ELSEVIER

Computer Communications 21 (1998) 1107–1123

computer communications

Tutorial

Internet/Intranet firewall security—policy, architecture and transaction services

Ray Hunt*

Department of Computer Science, University of Canterbury, Private Bag 4800, Christchurch, New Zealand

Received 22 September 1997; received in revised form 1 April 1998; accepted 3 April 1998

Abstract

The development of Internet/Intranet security is of paramount importance to organisations that plan to gain the economic benefits from interconnection with the Internet. This paper commences by examining firewall policy, focusing on both network service access policy and firewall design policy. Various firewall architectures, ranging from simple packet filters through to screened subnets and proxy gateways, are then discussed. Finally, the various mechanisms by which transactions can be secured over the Internet/Intranet are covered. These include encrypted tunnelling, IPv6, point-to-point tunnelling protocol, secure sockets layer, secure electronic transactions and secure multipart Internet mail encoding. © 1998 Elsevier Science B.V.

Keywords: Firewall design policy; Network service access policy; Packet filter/screening router; Dual-homed gateway; Screened host/subnet; Proxy gateway; Encrypted tunnelling; IPv6; PPTP; SSL; SET; S/MIME

1. Firewall policy

A firewall is a method of achieving security between trusted and untrusted networks and the choice, configuration and operation of a firewall is defined by the policy. The policy defines the services and type of access permitted between trusted and untrusted domains. Therefore, a firewall can be viewed as both a policy and the implementation of that policy in terms of network configuration, host systems, routers, encryption tunnels, authentication procedures and applications systems.

The definition of a firewall policy requires a clear explication of the security perimeter, since different firewall architectures provide different levels of guarantee against attack. Another important term is ''zone of risk'' that generally applies to TCP/IP capable networks, although networks using other protocols such as Netware/IPX, DECnet and SNA can also be vulnerable.

In principle, the zone of risk covers all networks and servers connected to the Internet, including the Internet backbone and related network infrastructures. The objective of the firewall policy is to minimise the organisation's zone of risk by removing the possibility of attack from an external network. In other words the firewall becomes the zone of risk for the trusted network.

It is widely accepted that there is a real risk from insider threats, and it is often stated that there are more insider attacks (for which a firewall is of little value) than external attacks [1]. Insiders usually have more direct access to the systems and the opportunity to abuse privileges. For example, many workstations can be easily reconfigured to grant privileged access and it is then a simple task to run a protocol analyser or decode software. Also, most standard TCP/IP applications, such as Telnet, FTP, rlogin, etc., have weak authentication control and passwords are transmitted in cleartext.

There are two levels of policy that influence the design, installation and use of a firewall:

- network service access policy (NSAP)
- firewall design policy (FDP).

1.1. NSAP

The NSAP defines which services are to be explicitly allowed or denied between trusted and untrusted networks, together with the way in which these services are to be used as well as any conditions for exception to this policy.

* Tel.: +64 3 3642347; fax: +64 3 3642569; e-mail: ray@cosc.canterbury.ac.nz

The NSAP should be an extension to existing business policy that will have already addressed the following issues:

- information value—what value does management places on information?
- responsibility—who is responsible for ensuring the protection of the organisation's information from untrusted networks?
- commitment—what is the organisation's commitment to protecting its information resources?
- domains—what domains should or should not be protected?

Further, business policy should already have implemented controls on such systems as

- virus scanning[1]
- physical security access
- floppy disk controls
- RAID back-up systems.

At the highest level the organisational policy might state:

- information is the strategic resource for the organisation;
- the availability, integrity, authenticity and confidentiality of the information will be protected by every cost-effective measure possible;
- ensuring the availability, integrity, authenticity and confidentiality of the information is a priority for all users at all levels of the company.

Below this level, specific policies are implemented which cover issues such as

- access to services (dial-in, dial-out)
- version controls
- user authentication
- trusted/untrusted network access.

It is at this level that the firewall's NSAP is formulated.

The NSAP must be drafted before the firewall is implemented. It must provide a balance between protecting the trusted network from known risks while providing users with convenient access to the untrusted network. Further, if a firewall denies access to certain services on an untrusted network, it is essential that the NSAP ensures that these controls are not circumvented or disabled. A typical NSAP might

- allow no access to applications or services on the trusted network from the Internet;
- as above, but allow access to a subset of applications or services by way of a secure server (e.g. bastion host);
- allow access from the Internet to selected applications on the trusted network (e.g. e-mail) in conjunction with strict authentication procedures (e.g. challenge/response and one-time password controls).

[1] Contrary to popular belief, firewalls can scan for viruses. This may require scanning at the application layer (mail or file headers). Most common antivirus products such as McAfee, F-Prot, Dr Solomon, Symantecs and Norton's AntiVirus can be configured to achieve firewall virus control.

1.2. FDP

FDP defines how the firewall implements restricted access and service filtering specified by the NSAP and addresses issues such as

- IP address filtering
- encryption tunnelling
- secure socket control to facilitate application access
- audit and accounting control.

This policy is specific to the firewall and defines the rules and procedures necessary to implement the NSAP, but it must take account of the capabilities and limitations of the particular firewall platform as well as the threats and vulnerabilities associated with TCP/IP. For example, if the NSAP forbids access to all applications on the trusted network, then implementing a firewall by way of a packet filtering router is extremely risky.

In principle a firewall can

- permit any service unless it is specifically disallowed
- deny any service unless it is specifically permitted.

However, in practice, only the latter option is used. The first option might unintentionally allow denied services to run on non-standard TCP/UDP ports. Further, some services such as FTP, RPC and X-Windows are difficult to filter [2].

Depending upon the various security and flexibility requirements, some firewalls are more appropriate than others, which means that the NSAP must be carefully designed before the firewall is implemented. For example, dual-homed gateways (Section 2.2.1) and screened subnets (Section 2.2.3) can both be used to implement a ''deny all'' firewall. However, the dual-homed gateway is cheaper but also less flexible than the screened subnet.

In order to arrive at a successful design policy, together with a platform that implements this policy, it is usual to start by restricting all access from the untrusted to the trusted network, and then to specify the following [3].

- What Internet services will the organisation use (e.g. e-mail, Telnet, FTP, World Wide Web (WWW))?
- Where will these services be used from (intra-company, between branches, on a mobile or dial-in basis, by subsidiary organisations, etc.)?
- What additional security features will be needed (e.g. one-time password control, authentication procedures, encryption tunnels, secure sockets, point-to-point encryption, dial-in/dial-back procedures. etc.)?
- What risks result from the provision of these services? E.g. is a 40-bit RSA [4] encryption key adequate for certain government or banking applications? Is dial-in access without formal authentication procedures an acceptable risk?
- What is the cost (e.g. financial, inconvenience) of providing these services? For example, how is key distribution handled? What is the cost of managing a dedicated authentication server?

- What is the balance between usability and security (e.g. if a particular service is too expensive or risky to use should its use be forbidden, thus creating great inconvenience)?

Some services that are inherently insecure may, with the addition of certain technologies, be secured to pose little or no risk. For example, a remote Telnet session can be very vulnerable to packet sniffing for passwords, and would pose a high risk when connecting a machine to a trusted network over an untrusted network such as the Internet. However, with the addition of encryption or strong authentication techniques this risk can be dramatically reduced.

Implementation of the firewall based upon these considerations requires careful use of risk analysis so that the calculated level of risk can be compared with that deemed to be acceptable according to overall company policy [5]. This may result in a change to the initial policy. For example, if the original NSAP denied all dial-in access, certain exceptions to this rule may need to be considered so as to achieve some overall organisational objective.

1.3. Sample policies

1.3.1. Remote access policy

As a specific example a ''remote user advanced authentication policy'' might address dial-in user access from the Internet as well as authorised users on travel or working from home. All such connections should use the advanced authentication service of the firewall to access systems at the site. Policy should reflect that remote users might not access systems through unauthorised modems placed behind the firewall, as it takes only one captured password or one uncontrolled modem line to enable a backdoor around the firewall.

Authorised users may also wish to have a dial-out capability to access those systems that cannot be reached through the Internet. These users need to recognise the vulnerabilities they may be creating if they are careless with modem access. A dial-out capability may easily become a dial-in capability if proper precautions are not taken.

Therefore, both dial-in and dial-out capabilities should be incorporated into the design of the firewall. Forcing outside users to go through the advanced authentication of the firewall should be strongly reflected in policy. Policy might also prohibit the use of unauthorised modems attached to host systems if the modem capability is offered through the firewall.

Since users could run point-to-point protocol (PPP) to create new network connections into a site protected by a firewall, it needs to be considered as part of the overall access policy. Such connections are potentially a backdoor around the firewall, and may be an even greater danger than a simple dial-in connection.

1.3.2. Information server policy

A site providing public access to an information server may wish to incorporate appropriate access controls into the firewall design and policy should reflect the idea that the security of the site will not be compromised in the provisioning of an information service. For example, a Web server that is intended to provide access for Internet users may not need to be behind the firewall at all, as the information provided by this server resides on that machine rather than being drawn from systems on the internal network. As long as the machine is regularly backed up it can operate unencumbered by a firewall and simply be restored if attacked.

It is useful to make a distinction between two fundamentally different types of traffic:

- information-server traffic (traffic concerned with retrieving information from an organisation's information server);
- business traffic, such as e-mail, file transfer, transaction services, etc.

The two types of traffic have their own risks and do not necessarily need to be mixed with each other. Screened subnet firewalls (Section 2.2.3) allow information servers to be located on a subnet and, therefore, to be isolated from other site systems. This reduces the chance that an information server could be compromised and then used to attack site systems.

1.4. Policy evolution

Two considerations drive the formation of an FDP with respect to Internet connections:

- the risk to the organisation's internal information and systems from external threats, e.g. denial of service attacks, IP spoofing, etc.;
- the risk of sensitive organisational information being disclosed as it is transmitted across the Internet, e.g. password file capture, information leakage attacks (e.g. finger), etc.

Once the FDP has been drafted, maintenance and review are important ongoing activities.

1.4.1. Maintenance of the FDP

Unlike many organisational policies, the FDP is not static and may need to change on a day-by-day basis depending upon new vulnerabilities which arise. For example JAVA was considered to be a great invention and the industry was assured by SUN that it was not a security risk. Therefore, as browsers evolved to become JAVA aware, JAVA code (applets) passed through firewalls. It is likely that JAVA never appeared in any company's firewall policy as it was probably considered to be part of the WWW. One large multinational decreed from the highest level that JAVA be disabled on all browsers, thus demonstrating the dynamic nature of an FDP. Other examples of policy maintenance

include changes to a network's filtering rules as well as rule changes resulting from the introduction of new services.

1.4.2. Review of the FDP

It is most important that the FDP remains under constant review to ensure that the policy reflects the current state of play. As a result of FDP maintenance, the original policy can become unrepresentative of reality and this can introduce security holes. Examples include: change of the systems expert; wrong versions of software being loaded following a system crash. In many of these cases problems may not be detected until after a security breach has occurred.

1.5. Installing and operating a firewall

Once the decision is made to use firewall technology to implement an organisation's security policy, it is then necessary to install a cost-effective firewall that provides an appropriate level of protection. In general, a firewall should provide the following levels of protection:

- support and not impose a security policy;
- support a ''deny all services except those specifically permitted'' design policy (even if this policy is not initially implemented);
- accommodate new facilities and services should an organisation's security policy change;
- contain advanced authentication measures, such as encryption, challenge/response systems, and should contain the hooks for installing these facilities;
- employ filtering techniques to permit or deny services to specified hosts as needed;
- use flexible and user-friendly IP filtering and be able to filter on as many attributes as possible, including source and destination IP address, source and destination TCP/UDP port, protocol type, and inbound/outbound interfaces;
- use proxy services for applications such as FTP and Telnet, so that advanced authentication measures can be utilised at the firewall.

It will also assist if the firewall supports proxies for services such as NTP, NNTP, X-Windows, Finger, HTTP, and certain web browser software (Section 2.2.4). The firewall should also have the ability to centralise simple mail transfer protocol (SMTP) access, thus reducing direct SMTP connections between site and remote systems. This results in centralised handling of site e-mail.

The firewall should have the ability to concentrate and filter dial-in access as well as logging suspicious activity. If the firewall requires an operating system such as UNIX or Windows NT, then a secured version of the operating system should be part of the firewall, with other security tools as necessary to ensure firewall host integrity. The operating system should have all patches installed and be developed in a manner that its strength and correctness is verifiable. It should be simple in design so that it can be understood and maintained.

Some organisations have the capability to put together their own firewalls, using available software components and equipment or by writing a firewall from scratch. Trusted Information Systems (TIS) Internet Firewall Toolkit [6] is a good example of a company that offers ''firewall construction kits''. At the same time, there are many vendors[2] [7] offering a wide range of services in firewall technology which include

- provision of the necessary hardware and software
- development of security policy and carrying out risk assessments
- security reviews and security training.

Consideration of the following questions may help an organisation decide whether or not it has the resources to install and operate a successful firewall.

- How will the firewall be tested?
- Who will verify that the firewall performs as expected?
- Who will perform general maintenance of the firewall, such as backups and repairs?
- Who will install updates to the firewall, such as for new proxy servers, new patches, and other enhancements?
- Can security-related patches and problems be corrected in a timely manner?
- Who will perform user support and training?

As a general rule it is desirable that sites:

- standardise operating system versions and software to make installation of patches and security fixes more manageable;
- institute a program for efficient, site-wide installation of patches and new software;
- use services to assist in centralising system administration, if this will result in better administration and better security;
- perform periodic scans and checks of host systems to detect common vulnerabilities and errors in configuration.

2. Firewall architecture

There are many different interpretations of the term firewall and this can be a source of confusion. One basis for defining a firewall is the OSI 7-layer model (Fig. 1) which provides a clearer picture than does the TCP/IP model.

[2] Examples include Digital's Alta Vista Firewall, Secure Computing's Firewall for NT, Eagle NMS (Raptor Systems), ANS Interlock Service 3.06 (ANS CO + RE Systems), Borderware Firewall Server, Firewall-1 v2.0 (Checkpoint Software), Black Hole 3.0 (Milkyway Networks Corp.), Sidewinder, Gauntlet (Data General), GFX Internet Firewall (Global Technology Associates).

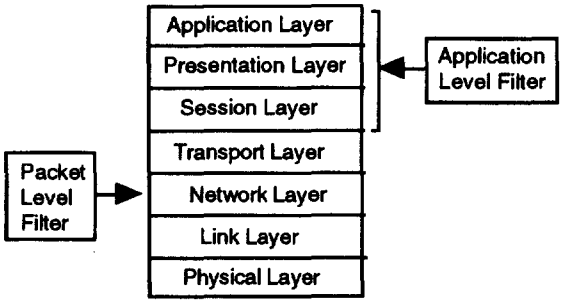

Fig. 1. Firewall architecture in relation to the OSI model.

It can be seen that firewall architecture consists of two levels:

- Packet-level firewalls that operate at the network (IP) and transport (TCP) layers. These are commonly referred to as screening routers or packet filters and block transmission of certain classes of traffic;
- application-level firewalls which operate at the session, presentation and application layers. They are usually implemented using dedicated hosts running specialised software and can also be referred to as bastion hosts or proxy servers, usually running under UNIX or Windows NT. They can also provide relay services to compensate for the effects of the filter(s).

Another important term often used in conjunction with a firewall is ''gateway'', and Internet firewalls are often referred to as secure Internet gateways. However, there is a more specific use of this term, as can be seen in Fig. 2. The network occupied by the gateway is often referred to as the demilitarised zone (DMZ) [8].

The gateway in the DMZ may consist of both an internal and external machine, as shown in Fig. 3. Normally these two gateways will have more open communication through the inside packet filter than the outside gateway has to other internal hosts. The outside packet filter can be used to protect the gateway from attack, while the inside gateway can be used to guard against the consequences of a compromised gateway.

2.1. Packet filters/screening routers

As packets pass through the router their filtering is based upon a set of rules established by the NSAP. Filtering based upon one of more of the following criteria are commonly applied:

- source IP address
- destination IP address
- TCP/UDP source port
- TCP/UDP destination port.

Not all packet filters can filter on TCP/IP port numbers. Some can examine which of the network interfaces a packet arrived at and then use this as further filtering criterion.

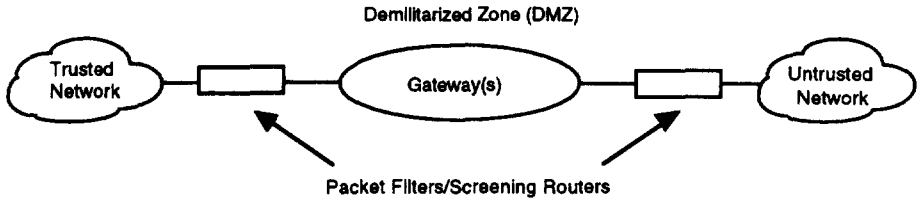

Fig. 2. Firewall design (filter/router and gateway).

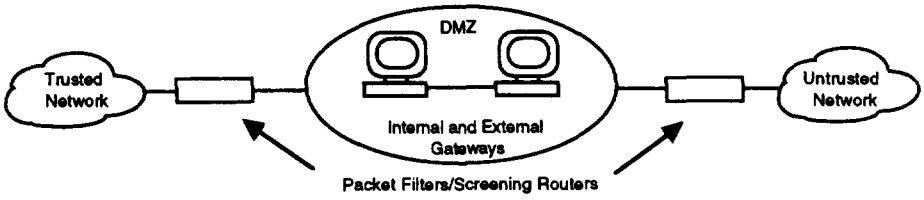

Fig. 3. Firewall design (filter/router with internal and external gateways).

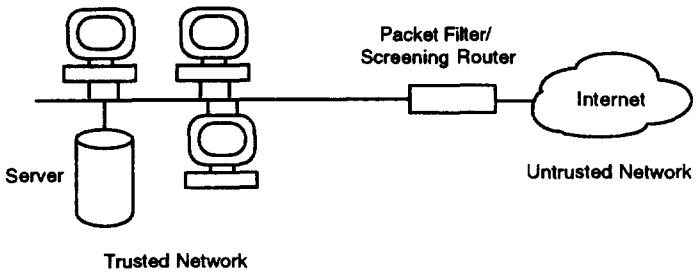

Fig. 4. Firewall using a packet filter/screening router.

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS
Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS
Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS
Sync your system to PACER to automate legal marketing.