protecting Your AS/400 from HARM on the Internet



Security Planner

IP filtering and network address translation

Secure Sockets Layer

Planning Internet security

When you develop your Internet use plans, you must plan for your Internet security needs.

You must gather detailed information about your Internet use plans and document your internal network configuration. Based on the information you gathered, you can accurately evaluate your security needs.

For instance, you need to document and describe the following information:

- Your current network configuration.
- Domain Name System (DNS) and e-mail server configuration information.
- Your connection to your Internet Service Provider (ISP).
- The services you want to use from the Internet.
- The services you want to provide to Internet users.

Documenting this type of information helps you determine where your security exposures are and what security measures you need to minimize these exposures.

| For example, you decide that you want to allow your internal users to use Telnet to connect to hosts at a
| special research location. Your internal users need this service to help them develop new products for
| your company; however, you might have some concerns about confidential data flowing unprotected
| across the Internet. If competitors capture and exploit the data, your company might face a financial risk.
| Having identified your usage needs (Telnet) and the associated risks (exposure of confidential

- | information), you can determine what additional security measures you must put into effect to ensure
- | data confidentiality for this usage (Secure Sockets Layer (SSL) enablement).

The layered defense approach to security

Your security policy defines what you want to protect and what you expect of your system users.

Your security policy provides a basis for security planning when you design new applications or expand your current network. It describes user responsibilities, such as protecting confidential information and creating nontrivial passwords.

Note: You need to create and enact a security policy for your organization that minimizes the risks to your internal network. The inherent security features of the i5/OS operating system, when properly configured, provide you with the ability to minimize many risks. When you connect your system to the Internet, however, you need to provide additional security measures to ensure the safety of your internal network.

Many risks are associated with using Internet access to conduct business activities. Whenever you create a security policy, you must balance providing services against controlling access to functions and data. With networking computers, security is more difficult because the communication channel itself is open to attack.

Some Internet services are more vulnerable to certain types of attacks than others. Therefore, it is critical that you understand the risks that are imposed by each service you intend to use or provide. In addition, understanding possible security risks helps you to determine a clear set of security objectives.

The Internet is home to a variety of individuals who pose threat to the security of Internet communications. The following list describes some of the typical security risks you might encounter:

- | • **Passive attacks**

In a passive attack, the perpetrator monitors your network traffic to try to learn secrets. Such attacks can be either network-based (tracing the communications link) or system-based (replacing a system component with a Trojan horse program that captures data insidiously). Passive attacks are the most difficult to detect. Therefore, you need to assume that someone is eavesdropping on everything you send across the Internet.

- | • **Active attacks**

In an active attack, the perpetrator is trying to break through your defenses and get into your network systems. There are several types of active attacks:

- In **system access attempts**, the attacker attempts to exploit security loopholes to gain access and control over a client or server system.
- In **spoofing** attacks, the attacker attempts to break through your defenses by masquerading as a trusted system, or a user persuades you to send secret information to him.
- In **denial of service attacks**, the attacker tries to interfere with or shut down your operations by redirecting traffic or bombarding your system with junk.
- In **cryptographic attacks**, the attacker attempts to guess or steal your passwords, or use specialized tools to try to decrypt encrypted data.

Multiple layers of defense

Because potential Internet security risks can occur at a variety of levels, you need to set up security measures that provide multiple layers of defense against these risks. In general, when you connect to the Internet, you should not wonder if you will experience intrusion attempts or denial of service attacks. Instead, you should assume that you will experience a security problem. Consequently, your best defense

is a thoughtful and proactive offense. Using a layered approach when you plan your Internet security strategy ensures that an attacker who penetrates one layer of defense will be stopped by a subsequent layer.

Your security strategy must include measures that provide protection across the following layers of the traditional network computing model. Generally, you need to plan your security from the most basic (system level security) through the most complex (transaction level security).

System level security

Your system security measures represent your last line of defense against an Internet-based security problem. Consequently, your first step in a total Internet security strategy must be to properly configure basic system security.

Network level security

Network security measures control access to your i5/OS operating system and other network systems. When you connect your network to the Internet, you need to ensure that you have adequate network level security measures in place to protect your internal network resources from unauthorized access and intrusion. A firewall is the most common means for providing network security. Your Internet service provider (ISP) can provide an important element in your network security plan. Your network security scheme needs to outline what security measures your ISP provides, such as filtering rules for the ISP router connection and public Domain Name System (DNS) precautions.

Application level security

Application level security measures control how users can interact with specific applications. In general, you should configure security settings for each application that you use. However, you should pay special attention to setting up security for those applications and services that you will use from or provide to the Internet. These applications and services are vulnerable to misuse by unauthorized users looking for a way to gain access to your network systems. The security measures that you decide to use need to include both server-side and client-side security exposures.

Transmission level security

Transmission level security measures protect data communications within and across networks. When you communicate across an untrusted network like the Internet, you cannot control how your traffic flows from source to destination. Your traffic and the data it carries flows through a number of different systems that you cannot control. Unless you set up security measures, such as configuring your applications to use the Secure Sockets Layer (SSL), your routed data is available for anyone to view and use. Transmission level security measures protect your data as it flows between the other security level boundaries.

When developing your overall Internet security policy, you should develop a security strategy for each layer individually. Additionally, you should describe how each set of strategies will interact with the others to provide a comprehensive security safety net for your business.

Related concepts

“Security levels for basic Internet readiness” on page 10

Before you connect to the Internet, you should decide what security level you need to adopt for protecting your system.

“Network security options” on page 11

To protect your internal resources, choose the appropriate network level security measures.

“Application security options” on page 16

You have some options to manage the security risks for a number of popular Internet applications and services.

“Transmission security options” on page 23

In order to protect your data when it flows across an untrusted network, such as the Internet, you should put the appropriate security measures into effect. These measures include the Secure Sockets Layer (SSL), iSeries Access for Windows, and virtual private network (VPN) connections.

“Security policy and objectives”

Your security policy defines what you want to protect and the security objectives are what to expect of users.

“E-mail security” on page 19

Using e-mail across the Internet or other untrusted network imposes security risks on your system, even though the system is under the protection of a firewall.

Virtual private network (VPN)

“FTP security” on page 21

File Transfer Protocol (FTP) provides the capability of transferring files between a client (a user on another system) and your server. You need to understand the security risks that you might encounter when you use FTP to ensure that your security policy describes how to minimize the risks.

Related reference

Security terminology

Related information



Lotus Domino Reference Library



Lotus Documentation



Lotus Notes and Domino R5.0 Security Infrastructure Revealed Redbook



Lotus Domino for AS/400 Internet Mail and More Redbook

Security policy and objectives

Your security policy defines what you want to protect and the security objectives are what to expect of users.

Your security policy

- | Each Internet service that you use or provide poses risks to your system and the network to which it is
- | connected. A security policy is a set of rules that apply to activities for the computer and communications
- | resources that belong to an organization. These rules include areas such as physical security, personnel
- | security, administrative security, and network security.

Your security policy defines what you want to protect and what you expect of your system users. It provides a basis for security planning when you design new applications or expand your current network. It describes user responsibilities, such as protecting confidential information and creating nontrivial passwords. Your security policy should also describe how you will monitor the effectiveness of your security measures. Such monitoring helps you to determine whether someone might attempt to circumvent your safeguards.

- | To develop your security policy, you must clearly define your security objectives. After you create a
- | security policy, you must take steps to put into effect the rules it contains. These steps include training
- | employees and adding necessary software and hardware to enforce the rules. Also, when you make
- | changes in your computing environment, you should update your security policy. This is to ensure that
- | you discuss any new risks that your changes might impose.

Your security objectives

When you create and carry out a security policy, you must have clear objectives. Security objectives fall into one or more of the following categories:

Resource protection

Your resource protection scheme ensures that only authorized users can access objects on the system. The ability to secure all types of system resources is a System i strength. You should

carefully define the different categories of users that can access your system. Also, you should define what access authorization you want to give these groups of users as part of creating your security policy.

Authentication

The assurance or verification that the resource (human or machine) at the other end of the session really is what it claims to be. Solid authentication defends a system against the security risk of impersonation, in which a sender or receiver uses a false identity to access a system. Traditionally, systems have used passwords and user names for authentication; digital certificates can provide a more secure method of authentication while offering other security benefits as well. When you link your system to a public network like the Internet, user authentication takes on new dimensions. An important difference between the Internet and your intranet is your ability to trust the identity of a user who signs on. Consequently, you should consider seriously the idea of using stronger authentication methods than traditional user name and password logon procedures provide. Authenticated users might have different types of permissions based on their authorization levels.

Authorization

The assurance that the person or computer at the other end of the session has permission to carry out the request. Authorization is the process of determining who or what can access system resources or perform certain activities on a system. Typically, authorization is performed in context of authentication.

Integrity

The assurance that arriving information is the same as what was sent out. Understanding integrity requires you to understand the concepts of data integrity and system integrity.

- **Data integrity:** Data is protected from unauthorized changes or tampering. Data integrity defends against the security risk of manipulation, in which someone intercepts and changes information to which he or she is not authorized. In addition to protecting data that is stored within your network, you might need additional security to ensure data integrity when data enters your system from untrusted sources. When data that enters your system comes from a public network, you need security methods so that you can perform the following tasks:
 - Protect the data from being sniffed and interpreted, typically by encrypting it.
 - Ensure that the transmission has not been altered (data integrity).
 - Prove that the transmission occurred (nonrepudiation). In the future, you might need the electronic equivalent of registered or certified mail.
- **System integrity:** Your system provides consistent and expected results with expected performance. For the i5/OS operating system, system integrity is the most commonly overlooked component of security because it is a fundamental part of i5/OS architecture. i5/OS architecture, for example, makes it extremely difficult for a hacker to imitate or change an operating system program when you use security level 40 or 50.

Nonrepudiation

The proof that a transaction occurred, or that you sent or received a message. The use of digital certificates and public key cryptography to sign transactions, messages, and documents supports nonrepudiation. Both the sender and the receiver agree that the exchange takes place. The digital signature on the data provides the necessary proof.

Confidentiality

The assurance that sensitive information remains private and is not visible to an eavesdropper. Confidentiality is critical to total data security. Encrypting data by using digital certificates and Secure Socket Layer (SSL) or virtual private network (VPN) connection helps ensure confidentiality when transmitting data across untrusted networks. Your security policy should conclude how you will provide confidentiality for information within your network as well as when information leaves your network.

Auditing security activities

Monitoring security-relevant events to provide a log of both successful and unsuccessful (denied)

access. Successful access records tell you who is doing what on your systems. Unsuccessful (denied) access records tell you either that someone is attempting to break your security or that someone is having difficulty accessing your system.

Related concepts

“System i and Internet security considerations” on page 2

Security issues related to the Internet are significant. This topic provides an overview of i5/OS security strengths and security offerings.

“The layered defense approach to security” on page 4

Your security policy defines what you want to protect and what you expect of your system users.

Configuring DCM

Secure Socket Layer (SSL)

“Scenario: JKL Toy Company e-business plans”

The typical scenario of JKL Toy Company, which has decided to expand its business objectives by using the Internet, might be helpful for you when you want to set your own e-business plans.

Scenario: JKL Toy Company e-business plans

The typical scenario of JKL Toy Company, which has decided to expand its business objectives by using the Internet, might be helpful for you when you want to set your own e-business plans.

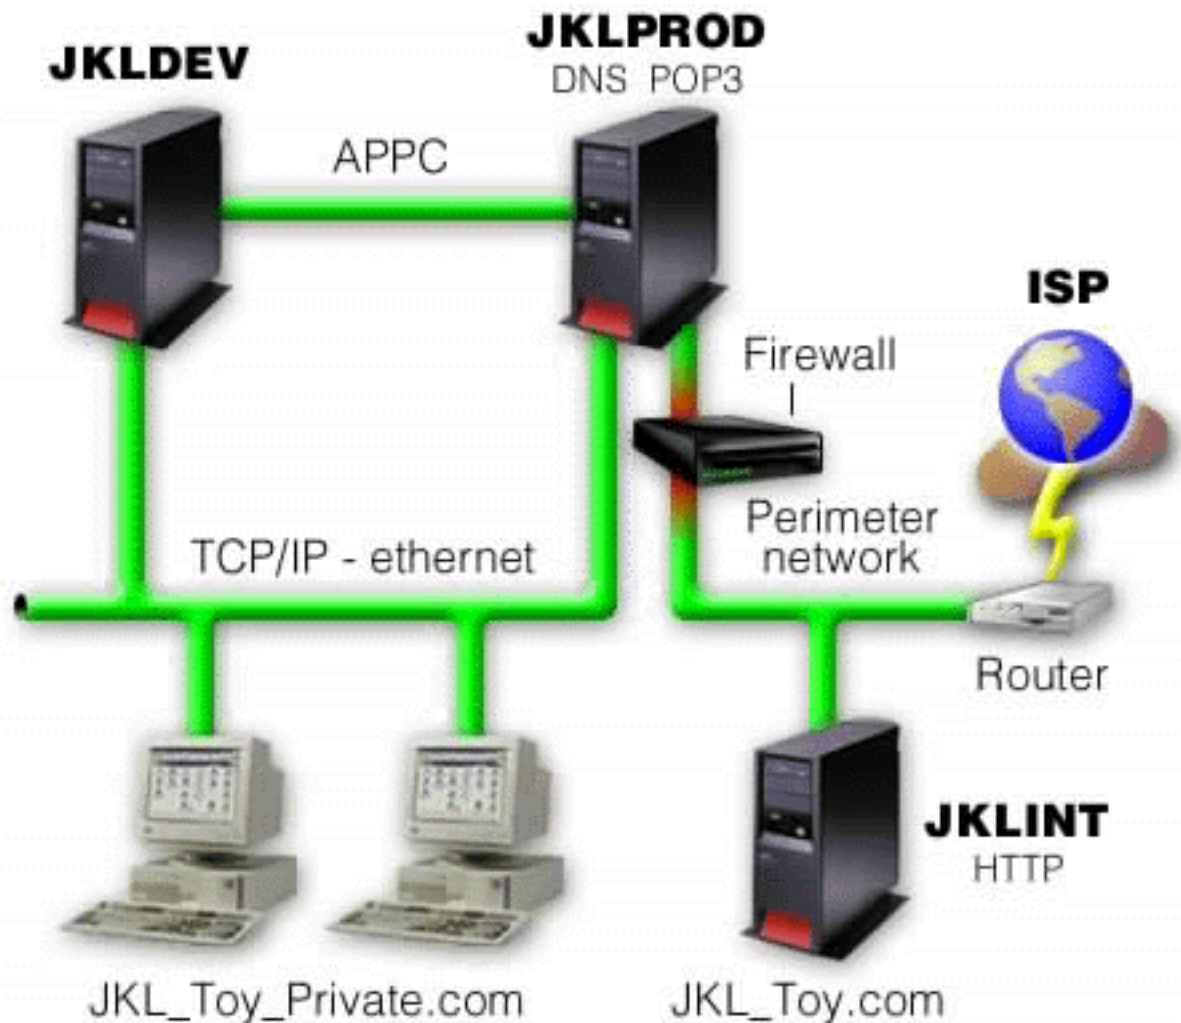
The JKL Toy Company is a small, but rapidly growing manufacturer of toys. The company president is enthusiastic about the growth of the business and how its new i5/OS operating system can ease the burdens of that growth. Sharon Jones, the accounting manager, is responsible for system administration and system security.

The JKL Toy Company has been successfully using its security policy for its internal applications for over a year. The company now has plans to set up an intranet to more efficiently share internal information. The company also has plans to begin using the Internet to further its business goals. Included in these goals are plans for creating a corporate Internet marketing presence, including an online catalog. They also want to use the Internet to transmit sensitive information from remote sites to the corporate office. Additionally, the company wants to allow employees in the design laboratory to have Internet access for research and development purposes. Eventually, the company wants to allow customers to use their Web site for direct online purchasing. Sharon is developing a report about the specific potential security risks for these activities and what security measures the company should use to minimize these risks. Sharon is responsible for updating the company’s security policy and putting into practice the security measures that the company decides to use.

The goals of this increased Internet presence are as follows:

- Promote general corporate image and presence as part of an overall marketing campaign.
- Provide an online product catalog for customers and sales staff.
- Improve customer service.
- Provide employee e-mail and World Wide Web access.

After ensuring that their system has strong basic system security, the JKL Toy Company decides to purchase and use a firewall product to provide network level protection. The firewall will shield their internal network from many potential Internet-related risks. The following figure describes the company’s Internet or network configuration.



As shown in the figure, JKL Toy Company has two primary systems. They use one system for development applications (JKLDEV) and one for production applications (JKLPROD). Both of these systems handle mission-critical data and applications. Consequently, they are not comfortable running their Internet applications on these systems. They have chosen to add a new system (JKLINT) to run these applications.

The company has placed the new system on a perimeter network and is using a firewall between it and the main internal network of the company to ensure better separation between their network and the Internet. This separation decreases the Internet risks to which their internal systems are vulnerable. By designating the new system as an Internet server only, the company also decreases the complexity of managing their network security.

- | The company will not run any mission-critical applications on the new system at this time. During this stage of their e-business plans, the new system provides a static public Web site only. However, the company wants to put security measures into effect to protect the system and the public Web site it runs to prevent service interruptions and other possible attacks. Consequently, the company will protect the system with packet filtering rules and network address translation (NAT) rules, as well as strong basic security measures.
- | As the company develops more advanced public applications (such as an e-commerce Web site or extranet access), they will put more advanced security measures into effect.

Related concepts

“System i and Internet security considerations” on page 2

Security issues related to the Internet are significant. This topic provides an overview of i5/OS security strengths and security offerings.

“Security policy and objectives” on page 6

Your security policy defines what you want to protect and the security objectives are what to expect of users.

“Network security options” on page 11

To protect your internal resources, choose the appropriate network level security measures.

“Transmission security options” on page 23

In order to protect your data when it flows across an untrusted network, such as the Internet, you should put the appropriate security measures into effect. These measures include the Secure Sockets Layer (SSL), iSeries Access for Windows, and virtual private network (VPN) connections.

Security levels for basic Internet readiness

Before you connect to the Internet, you should decide what security level you need to adopt for protecting your system.

Your system security measures represent your last line of defense against an Internet-based security problem. Your first step in a total Internet security strategy must be to properly configure i5/OS basic security settings. Do the following tasks to ensure that your system security meets the minimum requirements:

- Set the security level (QSECURITY system value) to 50. Security level 50 provides the highest level of integrity protection, which is suggested for protecting your system in high risk environments, such as the Internet.

Note: If you are currently running at a security level lower than 50, you might need to update either your operating procedures or your applications. You need to review the iSeries Security Reference before changing to a higher security level.

- Set your security-relevant system values to be at least as restrictive as the recommended settings. You can use the iSeries Navigator Security Wizard to configure the recommended security settings.
- Ensure that no user profiles, including IBM-supplied user profiles, have default passwords. Use the Analyze Default Passwords (ANZDFTPWD) command to check whether you have default passwords.
- Use object authority to protect your important system resources. Take a restrictive approach on your system. That is, by default, restricting everyone (PUBLIC *EXCLUDE) from system resources such as libraries and directories. Allow only a few users to access these restricted resources. Restricting access through menus is not sufficient in an Internet environment.
- You must set up object authority on your system. .

To help you configure these minimum system security requirements, you can use either the Security Planner or the Security Wizard, which is available from the iSeries Navigator interface. The Security Planner provides you with a set of security recommendations based on your answers to a series of questions. You can then use these recommendations to configure the system security settings that you need. Unlike the Security Planner, the wizard uses the recommendations to configure your system security settings for you.

The i5/OS inherent security features, when properly configured and managed, provide you with the ability to minimize many risks. When you connect your system to the Internet, however, you need to provide additional security measures to ensure the safety of your internal network. After ensuring that you have general system security in place, you are ready to configure additional security measures as part of your comprehensive security plan for Internet use.

Related concepts

“The layered defense approach to security” on page 4

Your security policy defines what you want to protect and what you expect of your system users.

Related reference

Security level system value

Related information



iSeries Security Reference PDF

Network security options

To protect your internal resources, choose the appropriate network level security measures.

- | When connecting to an untrusted network, your security policy must describe a comprehensive security scheme, including the security measures that you will put into effect at the network level. Installing a firewall is one of the best means of deploying a comprehensive set of network security measures.

Your Internet Service Provider (ISP) can provide an important element in your network security plan. Your network security scheme should outline what security measures your ISP will provide, such as filtering rules for the ISP router connection and public Domain Name System (DNS) precautions.

Although a firewall certainly represents one of your main lines of defense in your total security plan, it should not be your only line of defense. Because potential Internet security risks can occur at a variety of levels, you need to set up security measures that provide multiple layers of defense against these risks.

Consider using a firewall product as your main line of defense whenever you connect your system or your internal network to the Internet. Although you can no longer purchase the IBM Firewall for the i5/OS product and support for the product is no longer available, there are a number of other products that you can use.

- | Because commercial firewall products provide a full range of network security technologies, the JKL Toy Company chooses one to protect their network. Because the firewall that they choose does not protect their operating system, they add the additional security feature that comes from using the i5/OS packet rules. This allows them to create filter and NAT rules to control traffic for the Internet server.

Related concepts

“The layered defense approach to security” on page 4

Your security policy defines what you want to protect and what you expect of your system users.

“Scenario: JKL Toy Company e-business plans” on page 8

The typical scenario of JKL Toy Company, which has decided to expand its business objectives by using the Internet, might be helpful for you when you want to set your own e-business plans.

Intrusion detection

Related information



Redbook: All You Need to Know When Migrating from IBM Firewall for AS/400

Firewalls

A firewall is a blockade between a secure internal network and an untrusted network such as the Internet.

Most companies use a firewall to connect an internal network safely to the Internet, although you can also use a firewall to secure one internal network from another.

A firewall provides a controlled single point of contact (called a *chokepoint*) between your secure internal network and the untrusted network. The functions of the firewall are as follows:

- Enabling users in your internal network to use authorized resources that are located on the outside network
- Preventing unauthorized users on the outside network from using resources on your internal network

When you use a firewall as your gateway to the Internet (or other network), you reduce the risk to your internal network. Using a firewall also makes administering network security easier because firewall functions carry out many of your security policy directives.

How a firewall works

To understand how a firewall works, imagine that your network is a building to which you want to control access. Your building has a lobby as the only entry point. In this lobby, you have receptionists to welcome visitors, security guards to watch visitors, video cameras to record visitor actions, and badge readers to authenticate visitors who enter the building.

These measures might work well to control access to your building. But, if an unauthorized person succeeds in entering your building, you have no way to protect the building against this intruder's actions. If you monitor the intruder's movements, however, you have a chance to detect any suspicious activity from the intruder.

Firewall components

A firewall is a collection of hardware and software that, when used together, prevent unauthorized access to a portion of a network. A firewall consists of the following components:

• Hardware

Firewall hardware typically consists of a separate computer or device dedicated to running the firewall software functions.

• Software

Firewall software provides a variety of applications. In terms of network security, a firewall provides these security controls through a variety of technologies:

- Internet Protocol (IP) packet filtering
- Network address translation (NAT) services
- SOCKS server
- Proxy servers for a variety of services such as HTTP, Telnet, FTP, and so forth
- Mail relay services
- Split Domain Name System (DNS)
- Logging
- Real-time monitoring

Note: Some firewalls provide virtual private network (VPN) services so that you can set up encrypted sessions between your firewall and other compatible firewalls.

Using firewall technologies

You can use the firewall proxy servers, SOCKS server, or NAT rules to provide internal users with safe access to services on the Internet. The proxy and SOCKS servers break TCP/IP connections at the firewall to hide internal network information from the untrusted network. The servers also provide additional logging capabilities.

You can use NAT to provide Internet users with easy access to a public system behind the firewall. The firewall still protects your network because NAT hides your internal IP addresses.

A firewall also can protect internal information by providing a DNS server for use by the firewall. In effect, you have two DNS servers: one that you use for data about the internal network, and one on the firewall for data about external networks and the firewall itself. This allows you to control outside access to information about your internal systems.

When you define your firewall strategy, you might think it is sufficient to prohibit everything that presents a risk for the organization and allow everything else. However, because computer criminals constantly create new attack methods, you must anticipate ways to prevent these attacks. As in the example of the building, you also need to monitor for signs that, somehow, someone has breached your defenses. Generally, it is much more damaging and costly to recover from a break-in than to prevent one.

In the case of a firewall, your best strategy is to permit only those applications that you have tested and have confidence in. If you follow this strategy, you must exhaustively define the list of services you must run on your firewall. You can characterize each service by the direction of the connection (from inside to outside, or outside to inside). You should also list users who you will authorize to use each service and the machines that can issue a connection for it.

What a firewall can do to protect your network

You install a firewall between your network and your connection point to the Internet (or other untrusted network). Then you can limit the points of entry into your network. A firewall provides a single point of contact (called a chokepoint) between your network and the Internet. Because you have a single point of contact, you have more control over which traffic to allow into and out of your network.

A firewall appears as a single address to the public. The firewall provides access to the untrusted network through proxy or SOCKS servers or network address translation (NAT) while hiding your internal network addresses. Consequently, the firewall maintains the privacy of your internal network. Keeping information about your network private is one way in which the firewall makes an impersonation attack (spoofing) less likely.

- | A firewall allows you to control traffic into and out of your network to minimize the risk of attack to
- | your network. A firewall securely filters all traffic that enters your network so that only specific types of
- | traffic for specific destinations can enter. This minimizes the risk that someone might use Telnet or File
- | Transfer Protocol (FTP) to gain access to your internal systems.

What a firewall cannot do to protect your network

Though a firewall provides a tremendous amount of protection from certain kinds of attack, it is only part of your total security solution. For instance, a firewall cannot necessarily protect data that you send over the Internet through applications, such as Simple Mail Transfer Protocol (SMTP) mail, FTP, and Telnet. Unless you choose to encrypt this data, anyone on the Internet can access it when it travels to its destination.

i5/OS packet rules

You can use the i5/OS packet rules, which is an integrated function of the i5/OS operating system available from the iSeries Navigator interface, to protect your system.

You can use the packet rules to configure two core network security technologies to control the flow of TCP/IP traffic:

- Network address translation (NAT)
- IP packet filtering

Because NAT and IP filtering are integrated parts of your i5/OS operating system, they provide an economical way for you to secure your system. In some cases, these security technologies can provide

everything you need without any additional purchases. These technologies, however, do not create a true, functional firewall. You can use IP packet security alone or in conjunction with a firewall, depending on your security needs and objectives.

Note: The security of your system should take precedence over cost. To ensure that you provide maximum protection for your production system, consider using a firewall.

Network address translation and IP packet filtering

Network address translation (NAT) changes the source or the destination IP addresses of packets that flow through the system. NAT provides a more transparent alternative to the proxy and SOCKS servers of a firewall. NAT can also simplify network configuration by enabling networks with incompatible addressing structures to connect to each other. Consequently, you can use NAT rules so that an i5/OS operating system can function as a gateway between two networks that have conflicting or incompatible addressing schemes. You can also use NAT to hide the real IP addresses of one network by dynamically substituting one or more addresses for the real ones. Because IP packet filtering and NAT complement each other, you will often use them together to enhance network security.

Using NAT can also make it easier to operate a public Web server behind a firewall. Public IP addresses for the Web server translate to private internal IP addresses. This reduces the number of registered IP addresses that are required and minimizes impacts to the existing network. It also provides a mechanism for internal users to access the Internet while hiding the private internal IP addresses.

IP packet filtering provides the ability to selectively block or protect IP traffic based on information in the packet headers. You can use the Internet Setup Wizard in iSeries Navigator to quickly and easily configure basic filtering rules to block unwanted network traffic.

You can use IP packet filtering to do the following tasks:

- Create a set of filter rules to specify which IP packets are allowed or denied the access to your network. When you create filter rules, you apply them to a physical interface (for example, a token ring or Ethernet line). You can apply the rules to multiple physical interfaces, or you can apply different rules to each interface.
- Create rules to either permit or deny specific packets that are based on the following header information:
 - Destination IP address
 - Source IP address Protocol (for example, TCP, UDP, and so forth)
 - Destination port (for example, it is port 80 for HTTP)
 - Source port
 - IP datagram direction (inbound or outbound)
 - Forwarded or Local
- Prevent unwanted or unnecessary traffic from reaching applications on the system. Also, you can prevent traffic from forwarding to other systems. This includes low-level Internet Control Message Protocol (ICMP) packets (for example, PING packets) for which no specific application server is required.
- Specify whether a filter rule creates a log entry with information about packets that matches the rule in a system journal. After the information is written to a system journal, you cannot change the log entry. The log is an ideal tool for auditing network activity.

With the packet filter rules, you can protect your computer systems by rejecting or accepting IP packets according to criteria that you define. NAT rules allow you to hide your internal system information from external users by substituting one public IP address for your internal IP address information. Although IP packet filter and NAT rules are core network security technologies, they do not provide the same level of security that a fully functional firewall product does. You should carefully analyze your security needs

and objectives when deciding between a complete firewall product and the i5/OS packet rules feature.

Related concepts

Network address translation (NAT)

IP packet filtering

Intrusion detection

Intrusion detection involves gathering information about unauthorized access attempts and attacks coming in via the TCP/IP network.

The term *intrusion detection* is used two ways in iSeries documentation. In the first sense, intrusion detection refers to the prevention and detection of security exposures. For example, a hacker might be trying to break into the system using an invalid user ID, or an inexperienced user with too much authority might be altering important objects in system libraries.

In the second sense, intrusion detection refers to the new intrusion detection function that uses policies to monitor suspicious traffic on the system. You can create an intrusion detection policy that audits suspicious intrusion events that come in through the TCP/IP network.

Choosing i5/OS network security options

You need to choose the network security options according to your Internet use plans.

Network security solutions that guard against unauthorized access generally rely on firewall technologies to provide the protection. To protect your system, you can use a full-capability firewall product or put into effect specific network security technologies as part of the i5/OS TCP/IP implementation. This implementation consists of the packet rules feature (which includes IP filtering and NAT) and the HTTP for i5/OS, which is a proxy server licensed program.

Choosing to use either the packet rules feature or a firewall depends on your network environment, access requirements, and your security needs. You must consider using a firewall product as your main line of defense whenever you connect your system or your internal network to the Internet or other untrusted network.

A firewall is preferable in this case because a firewall typically is a dedicated hardware and software device with a limited number of interfaces for external access. When you use the i5/OS TCP/IP technologies for Internet access protection, you are using a general purpose computing platform with a myriad number of interfaces and applications open to external access.

Note: You might want to use both a firewall and integrated i5/OS network security technologies. This helps protect your system from internal attacks (from behind your firewall) and any attacks that might penetrate your firewall because of misconfiguration or other means.

- | The difference is important for a number of reasons. For example, a dedicated firewall product does not
- | provide any other functions or applications beyond those that comprise the firewall itself. Consequently,
- | if an attacker successfully circumvents the firewall and gains access to it, the attacker cannot do much.
- | Whereas, if an attacker circumvents the TCP/IP security functions on your system, the attacker
- | potentially might have access to a variety of useful applications, services, and data. The attacker can then
- | use these to destroy the system itself or to gain access to other systems in your internal network.

As with all the security choices that you make, you must base your decision on the cost versus benefit trade-offs that you are willing to make. You must analyze your business goals and decide what risks you are willing to accept versus the cost you want to pay for security to minimize these risks. The following table provides information about when it is appropriate to use TCP/IP security features versus a fully functional firewall device. You can use this table to determine whether you need to use a firewall, TCP/IP security features, or a combination of both to provide your network and system protection.

Security technology	Best use of i5/OS TCP/IP technology	Best use of a fully functional firewall
IP packet filtering	<ul style="list-style-type: none"> • To provide additional protection for a single i5/OS operating system, such as a public Web server or an intranet system with sensitive data. • To protect a subnetwork of a corporate intranet when the i5/OS operating system is acting as a gateway (casual use router) to the rest of the network. • To control communication with a somewhat trusted partner over a private network or extranet where the i5/OS operating system is acting as a gateway. 	<ul style="list-style-type: none"> • To protect an entire corporate network from the Internet or other untrusted network to which your network is connected. • To protect a large subnetwork with heavy traffic from the remainder of a corporate network.
Network Address Translation (NAT)	<ul style="list-style-type: none"> • To enable the connection of two private networks with incompatible addressing structures. • To hide addresses in a subnetwork from a less trusted network. 	<ul style="list-style-type: none"> • To hide addresses of clients accessing the Internet or other untrusted network. To use as an alternative to Proxy and SOCKS servers. • To make services of a system in a private network available to clients on the Internet.
Proxy server	<ul style="list-style-type: none"> • To proxy at remote locations in a corporate network when a central firewall provides access to the Internet. 	<ul style="list-style-type: none"> • To proxy an entire corporate network when accessing the Internet.

Related reference

IP filtering and network address translation



HTTP Server for i5/OS

Related information



AS/400 Internet Scenarios: A Practical Approach

Application security options

You have some options to manage the security risks for a number of popular Internet applications and services.

Application level security measures control how users can interact with specific applications. In general, you must configure security settings for each application that you use. However, you need to take special care to set up security for those applications and services that you will use from or provide to the Internet. These applications and services are vulnerable to misuse by unauthorized users looking for a way to gain access to your network systems. The security measures that you use need to include both server-side and client-side security exposures.

Though it is important to secure each application that you use, the security measures play a small part in your overall security policy implementation.


Related concepts

“The layered defense approach to security” on page 4

Your security policy defines what you want to protect and what you expect of your system users.

Web serving security

When you provide access for visitors to your Web site, do not expose your viewers to information about how the site is set up and the coding that is used to generate the page. Their visit to your page needs to be easy, fast, and smooth, with all the work being done behind the scenes.

- | As an administrator, you need to ensure that your security practices do not negatively affect your Web site, and that they do implement the security models you have chosen. To achieve this, you need to choose among the security features that are built into the IBM HTTP Server for i5/OS.
- | The chapter about deploying security of IBM HTTP Server (powered by Apache) Redbook  describes how to use authentication, access control, and encryption to implement the security features.
- | Hypertext Transfer Protocol (HTTP) provides you with the capability to display data, but not alter data in a database file. However, sometimes you might need to write some applications that need to update a database file. For example, you might want to create forms that, after users complete them, update an i5/OS database. You can use the Common Gateway Interface (CGI) programs to do this.
- | Proxy server is another security feature that you can use. It receives requests intended for other servers, and then it fulfils, forwards, redirects, or rejects the requests.
- | The HTTP Server provides an access log that you can use to monitor both accesses and attempted accesses through the server.
- | In addition to using CGI programs in your Web pages, you can also use Java™ programming. You need to understand Java security before you add Java to your Web pages.

Related concepts

“Java Internet security”

Java programming is becoming increasingly widespread in today’s computing environments. You should prepare to deal with the security factors that are associated with Java.

Related information

Proxy server types and uses for HTTP Server (powered by Apache)

Security tips for HTTP Server

Common Gateway Interface

Java Internet security

Java programming is becoming increasingly widespread in today’s computing environments. You should prepare to deal with the security factors that are associated with Java.

Although a firewall is a good defense against most general Internet security risks, it does not provide protection for many risks that using Java presents. Your security policy should include details for protecting your system against three areas of concern for Java: applications, applets, and servlets. Also, you should understand how Java and resource security interact in terms of authentication and authorization for Java programs.

Java applications

As a language, Java has some characteristics that protect Java programmers from unintentional errors that can cause integrity problems. (Other languages that are commonly used for PC applications, such as C or C++, do not protect the programmers from unintentional errors as strongly as Java does.) For example, Java uses strong typing, the strict enforcement of type rules with no exceptions, to protect the programmer from using objects in unintended ways. Java does not allow pointer manipulation, which protects the programmer from accidentally going outside the memory boundaries of the program. From an application development perspective, you can view Java as you do other high-level languages. You

need to apply the same security rules for application design that you apply with other languages on your system.

Java applets

Java applets are small Java programs that you can include in HTML pages that run on the client but have the potential to access your i5/OS operating system. An Open Database Connectivity (ODBC) program or an advanced program-to-program communications (APPC) program that operates on a PC in your network can also potentially access your operating system when, for example, your system is being used to serve applications or is being used as a Web server. In general, Java applets can establish a session only with the i5/OS operating system from which the applet originated. Therefore, a Java applet can access your i5/OS operating system from a connected PC only when the applet comes from that i5/OS operating system.

| An applet can attempt to connect to any TCP/IP port on a system. It does not need to talk to a software
| server that is written in Java. But, for systems that are written with the IBM Toolbox for Java, the applet
| must provide a user ID and password when it establishes connections back to the system. In this
| material, the systems described are all i5/OS operating systems. (A Java application server does not need
| to use the IBM Toolbox for Java.) Typically, the IBM Toolbox for Java class prompts the user for a user ID
| and password for the first connection.

The applet can perform functions on the i5/OS operating system only if the user profile has authorization to those functions. Therefore, a good resource security scheme is essential when you begin to use Java applets to provide new application functions. When the system processes the requests from applets, it does not use the limited capability value that is specified in the user profile.

The applet viewer allows you to test an applet on the i5/OS operating system; however, it is not subject to browser security restrictions. Therefore, you need to use the applet viewer to test only your own applets, never to run applets from outside sources. Java applets often write to the PC drive of the user, which might provide the applet with the opportunity to perform a destructive action. However, you can use a digital certificate to sign a Java applet to establish its authenticity. The signed applet can write to the PC's local drives, even though the default setting for the browser prevents it. The signed applet can also write to mapped drives on your system because they appear to the PC to be local drives.

For Java applets that originate from your system, you might need to use signed applets. However, you need to instruct your users, in general, not to accept signed applets from unknown sources.

Beginning with V4R4, you can use the IBM Toolbox for Java to set up a Secure Sockets Layer (SSL) environment. You can also use the IBM Developer Toolkit for Java to make a Java application secure with SSL. Using SSL with your Java applications ensures encryption of the data, including the user IDs and passwords that pass between the client and server. You can use Digital Certificate Manager (DCM) to configure registered Java programs to use SSL.

Java servlets

Servlets are server-side components that are written in Java, which dynamically extend the function of a Web server without changing Web server code. The IBM WebSphere® Application Server that is included in IBM Web Enablement for i5/OS provides support for using servlets on the i5/OS operating systems.

You must use resource security on servlet objects that the system uses. However, applying resource security to a servlet does not sufficiently secure it. After a Web server loads a servlet, resource security does not prevent others from running it too. Consequently, you need to use resource security in addition to using HTTP server security controls and directives. For example, do not allow servlets to run under the profile of the Web server only. You also need to use the security features provided by your servlet development tools, such as those found in the WebSphere Application Server for i5/OS.

Review these resources to know more about general security measures for Java:

- IBM Developer Kit for Java: Java security.
- IBM Toolbox for Java: Security classes.
- Security considerations for Internet browsers.

Java authentication and authorization to resources

IBM Toolbox for Java contains security classes to provide verification of the identity of the user and optionally to assign that identity to the operating system thread for an application or servlet that is running on an i5/OS operating system. Subsequent checks for resource security occur under the assigned identity.

The IBM Developer Kit for Java provides support for the Java Authentication and Authorization Service (JAAS), which is a standard extension to the Java 2 Software Development Kit (J2SDK), Standard Edition. Currently, J2SDK provides access controls that are based on where the code originated and who signed the code (code source-based access controls).

Securing your Java applications with SSL

You can use Secure Sockets Layer (SSL) to secure communications for i5/OS applications that you develop with IBM Developer Kit for Java. Client applications that use IBM Toolbox for Java can also take advantage of SSL. The process for enabling SSL for your own Java applications is somewhat different from the process for enabling SSL for the other applications.

Related concepts

“Web serving security” on page 17

When you provide access for visitors to your Web site, do not expose your viewers to information about how the site is set up and the coding that is used to generate the page. Their visit to your page needs to be easy, fast, and smooth, with all the work being done behind the scenes.

Configuring DCM

Authentication services

Related tasks

Using SSL (JSSE, version 1.0.8)

Related information

Java Authentication and Authorization Service

Secure Sockets Layer (SSL)

E-mail security

Using e-mail across the Internet or other untrusted network imposes security risks on your system, even though the system is under the protection of a firewall.

You must understand these risks to ensure that your security policy describes how you will minimize these risks.

E-mail is like other forms of communication. It is important to use discretion before sending any confidential information through e-mail. Because your e-mail travels through many systems before you receive it, it is possible for someone to intercept and read your e-mail. Consequently, you might want to use security measures to protect the confidentiality of your e-mail.

Common e-mail security risks

These are some risks associated with using e-mail:

- **Flooding** (a type of denial of service attack) occurs when a system becomes overloaded with multiple e-mail messages. It is relatively easy for an attacker to create a simple program that sends millions of e-mail messages (including empty messages) to a single e-mail server to attempt to flood the server. Without the correct security, the target server can experience a denial of service because the server's storage disk fills with useless messages. The system can also stop responding because all system resources become involved in processing the mail from the attack.
- **Spamming** (junk e-mail) is another type of attack common to e-mail. With increasing numbers of businesses providing e-commerce over the Internet, there has been an explosion of unwanted or unrequested for business related e-mail. This is the junk mail, that is being sent to a wide distribution list of e-mail users, filling the e-mail box of each user.
- **Confidentiality** is a risk associated with sending e-mail to another person through the Internet. This e-mail passes through many systems before it reaches your intended recipient. If you have not encrypted your message, a hacker can intercept and read your e-mail at any point along the delivery route.

E-mail security options

To guard against flooding and spamming risks, you must configure your e-mail server appropriately. Most server applications provide methods for dealing with these types of attacks. Also, you can work with your Internet Service Provider (ISP) to ensure that the ISP provides some additional protection from these attacks.

What additional security measures you need depend on the level of confidentiality that you need, as well as what security features your e-mail applications provide. For example, is keeping the contents of the e-mail message confidential sufficient? Or do you want to keep all information associated with the e-mail, such as the originating and target IP addresses, confidential?

Some applications have integrated security features that might provide the protection you need. Lotus Notes® Domino®, for instance, provides several integrated security features including encryption capability for an entire document or for individual fields in a document.

In order to encrypt mail, Lotus Notes Domino creates a unique public and private key for each user. You use your private key to encrypt the message so that the message is readable to only those users that have your public key. You must send your public key to the intended receivers of your note so that they can use it to decipher your encrypted note. If someone sends you encrypted mail, Lotus Notes Domino uses the public key of the sender to decipher the note for you.

You can find information about using these Notes encryption features in the online help files for the program.

When you want to provide more confidentiality for e-mail or other information that flows between branch offices, remote clients, or business partners, you have a couple of options.

If your e-mail server application supports it, you can use Secure Sockets Layer (SSL) to create a secure communications session between the server and e-mail clients. SSL also provides support for optional client-side authentication, when the client application is written to use it. Because the entire session is encrypted, SSL also ensures data integrity while the data is in transit.

Another option available to you is to configure a virtual private network (VPN) connection. You can use your system to configure various VPN connections, including connections between remote clients and your system. When you use a VPN, all traffic that flows between the communicating endpoints is encrypted, ensuring both data confidentiality and data integrity.

Related concepts

Virtual private network (VPN)

“FTP security”

File Transfer Protocol (FTP) provides the capability of transferring files between a client (a user on another system) and your server. You need to understand the security risks that you might encounter when you use FTP to ensure that your security policy describes how to minimize the risks.

“The layered defense approach to security” on page 4

Your security policy defines what you want to protect and what you expect of your system users.

Related reference

Security terminology

Related information



Lotus Domino Reference Library



Lotus Documentation



Lotus Notes and Domino R5.0 Security Infrastructure Revealed Redbook



Lotus Domino for AS/400 Internet Mail and More Redbook

FTP security

File Transfer Protocol (FTP) provides the capability of transferring files between a client (a user on another system) and your server. You need to understand the security risks that you might encounter when you use FTP to ensure that your security policy describes how to minimize the risks.

You can also use the remote command capability to submit commands to the server. Consequently, FTP is useful for working with remote systems or moving files between systems. However, the use of FTP across the Internet, or across other untrusted networks, exposes you to certain security risks. To understand these risks helps you secure your system.

- Your object authority scheme might not provide enough protection when you allow FTP on your system.

For example, the public authority for your objects might be *USE, but today you are preventing most users from accessing those objects by using menu security. (Menu security prevents users from doing anything that is not one of their menu options.) Because FTP users are not restricted to menus, they can read all objects on your system.

Here are some options for controlling this security risk:

- Put into effect full i5/OS object security on the system (in other words, change the system’s security model from menu security to object security. This is the best and most secure option).
- Write exit programs for FTP to restrict access to files that might be transferred through FTP. These exit programs need to provide security that is at least the equivalent as the security that the menu program provides. You might want to make the FTP access controls even more restrictive. This option only covers FTP, not other interfaces such as ODBC, DDM, or DRDA®.

Note: *USE authority to a file allows the user to download the file. *CHANGE authority to a file allows the user to upload the file.

- A hacker can mount a denial of service attack with your FTP server to disable user profiles on the system. This is done by repeatedly attempting to log on with an incorrect password for a user profile until the user profile is disabled. This type of attack disables the profile if it reaches the maximum sign on count of three.

What you can do to avoid this risk involves analyzing the trade-offs that you are willing to make to increase security to minimize the attack versus providing users with ease of access. The FTP server normally enforces the QMAXSIGN system value to prevent a hacker from having unlimited attempts to guess a password and therefore mount password attacks. Here are some options that you need to consider using:

- Use an FTP server logon exit program to reject logon requests by any system user profiles and those user profiles that you designate not be allowed FTP access. (When using such an exit program, logon attempts rejected by the server logon exit point for the user profiles that you block do not get counted against the profile's QMAXSIGN count.)
- Use an FTP server logon exit program to limit the client machines from which a given user profile is allowed to access the FTP server. For example, if a person from Accounting is allowed FTP access, only allow that user profile FTP server access from computers that have IP addresses in the Accounting department.
- Use an FTP server logon exit program to log the user name and IP address of all FTP logon attempts. Review these logs regularly, and whenever a profile is disabled by maximum password attempts, use the IP address information to identify the perpetrator and take appropriate measures.
- Use the intrusion detection system to detect denial of service attacks on the system.

Additionally, you can use FTP server exit points to provide an anonymous FTP function for guest users. Setting up a secure, anonymous FTP server requires exit programs for both the FTP server logon and FTP server request validation exit points.

You can use the Secure Sockets Layer (SSL) to provide secure communications sessions for your FTP server. Using SSL ensures that all FTP transmissions are encrypted to maintain confidentiality for all data that passes between the FTP server and the client, including user names and passwords. The FTP server supports the use of digital certificates for client authentication also.

In addition to these FTP options, you might want to consider using anonymous FTP to provide a convenient way for users to access non-confidential material easily. Anonymous FTP enables unprotected access (no password required) to selected information about a remote system. The remote site determines what information is made available for general access. Such information is considered to be publicly accessible and can be read by anyone. Before configuring anonymous FTP, weigh the security risks and consider securing your FTP server with exit programs.

Related concepts

"E-mail security" on page 19

Using e-mail across the Internet or other untrusted network imposes security risks on your system, even though the system is under the protection of a firewall.

Virtual private network (VPN)

"The layered defense approach to security" on page 4

Your security policy defines what you want to protect and what you expect of your system users.

Related tasks

Configuring anonymous File Transfer Protocol

Managing access using File Transfer Protocol exit programs

Related reference

Security terminology

Related information



Lotus Domino Reference Library



Lotus Documentation



Lotus Notes and Domino R5.0 Security Infrastructure Revealed Redbook



Lotus Domino for AS/400 Internet Mail and More Redbook

Securing FTP

Using SSL to secure the FTP server

Transmission security options

In order to protect your data when it flows across an untrusted network, such as the Internet, you should put the appropriate security measures into effect. These measures include the Secure Sockets Layer (SSL), iSeries Access for Windows, and virtual private network (VPN) connections.

Remember that the JKL Toy company scenario has two primary systems. They use one for development and the other for production applications. Both of these systems handle mission-critical data and applications. Consequently, they choose to add a new system on a perimeter network to handle their intranet and Internet applications.

Establishing a perimeter network ensures that they have some physical separation between their internal network and the Internet. This separation decreases the Internet risks to which their internal systems are vulnerable. By designating the new system as only an Internet server, the company also decreases the complexity of managing their network security.

- | Because of the pervasive need for security in an Internet environment, IBM is continually developing
- | security offerings to ensure a secure networking environment for conducting e-business on the Internet.
- | In an Internet environment you must ensure that you provide both system-specific and
- | application-specific security. However, moving confidential information through a company intranet or
- | across an Internet connection further increases the need to enact stronger security solutions. To combat
- | these risks, you need to put security measures into effect that protect the transmission of data while it
- | travels over the Internet.

You can minimize the risks associated with moving information across untrusted systems with two specific transmission level security offerings for the i5/OS operating system: SSL secure communications and VPN connections.

The SSL protocol is an industry standard for securing communication between clients and servers. SSL was originally developed for Web browser applications, but an increasing number of other applications are now able to use SSL. For the i5/OS operating system, these include:

- IBM HTTP Server for i5/OS (original and powered by Apache)
- FTP server
- Telnet server
- Distributed relational database architecture (DRDA) and distributed data management
- (DDM) server
- Management Central in iSeries Navigator
- Directory Services Server (LDAP)
- iSeries Access for Windows applications, including iSeries Navigator, and applications that are written to the iSeries Access for Windows set of application programming interfaces (APIs)
- Programs developed with Developer Kit for Java and client applications that use IBM Toolkit for Java
- Programs developed with Secure Sockets Layer (SSL) Application Programmable Interfaces (APIs), which can be used to enable SSL on applications. See the Secure Sockets Layer APIs for more information about how to write programs that use SSL.

Several of these applications also support the use of digital certificates for client authentication. SSL relies on digital certificates to authenticate the communication parties and to create a secure connection.

Virtual Private Network (VPN)

You can use your VPN connections to establish a secure communications channel between two endpoints. Like an SSL connection, the data that travels between the endpoints can be encrypted, thereby providing both data confidentiality and data integrity. VPN connections, however, allow you to limit the traffic flow

to the endpoints that you specify and to restrict the type of traffic that can use the connection. Therefore, VPN connections provide some network level security by helping you to protect your network resources from unauthorized access.

Which method should you use

Both SSL and VPN address the need for secure authentication, data confidentiality, and data integrity. Which of these methods you should use depends on several factors. You need to consider who you are communicating with, what applications you use to communicate with them, how secure you need the communication to be, and what trade-offs in cost and performance you are willing to make to secure this communication.

Also, if you want to use a specific application with SSL, that application must be set up to use SSL. Although many applications cannot take advantage of SSL yet, many others, like Telnet and iSeries Access for Windows, have added SSL capability. VPNs, however, allow you to protect all IP traffic that flows between specific connection endpoints.

For example, you can use HTTP over SSL currently to allow a business partner to communicate with a Web server on your internal network. If the Web server is the only secure application that you need between you and your business partner, then you might not want to switch to a VPN connection. However, if you want to expand your communications, you might want to use a VPN connection instead. You might also meet a situation, in which you need to protect traffic in a portion of your network, but you do not want to individually configure each client and server to use SSL. You can create a gateway-to-gateway VPN connection for that portion of the network. This can secure the traffic, but the connection is transparent to individual servers and clients on either side of the connection.

Related concepts

“The layered defense approach to security” on page 4

Your security policy defines what you want to protect and what you expect of your system users.

“Scenario: JKL Toy Company e-business plans” on page 8

The typical scenario of JKL Toy Company, which has decided to expand its business objectives by using the Internet, might be helpful for you when you want to set your own e-business plans.

Related reference

Secure sockets APIs

Related information

Secure Sockets Layer (SSL)

Using digital certificates for SSL

Digital certificates provide the foundation for using the Secure Sockets Layer (SSL) for secure communications and as a stronger means of authentication.

The i5/OS operating system provides you with the ability to easily create and manage digital certificates for your systems and users with Digital Certificate Manager (DCM), an i5/OS integrated feature.

Additionally, you can configure some applications, such as the IBM HTTP Server for i5/OS, to use digital certificates for a stronger method of client authentication instead of user name and passwords.

What is a digital certificate

A digital certificate is a digital credential that validates the identity of the certificate owner, much as a passport does. A trusted third party, called a certificate authority (CA), issues digital certificates to users and servers. The trust in the CA is the foundation of trust in the certificate as a valid credential.

Each CA has a policy to determine what identifying information the CA requires to issue a certificate. Some Internet CAs might require little information, such as only requiring a distinguished name. This is

the name of the person or system to whom a CA issues a digital certificate address and a digital e-mail address. A private key and a public key are generated for each certificate. The certificate contains the public key, while the browser or a secure file stores the private key. The key pairs associated with the certificate can be used to sign and encrypt data, such as messages and documents, sent between users and servers. Such digital signatures ensure the reliability of an item's origin and protect the integrity of the item.

Although many applications cannot take advantage of SSL yet, many others, like Telnet and iSeries Access for Windows, have added SSL capability.

Related concepts

Configuring DCM

Secure Sockets Layer (SSL)

Related reference

Security terminology

Secure Sockets Layer for secure Telnet access

You can configure your Telnet server to use the Secure Sockets Layer (SSL) to secure Telnet communications sessions.

To configure your Telnet server to use SSL, you must use Digital Certificate Manager (DCM) to configure the certificate for the Telnet server to use. By default the Telnet server handles both secure and non-secure connections. However, you can configure Telnet so that it allows only secure Telnet sessions. Additionally, you can configure the Telnet server to use digital certificates for stronger client authentication.

- | When you choose to use SSL with Telnet, you gain some strong security benefits. For Telnet, besides
- | server authentication, the data is encrypted before any Telnet protocol data flows. After the SSL session is
- | established, all Telnet protocols including user ID and password exchange are encrypted.

The most important factor to consider when using the Telnet server is the sensitivity of the information that you use in a client session. If the information is sensitive or private, then you may find it beneficial to set up your Telnet server using SSL. When you configure a digital certificate for the Telnet application, the Telnet server is able to operate with both SSL and non-SSL clients. If your security policy requires that you always encrypt your Telnet sessions, you can disable all non-SSL Telnet sessions. When there is no need for you to use the SSL Telnet server, you can turn off the SSL port. You can control the use of SSL for Telnet sessions using the Change Telnet Attributes (CHGTELNA) command Allow Secure Socket Layer (ALWSSL) parameter. To ensure no applications can use the SSL or Non-SSL ports as appropriate, you can also restrict this by using the Add TCP/IP Port Restriction (ADDTCPPORT) command.

- | To learn more about Telnet and about security tips for Telnet with and without SSL, the IBM Systems
- | Software Information Center topic on Telnet provides the information that you need to use Telnet on your i5/OS operating system.

Related concepts

Telnet scenario: Securing Telnet with SSL

Planning for DCM

Related information

Telnet

Secure Sockets Layer for secure iSeries Access for Windows

To secure iSeries Access for Windows communications sessions, you can configure your iSeries Access for Windows to use the Secure Sockets Layer (SSL).

Using SSL ensures that all traffic for the iSeries Access for Windows sessions are encrypted. This keeps data from being read while it is in transit between the local and remote hosts.

Related information

Secure Sockets Layer administration

Java security

Security classes

Virtual private network for secure private communications

Virtual private network (VPN), an extension of a company's intranet over the existing framework of either a public or private network, can help you communicate privately and securely within your organization.

| With the rise in the use of VPN and the security they provide, JKL Toy company is exploring options to
| transmit data over the Internet. They have recently acquired another small toy manufacturing company
| that they intend to operate as a subsidiary of themselves. JKL will need to pass information between the
| two companies. Both companies use the i5/OS operating system and a VPN connection that can provide
| the security they need to communicate between the two networks. Creating a VPN is more cost-effective
| than using traditional non-switched lines.

| These are some of the users who can benefit from using VPNs for connectivity:

- Remote and mobile users.
- Home office to the branch office or other off-site locations.
- Business-to-business communications.

| Security risks occur if you do not limit user access to sensitive systems. Without limiting who can access
| a system, you may increase the chances that company information is not kept confidential. You need a
| plan that will allow only those who need to share information about a system to access that system. A
| VPN allows you to control network traffic while providing important security features such as
| authentication and data privacy. Creating multiple VPN connections allows you to control who can access
| which systems for each connection. For example, Accounting and Human Resources may link through
| their own VPN.

When you allow users to connect to system over the Internet, you may be sending sensitive corporate data across public networks, which might expose this data to attack. One option for protecting transmitted data is to use encryption and authentication methods for ensuring privacy and security from outsiders. VPN connections provide a solution for a specific security need: securing communications between systems. VPN connections provide protection for data that flows between the two endpoints of the connection. Additionally, you can use packet rules security to define what IP packets are allowed across the VPN.

You can use VPN to create secure connections to protect traffic that flows between controlled and trusted endpoints. However, you still must be wary about how much access you provide to your VPN partners. A VPN connection can encrypt data while it travels over public networks. But, depending on how you configure it, data flowing across the internet may not be transported through a VPN connection. In such a case, the data would not be encrypted as it flows across the internal networks that communicate through the connection. Consequently, you should carefully plan how to set up each VPN connection. Ensure that you give your VPN partner access to only those hosts or resources on your internal network that you want them to access.

For instance, you might have a vendor that needs to obtain information about what parts you have in stock. You have this information in a database that you use to update Web pages on your intranet. You want to allow this vendor to access these pages directly through a VPN connection. But you do not want the vendor to be able to access other system resources, such as the database itself. You can configure your VPN connection such that traffic between both endpoints is restricted to port 80. Port 80 is the default port that HTTP traffic uses. Consequently, your vendor can send and receive HTTP requests and responses over that connection only.

Because you can restrict the type of traffic that flows across the VPN connection, the connection provides a measure of network level security. However, VPN does not work in the same manner that a firewall does to regulate traffic into and out of your system. Also, a VPN connection is not the only means available to secure communications between your i5/OS operating system and other systems. Depending on your security needs, you might find that using SSL is a better fit.

Whether a VPN connection provides the security that you need depends on what you want to protect. Also, it depends on the trade-offs that you are willing to make to provide that security. As with any decision that you make about security, you should consider how a VPN connection supports your security policy.

Related concepts

Virtual private networks (VPN)

“System i and Internet security considerations” on page 2

Security issues related to the Internet are significant. This topic provides an overview of i5/OS security strengths and security offerings.

Appendix. Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

- I IBM Corporation

| Software Interoperability Coordinator, Department YBWA
| 3605 Highway 52 N
| Rochester, MN 55901
| U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

| The licensed program described in this information and all licensed material available for it are provided
| by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement,
| IBM License Agreement for Machine Code, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_. All rights reserved.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Programming Interface Information

This System i and Internet security publication documents intended Programming Interfaces that allow the customer to write programs to obtain the services of IBM i5/OS.

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

- | AS/400
- | Domino
- | DRDA
- | i5/OS
- | IBM
- | IBM (logo)
- | iSeries
- | Lotus Notes
- | Notes
- | System i
- | WebSphere

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Other company, product, and service names may be trademarks or service marks of others.

Terms and conditions

Permissions for the use of these publications is granted subject to the following terms and conditions.

Personal Use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these publications, or any portion thereof, without the express consent of IBM.

Commercial Use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER

EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.



Printed in USA