SECURITY-7 WHITE PAPER

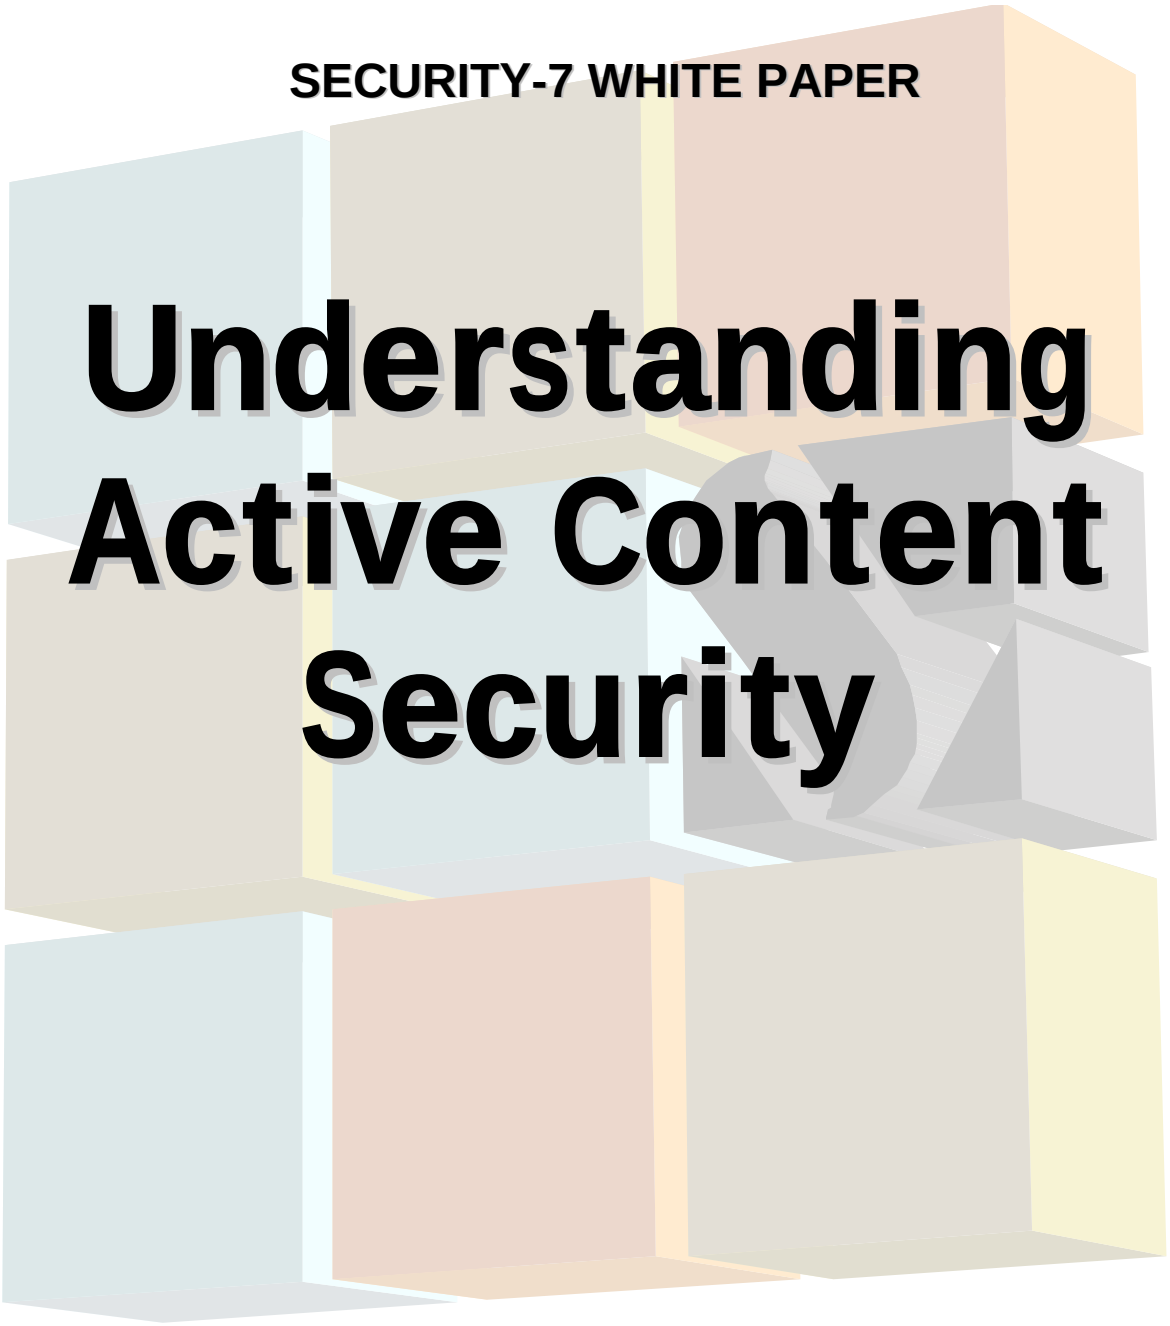

Understanding Active Content Security

Security-7 Software Inc.
10 Tower Office Park, Suite 214 • Woburn • MA 01801-2120
Tel: 781 932-9300 • Fax: 781 932-9395
Email: info@us.security7.com • Web: www.security7.com

FIREWALLS AND ACTIVE CONTENT SECURITY

In order to understand why customers need an Active Content Security solution in addition to their firewall, one must understand the limitations of firewalls.

In order to understand why customers need an Active Content Security solution in addition to their firewall, one must understand two fundamentals:

1. There is a difference between how a firewall controls IP data connections, and the data itself within the connection, and
2. The practical risks associated with Active Content.

Firewalls are the first line of defense typically employed by a company when connecting its internal network to the Internet. They do this by specifying rules about *who* can access *what* network resources using well-known network services (e.g. browsing, e-mail, etc). In short, they control the connections between users and resources by service type.

Additional measures can be layered on top of the firewall to further secure the connection. For example, authentication works in conjunction with the firewall to better identify users rather than relying on IP addresses. Encryption scrambles the data passing through the connection between the user and the firewall, preventing third parties from stealing and deciphering it.

However, none of these systems can look inside the connection and determine if the data itself might violate some security rules of the enterprise. This data, in the form of executable code like Java Applets, can therefore enter the network through the firewall and act with potentially malicious behavior and damaging results.

Typically, many enterprise network administrators will configure the firewall to allow their internal users (all internal client IPs) to access all Internet World Wide Web locations (all external IPs) using their browsers (HTTP protocol).

Once an HTTP connection has been established, Active Content might enter the network unrestricted, running inside Web browsers and e-mail clients and performing hostile activities only possible by it working from inside the network.

ACTIVE CONTENT

Java Applets and ActiveX Controls deliver the level of interactivity that is required by business applications, e-commerce, multimedia and other advanced computer applications.

Active Content is made of static Web content – HTML pages and images – embedded with a variety of active components: Java Applets, ActiveX Controls, plug-ins and scripts.

Per se, HTML pages are harmless, static content that is located on some remote server and brought over to the end-user's computer so it can be viewed with a Web browser, e-mail client or similar applications.

Java Applets, ActiveX Controls and plug-ins are the common forms of active code components that deliver the level of interactivity that is required by business applications, e-commerce, multimedia and other forms of advanced computer applications.

Like other interactive computer software, they run on the end-user's computer and have access to all its resources – both a benefit and a potential security risk.

The Benefits of Active Content

Different forms of Active Content, whether they are Java Applets, ActiveX Controls or plug-ins, are widely popular on the Internet because they help bring Web pages to life. They are becoming even more popular inside corporate Intranets by enabling business applications and services to be built using simple and ubiquitous Web technologies.

In a typical application, a user will use a Java Applet from a corporate Intranet site to query a database, refine the query with filters and sort orders, then export the information into an Excel worksheet. Or the user will prepare a report with Word and then use the browser to file the report in an on-line discussion form or a workflow application.

It is no doubt that businesses are so interested in exploiting these new technologies.

Different Types of Active Content

Java Applets were devised by Sun as a solution for enhancing Web pages. ActiveX Controls are specific to Internet Explorer and Windows.

The term Active Content relates mainly to Java Applets and ActiveX Controls.

Java Applets

Java Applets were devised by Sun as a solution for enhancing Web pages and allowing for application development, and have since been adopted by software houses around the world.

The term “Applets” was coined to describe the small applications that run inside the Web browser. They are automatically downloaded from a server, and run by a dedicated Java Virtual Machine that is an integral part of each Web browser.

The best source for information about Java, its capabilities and how to start using it is JavaSoft itself, at www.javasoft.com.

ActiveX Controls

ActiveX Controls are specific to Internet Explorer and the Windows operating system. Similar to Java Applets, ActiveX Controls are small software components that run inside the Web browser.

Unlike Java Applets, ActiveX Controls are developed in more traditional languages such as C++ and Visual Basic. ActiveX Controls are, in fact, Windows applications adapted to load on demand from the Internet and run inside Internet Explorer from specially designed Web pages.

Since ActiveX Controls are native executables, they are more prone to a variety of security risks and holes. Adoption on the Internet has been slow. ActiveX Controls are mostly popular inside corporate Intranets, where they are easier to control.

Plug-ins

Plug-ins are slowly being out-phased by Java Applets and ActiveX Controls. Still, plug-ins have the same security model as ActiveX: when you download a plug-in, you are trusting it to be harmless. All of the warnings about ActiveX programs apply to plug-ins too.

Exposure to Hostile Code

The whole process of requesting the Active Content and running it occurs almost instantaneously, without the user even being aware of it or having a chance to intervene.

As a rule of thumb, the more business productive Active Content is, the more of a security breach it may become.

The whole process of requesting the Active Content, initializing it, and running it occurs almost instantaneously, without the user even being aware of it or having a chance to intervene.

The previous example detailed two useful business scenarios involving Active Content. Imagine a similar Java Applet that uses the exact same technologies, but instead of being productive it becomes hostile: it steals information from the database, it misplaces information inside Excel worksheets, it modifies Word documents and gains access to sensitive e-mails.

Using JDBC, a Java Applet can access any database server (ActiveX uses a similar technology called ODBC). The database may be accessed to query information and perform transactions on behalf of the user, but a malicious Applet might attempt to access the database in exactly the same manner and perform uncalled-for changes or retrieve sensitive information.

The `java.io` package allows Java Applets to access the local disk storage, for example, for reading documents or storing a user's preferences (ActiveX achieves the same using the Windows API). A hostile Applet would use `java.io` to read information it's not supposed to access.

The `java.net` package allows Java Applets to interact with remote servers. In a productive way, the Applet could check inventory stocks, fill and submit orders, and e-mail invoices. A hostile Applet might follow a variant of the theme, filling an order for a non-paying customer, or stealing credit card numbers right off the server.

A well-devised hostile Applet or ActiveX would perform such actions in the background, appearing to be useful or entertaining to the user, as it completes its malicious act. The user falls under a false sense of security watching the Applet doing something useful, unaware of the actions taken in the background.

As far as the end-user is concerned, both forms of Active Content use the same technologies, access the same resources, require the same permissions – but to different ends.

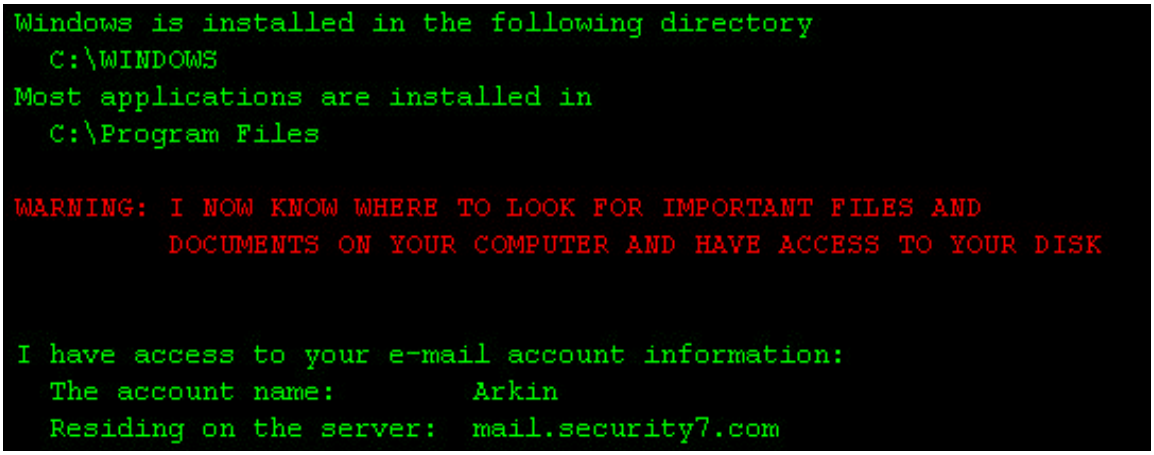

Figure 1: An example of a hostile Java Applet from www.withinreach.co.il/hostiles

Since firewalls are set to leave the HTTP protocol open for browsing the Web, intrusions from the outside might be difficult, but intrusions from the inside, aided by Active Content running inside the corporate network, are all too easy.

An interesting trait of the Web is that, while being the ideal medium for mass distribution, it is also a perfect medium for launching a well-focused attack. Viruses spread by using the most common denominator: disks and Word documents. But with the advent of the Internet, hostile code can be crafted to attack a particular database, application server, or company's IT system.

METHODS OF PROTECTION

Protection at the Browser

Web browsers come with security features. They have one weak spot, though – that being the users themselves.

Web browsers do not come without security features. The major browsers provide a variety of technologies to help protect their users. They have one weak spot, though – that being the users themselves.

In order to be protected, the user must be fully aware of what risks a particular form of Active Content can inflict, and constantly be on the alert to prevent it from doing so.

Figure 2: Browser security is based on understanding the content of this message box and clicking the correct button for each Active Content being used throughout the day.

A user is asked by the browser to decide whether or not to allow a particular form of Active Content to run and the decision is left at the sole discretion of the end-user. Yet, end-users cannot be expected to constantly check every form of Active Content coming over the network, or even to understand the risks involved and how to make a sound judgment.

In fact, this security approach has the effect of improving a user's ability to click on the *Yes* button in response to a browser's security message box.

Protection at the Firewall

Organizations that are aware of the risks involved in Active Content have found firewalls to be their first aid. Most firewalls in the market allow Active Content to be completely shut off.

The major shortcoming of firewalls is the total lack of discretion. They either allow or block all forms of Active Content, either preventing users from utilizing productive forms of Active Content that lie outside the corporate network, or leaving open a hole through which Active Content can penetrate the network.

Protection at the Gateway

The best form of protection would not involve the user at all, but rather rely on the knowledge and expertise of the security manager to balance productivity needs and the security risks involved.

A gateway enterprise solution, such as SafeGate, offers the security manager the ability to control access to Active Content down to the level of a function call or up to the level of corporate departments, relieving end-users of that responsibility.

By working at the gateway, it proactively blocks unwanted Active Content from ever entering the network. Active Content is analyzed based on its source of origin, the level of trust associated with that particular source, and an understanding of risks involved with what this Active Content is attempting to do.

Furthermore, a product such as SafeGate is capable of inspecting content delivered in a variety of ways, whether by surfing the Internet, delivered through a push c hannel, or sent as an e-mail.

DIGITAL SIGNATURES AND CERTIFICATES

Digital signatures were introduced to improve the security of Active Content, but at the same time allow for security restrictions to be lifted.

Digital signatures were introduced to improve the security of Active Content by providing a tamper-proof measure of identification.

Digitally signing an Active Content, associates that content with a one-way hash code that could only be produced by the creator of that Active Content. Any attempt to modify the Active Content, for example, the classical man-in-the-middle attack, would render it invalid.

Attached to the digitally signed object is a certificate – a proof of ownership. Like the signature, sophisticated encryption technology is used to create the certificate so it can be proven beyond doubt that this particular certificate belongs to its designated owner.

Risks Involved in Signed Active Content

Because digital signatures improved security through accountability, some security restrictions were lifted.

Java Applets execute within a confined environment, dubbed the “Java sandbox”, that prevents Applets from accessing the local disk, system facilities or other vulnerable resources.

A digitally signed Java Applet can extend beyond the confinements of the sandbox by the mere fact of having a digital signature. Whether signed by a respectable Fortune 500 corporation or by a hacker, the signed Applet is presumed to be a trusted executable.

A signed Java Applet can gain access to valuable system resources, modify files on the local disk as well as remote machines, access databases and application servers, in fact, every form of attack accomplished by a computer software.

Certificate-Based Security

Certificates are a double-edge sword: a reliable security mechanism on one hand, and a major security risk on the other. Because they allow Java Applets to break out of the limiting sandbox and access vulnerable resources, relying on the discretion of the unsuspecting end-user, they pave the high road for attacks on corporate networks.

At the same time, relying on public-private key digital encryption, they provide the most adequate measure of security when used with the proper tools.

While a hacker can easily create a digitally signed hostile Active Content, the hacker cannot impersonate the certificate, or steal the signature of another body. The security manager can easily identify a known and trusted source, such as a software vendor or business partner, and grant access to sensitive resources, while preventing such access from potentially harmful Active Content.

Some corporations use their own certificates for in-house development, some even use them as a way to distinguish between risky (in development stage) and safe (in production) in-house Active Content.

MANAGING SECURITY WITH TRUST LEVELS

Trust levels are the key to secure and protective control of Active Content. Different levels of trust can be associated with different sources of Active Content.

Trust levels are the key to secure and protective control of Active Content. The fact that a certificate is valid and has not been tampered with does little to prove that the certificate owner had no intention of malice.

Managing Trust Levels

Centrally managed products, such as SafeGate, allows security managers to manage security by defining a variety of security policies, or plans, based on the level of trust associated with the origin of Active Content.

In that respect, they provide better control and manageability than any browser-based security, and fine grain control not available from a firewall.

A source that is fully trusted, such as a major software vendor or Web content/service provider might be allowed to perform a wide range of activities even to the level of replacing operating system and application components.

Wary system administrators might want to attach their own signature to commercial code as way of creating different trust levels for different commercial software products.

If the source originated within the company, one of its business partners, or primary software vendors, it might be enough to grant access to valuable resources used in daily business activities, such as databases and application servers.

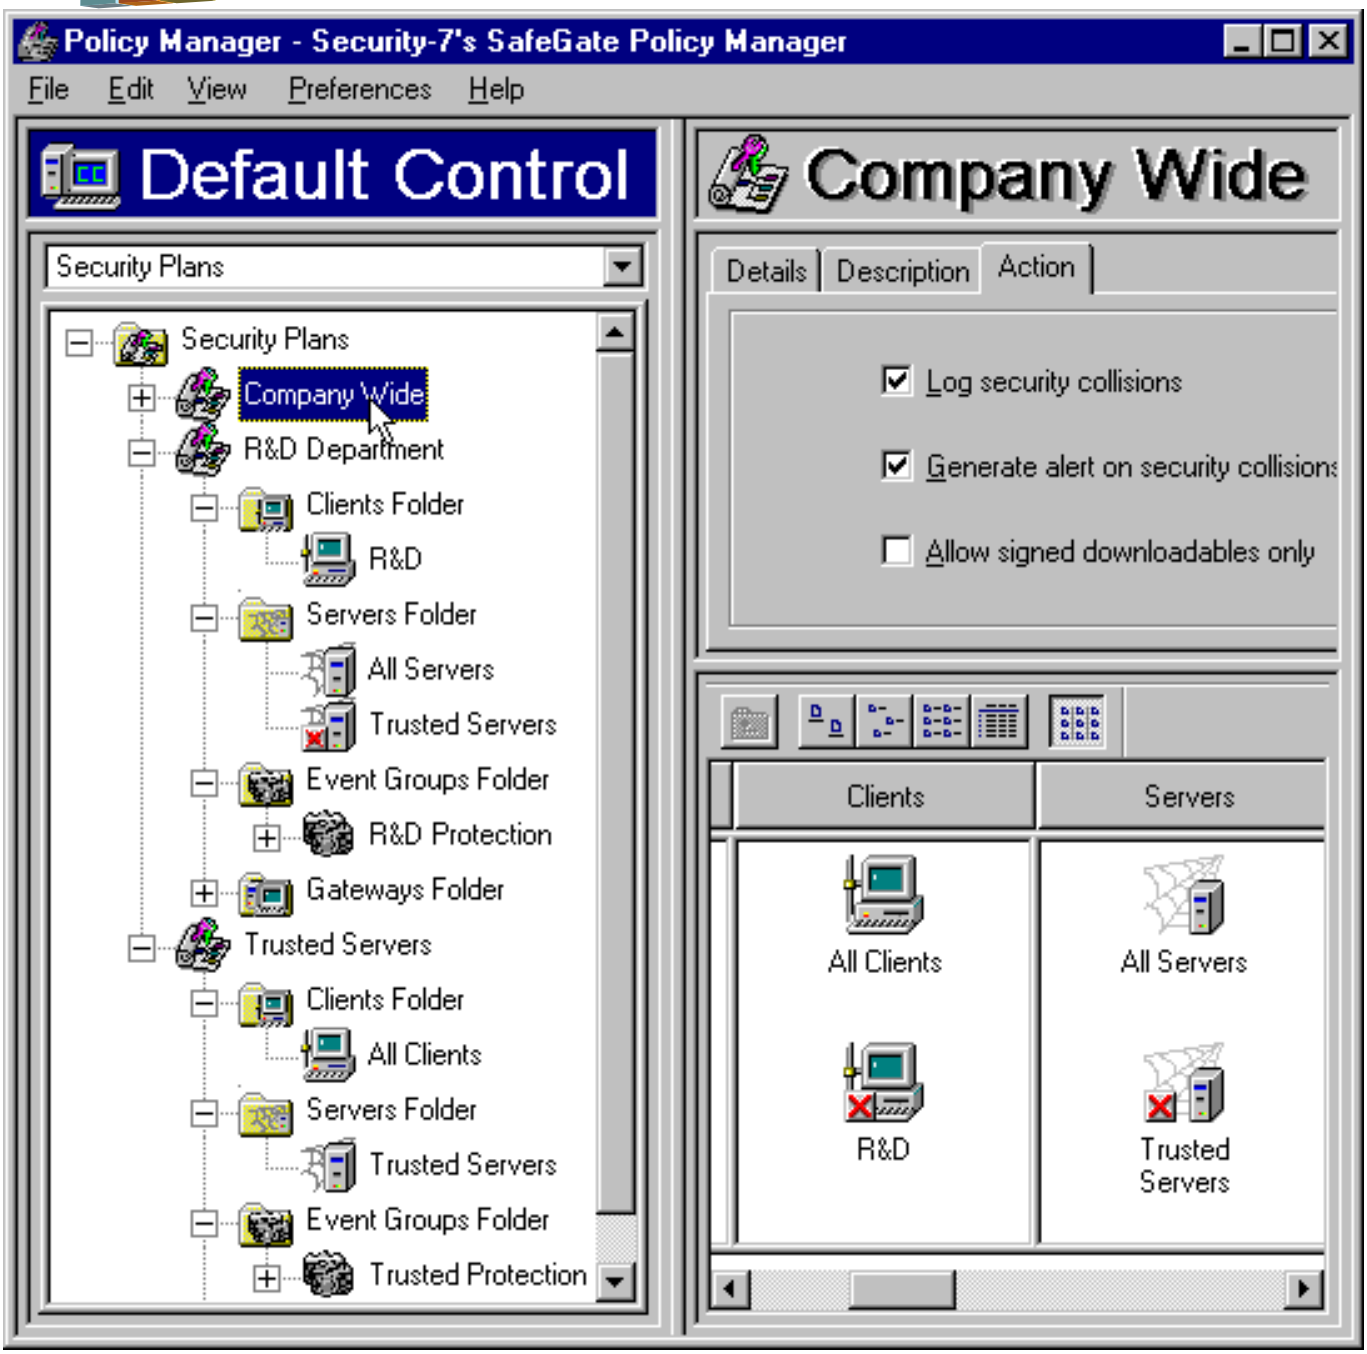

Figure 3: Example of three security plans based on different users and trust-levels. (Screen shot from SafeGate Policy Manager)

Again, the system administrator might create different trust levels inside the organization, for example, making sure that accounting never has access to Active Content from the R&D network, as a way to protect them from probable bugs occurring during the development phases.

If the source is considered trustworthy to some extent, such as a certificate issued to a software vendor by a certificate authority, it may be enough to grant access to minor – and certainly not vulnerable – resources that are necessary to improve business productivity.

Sources that are not trustworthy, or their trust level cannot be decided easily, should be prevented from any form of access to any resource that might put a computer at jeopardy – even if that means they won't be able to run at all. And granting access to such resources should not even be considered until the source of origin can be proven to be trusted.

The flexibility of the security plan means that risk can be managed based on the tradeoffs between the possibility of a hazard and the advantage of business productivity. Active Content Security is no longer an on/off switch or a shot in the dark.

APPENDIX A: STEPS TO PROTECT THE ORGANIZATION

Four Steps to Follow

Follow these four steps to keep your corporate network covered at all times.

Step 1

Install a firewall to protect the network against unwanted attacks. Set the firewall not to allow any Active Content (mostly Java Applets and ActiveX Controls) to enter the network until Active Content Security is installed and is fully functional.

Step 2

Upgrade all Web browsers to the latest version. Older versions of Web browsers contain some bugs that can lead to security holes or unexpected behavior.

Step 3

Install an Active Content Security product, such as SafeGate, on all gateways connecting the network to the outside world, whether the Internet or Intranets of business partners. Set-up an inclusive set of security plans, blocking out known hostile Active Content and Web sites.

We recommend not allowing any Active Content to enter machines that act as database, application or e-mail servers.

We recommend allowing only signed ActiveX Controls, and removing support for weak digital signatures (e.g. VeriSign Class-2 certificates which are too easy to obtain).

Step 4

Obtain a digital certificate and use it to sign in-house developed Active Content. Use it to sign Active Content that is intended for business activities and is known to be trusted.

Setting Trust Levels

Trust levels: Active Content is subject to different restrictions based on the source of origin and the level of trust associated with each source.

Figure 3 illustrated how different security plans are created, in this example using SafeGate's Policy Manager tool, to define different trust levels for different users and sources of Active Content.

The *Company Wide* security plan provides a common level of protection for all users in the organization (using the *Common Protection* pre-built event group), except for the R&D department. It protects everyday users against hostile Active Content that might originate from unknown sources.

The R&D department is covered by the *R&D Department* plan, which provides lowered protection – these guys know what they're doing – but it does validate digital signatures and blocks known hostiles.

The *Trusted Servers* plan covers all the users in the organizations, but only Active Content originating from the designated servers (*Trusted Servers*). Since these sources (company's and business partners' Intranets, software vendors) are known to be trusted, Active Content are allowed storage and network access that would be blocked by the other two security plans.

Digital signatures are validated to make sure Active Content originating from *Trusted Servers* were indeed developed and signed by trusted sources and have not been tampered with. Notice that only digitally signed Active Content is allowed to enter from these servers.

APPENDIX B: ACTIVE CONTENT RISKS

Breaking the Firewall

To learn how vulnerable your network is, check out the hostile demonstration at www.withinreach.co.il/hostiles

A Java Applet or ActiveX Control arriving from the Internet runs on the end-user's machine, behind the firewall, where it can gain access to all network resources.

The Applet or ActiveX can send the information back to the machine from which it came using the same HTTP protocol that was used in fetching it from the Web. The firewall offers no protection for such intrusion attacks.

Installing A Trojan Horse

Once a hostile Active Content gains access to the local disk it can read, modify and delete files. Just as easily, it can install a virus or a Trojan horse on the computer.

Trojan horses for stealing passwords, gaining access to administrative accounts, and retrieving database information are circulating freely on the Internet.

Sending Hostiles by E-Mail

Web surfers might feel immune by using their discretion as to which sites they visit. But hackers can go a step further and send a hostile Active Content to users that are the target of their attack. All it takes is a Web browser and an e-mail client.

Try it yourself. Open one of the hostiles at www.withinreach.co.il/hostiles, then use the *Send Page By E-mail* option of the Web browser to send that hostile to an e-mail address of your choice. You can see that the recipient has little choice in preventing the hostile from running.

Vulnerable e-mail clients: Outlook Express, Outlook 98, Netscape Communicator 4.0, Eudora 4.0 and many others.