Security · Networking · VPN Clients

Documentation
Printed guides
HowTo
Q & A
Interoperability
Reference
Encyclopedia
Other resources
Training
Self-test courses
Download
Software
Service Tool
Support
Contact

Firewall security

Firewall security – an overview

This overview of Firewall security will give you a basic understanding of the role of firewalls in the protection of local networks, as well as a brief description of how firewalls work and why it is important that a company has a carefully planned security policy.

For a more thorough walkthrough of the actual techniques behind firewalls you may want to turn to **'More about Firewall security'**.

What is a firewall?

A firewall protects your local network from unwanted visits or attacks from the Internet.

Having your company's local network connected to the Internet will bring you a long list of benefits. Suddenly the whole world can be reached from the office. But there's a downside: the whole world may be able to reach your office too…

Usually, visits from the outside world are welcome. Business is about establishing contacts, and it would make little sense to turn an interested visitor or a prospective customer away.

However, there are always a few interested visitors that businesses can do without. On the Internet thieves and spies come in the form of hackers, so businesses need to put a lock on their 'door' to the Internet – which is where the firewall comes into the picture.

Applying a firewall ensures that there is only one point through which data can enter the business' local network, and – importantly – that traffic through that point can be monitored.

The firewall can restrict access to your local network from the Internet by effectively hiding all or parts of the local network from outsiders, and by

applying various filtering and relaying mechanisms. Thus, you are able to let in the people you want and at the same time bar those that you don't want from gaining access.

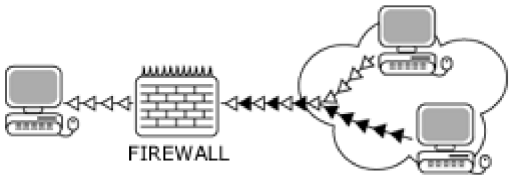

The firewall filters away unwanted data

to the top

How do firewalls work?

Firewalls are clever mechanisms – typically a combination of hardware and software. Myriads of such possible combinations exist, allowing a great degree of flexibility. The filtering mechanism of a firewall looks at things such as where the visitor data comes from, its destination, its contents or the like.

All computers connected to the Internet use the TCP/IP standard for their communication. That way they share a common 'language'. Using TCP/IP, the message to be sent – it could be an e-mail or an HTML page – is split into smaller segments, packets, for convenient transmission. It is these packets that the firewall examines and, provided they're acceptable, relays.

A firewall can also simply – but efficiently – hide the structure and contents of your local network from outsiders. Traffic *from* the local network to the Internet can easily get out, whereas traffic *to* the local network from outsiders on the Internet cannot get in unless specifically invited.

This process of hiding is known as Network Address Translation, NAT. NAT is just one of a range of firewall techniques, but it provides a high level of security, as it is no easy task for a hacker to work his way into something he does not even know is there.

Security policies

Exactly what a firewall does and how it does it depends on the security policy of the particular business. Whereas one firewall may rigidly block most incoming traffic in one business, another business may have a firewall that only employs a subtle filtering.

In other words: the firewall is an implementation of the security policy. Thus, for the firewall to serve a clear purpose, the security policy must have clear and realistic goals and match the needs of the business.

It is the business' system administrator that implements the security policy and its access control measures, so that the firewall is set up according to the communication needs of the specific business. Some businesses will be interested in combining the firewall's access control with an advanced access log to keep track of all access attempts, some in combining it with alarm functions that will alert the system administrator of suspicious messages.

The access control settings that define exactly what is allowed to pass through the firewall are known as firewall rules. Exactly how implementation takes place can range from applying the system administrator's own firewall rules to installing and running an all-in-one tailor-made professional product.

The most basic – and safest – policy of all would be '*block all traffic*'. Such a restrictive and inflexible policy would, however, defeat the object of having an Internet connection. The rules to be implemented must therefore ideally offer a high degree of flexibility as well as a high degree of security.

A professional firewall package often helps the system administrator in finding a suitable balance by providing a flexible firewall set-up interface which supports a wide range of possibilities. Typically, rules can be defined covering different types of traffic (such as incoming, outgoing, forwarding, etc.) and different ways for the firewall to scan messages (on the basis of source, destination, contents, etc.).

The human factor

Knowing that a firewall provides a high degree of security, it is still worth noticing that even the best firewall in the world cannot protect your business from certain types of security leaks.

Firewalls can have difficulties detecting virus-infected programs sent as e-mail attachments. Such programs are usually concealed as something innocuous and are furthermore compressed, making it hard for the firewall to scan them. If receiving such attachments from distrusted sources, opening and running them should be done with an amount of caution. A firewall, therefore, should be supplemented by anti-viral software.

Straightforward as it may seem, it is also still worth remembering that a firewall does not protect against attacks that do not go through the firewall…

A firewall, for instance, cannot protect your local network from harmful files brought in from the outside world on a disk or other data storage medium. Neither will the presence of a firewall in itself prevent employees from copying valuable data onto a disk and taking it out of the secure environment – even though the copying may be with the best of intentions.

Similarly, a firewall which is too strict and inflexible may lead even the most conscientious employees to wish to circumvent it by using their own modems or the like. Again, the motives may be perfectly noble, such as a desire to work faster and more effectively, but internal network security could be severely compromised.

Consequently, it is important that the company's security policy is known and understood by all employees using the internal network in order to minimise such other risks.

to the top

More about Firewall security

This section outlines the techniques behind firewalls, with a brief description of the most widely used data filtering approaches.

Main types of firewall

All computers connected to the Internet use the TCP/IP standard for their communication. That way they share a common 'language'. Using TCP/IP the message to be sent – it could be an e-mail or a HTML page – is split into smaller segments, packets, for convenient transmission. It is these packets that a firewall examines and, provided they're acceptable, relays.

Basically, two main types of firewall exist, although in everyday use the difference is rather subtle as the techniques behind the two types have more or less merged of late. The two types are: application firewalls and network firewalls, each focussing on a different part of the TCP/IP data packets.

Network firewalls generally accept or reject messages according to their source or destination. This way the firewall can block messages *from* known unreliable sources, e.g. a specific domain, and *to* local network addresses that outsiders aren't meant to access. In short, they'll only relay acceptable messages.

The general approach taken by application firewalls differs only slightly, although they are often set up as special proxy servers that refuse specified traffic between networks. They are sometimes – somewhat crudely – likened to guest lists for parties: basically, only those whose names are on the list are let in and allowed to circulate at the party.

The biggest difference between the two types of firewall is that they focus their analysis on different parts, layers, of the TCP/IP data packets.

Looking at layers

A data packet is made up of several layers, each containing different elements of the sent message. Most of the layers have actually got to do with delivering the data packet safely to the receiver: one layer is responsible for establishing and keeping open the communications channel, another layer is responsible for routing the data from point to point, etc.

If we look at the commonly used so-called OSI model, there are seven layers to a data packet, but TCP/IP data packets make use of only four. Thus, there exists another conceptual model, the so-called DARPA model (named after the agency that developed TCP/IP), which contains only the four layers used by TCP/IP:

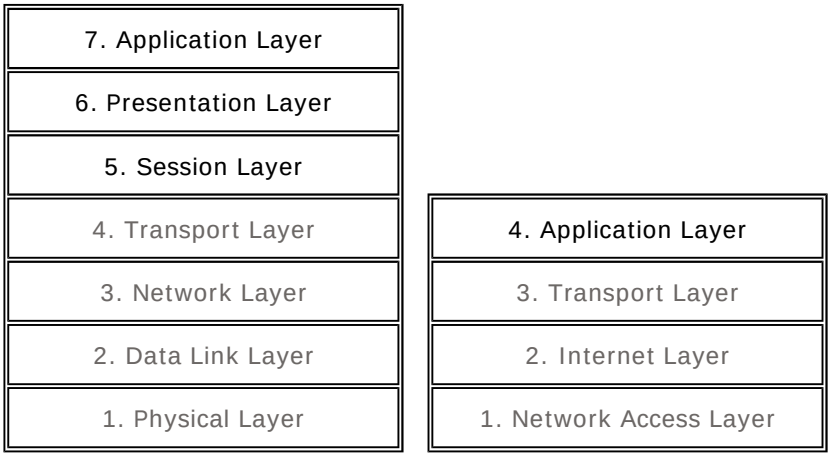

The 7-layer OSI model (left) and the 4-layer DARPA model (right)

Basically, the top three layers of the data packet (the top layer in the

DARPA TCP/IP model) deal with what the user actually gets to see: the program-to-program communication, etc. The bottom four layers (three in the DARPA TCP/IP model) deal with how the operating systems handle the data transfer.

Looking at the models in which the layers are placed like bricks, one upon another, it probably comes as no surprise that these constructs are also known as 'stacks'.
When data is sent, it is passed from one layer to the next down through the stack before being transmitted over the physical network. When the data is received at the remote end, it is passed from one layer to the next upwards through the stack to the receiving application.

When data is sent, each of the stack´s layers add information (such as a destination address, a checksum, or the like) to the data packet in order to ensure safe delivery. This information is enclosed in a so-called 'headers' and/or 'trailers' placed in front of and/or behind the layer data - a process known as 'encapsulation'.
As each layer sees all the information it gets from the layer above it as data, it places its header and/or trailer wrapped around it. This means that when the original message exits onto the physical network for transmission it is enveloped in multiple wrappers, one for each layer the data has passed on its way down through the stack.

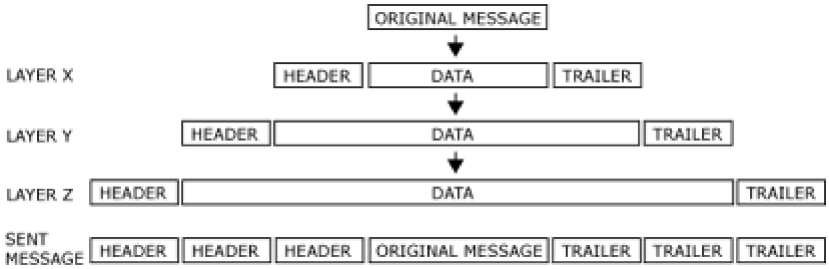

Example of data encapsulation for network transmission: each layer adds information wrapped around the original message. Later, when the data is received by the remote party, the process is reversed (decapsulation).

When data is received the process reverses. Each layer interprets and then cuts off its header and/or trailer information as the data passes upwards through the stack. This process of stripping off headers and/or trailers is know as 'decapsulation'.

Firewalls and layers

As the name implies, the network firewall concentrates on the layers of the TCP/IP data packet that have to do with networking: the parts responsible for routing the data from point to point, etc., notably the Internet layer (see illustration of layer models in '**Looking at layers**' above).
At its most basic level a network firewall will always look at the data packet and ask itself: is it possible to route (relay) the data packet, and is it safe to do it? Whether it is safe to route the packet or not depends on the rules laid down by the system administrator. If the rules do not allow routing of the particular packet, the packet is simply discarded.

An application firewall, in turn, concentrates its analysis on the application layer of the TCP/IP data packet: the layer that enables program-to-program communication. This means that in order to detect whether something suspicious is going on, the application firewall – which typically is a proxy server – has to know the application that is to be used to handle the message (e.g. FTP, HTTP). This often means that a relatively complex

configuration is needed, but on the other hand the complexity offers the possibility of e.g. more detailed audit reports.

Both firewall approaches have their advantages, and it is increasingly possible to bring together the best of both worlds.

Filtering data packets

Firewalls examine incoming and outgoing data packets and, according to the firewall rules defined by the system administrator, either let the data packets through or block them out. In other words, the firewall works as a filter.

The process of filtering data packets is relatively complicated and to a large extent depends on the system administrator´s settings. It is, however, possible to describe in some detail the process taking place when a data packet is filtered.

Looking at the firewall as a filter, the filter works by checking data packets against a list of the firewall rules specified by the system administrator. In fact the filter contains no less than three lists with the firewall rules on them, known as firewall chains, against which it repeatedly checks the data packets. The three firewall chains are called Input, Forward and Output.

Each firewall chain works as an independent checklist of firewall rules. The rules are looked at hierarchically, with the more specific rules applied first. The more general and all-encompassing a rule is, the later in the sequence it is applied. Thus, if the first (probably very specific) rule does not apply for the packet, the next rule in the chain´s hierarchy is consulted, and so forth.

When a data packet reaches the firewall, the Input chain is used to decide what is going to happen to the packet. If the packet survives that stage, the system will decide where to send (route) the packet next. In case the packet is meant for another computer, the Forward chain is used. Finally, before the firewall lets go of the packet, it will consult its Output chain.

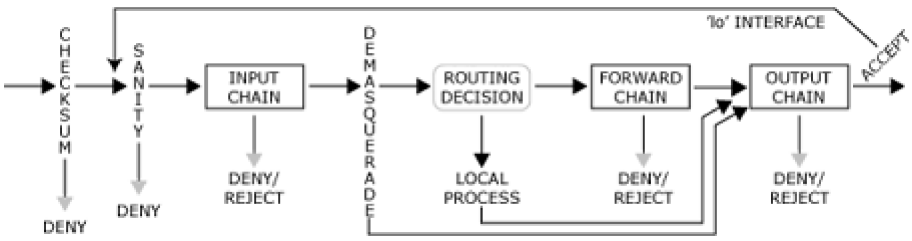

The path of a data packet coming into a firewall

The illustration above outlines the path of a packet coming into a firewall. The packet enters on the left hand side, and passes an array of tests on its way:

Checksum: a test to ensure the packet has not been corrupted. If it has, it is denied further access.

Sanity: if packets are for some reason malformed, they may be difficult to check against the firewall rules, and are denied further access at this stage. Sanity checks are carried out before each of the firewall chains, but the sanity check before the Input chain is by far the most important.

Input chain: the first of the three firewall chains. The packet is tested

against the firewall rules. If the result is not 'deny' or 'reject', the packet is able to continue.

Demasquerade: in case the packet is a reply to a previously masqueraded packet (one in which the sender´s IP address details have been changed by NAT, Network Address Translation), the packet is demasqueraded and will take the fast lane straight to the output chain.

Routing decision: at this stage the packet´s destination field is examined in order to decide whether the packet should go to a local process (see below) or be forwarded to a remote computer via the Forward chain.

Local process: should the packet be destined for the computer running the firewall itself, the computer is able to receive packets immediately after the routing decision step. The computer is of course also able to send packets originating from its own local processes. These packets will pass through the routing decision step and on towards the Output chain.

'lo' interface: this is for packets concerning local processes (see above). When a packet originating from a local process is destined for a local process, it will pass through the Output chain with a so-called 'lo' (loopback) interface, and then return via the Input chain, also with a 'lo' interface.

Forward chain: the second of the three firewall chains. This chain has to be passed by any packet (unless created by a local process) attempting to pass through the firewall on its way to another computer. Again, the packet is tested against the firewall rules. If the result is not 'deny' or 'reject', the packet is able to continue.

Output chain: the third and final of the firewall chains. All packets have to pass this chain before being sent out. If the result is not 'deny' or 'reject', the packet is able to continue.

The firewall rules specified by the system administrator thus play a very important role in the protection of the business´ network. It is, however, again worth remembering the great impact firewall rules can have on a business´ communication. The above example of a data packet´s path through the firewall illustrates how easily a data packet can be discarded if it does not meet the firewall´s criteria. This stresses the effectiveness of the firewall, but it also stresses the need for the system administrator to be aware of the firewall´s implications on the business´ ablity to communicate.

to the top

Test your knowledge

1.

The access control settings that define exactly what is allowed to pass through the firewall are known as...?

- Firewall groups
- Firewall rules
- Guest lists

2.

What role should an organisation´s security policy play when setting up a firewall?

- No particular role. Firewalls and security policies are separate things.
- A limited role. Security policies do not really match the technical nature of a firewall.
- A great role. The firewall is an implementation of the security policy.

3.

What does NAT do?

- NAT hides the structure and contents of a local network from outsiders by translating network addresses.
- NAT, Network Attack Termination, is a protocol that hunts down and destroys malicious data packets.
- NAT allows traffic from the Internet to gain easy access to a local network, whereas traffic to the Internet from the local network cannot get out unless specifically invited.

4.

Should your company still use anti-virus software when it has set up its firewall?

- No, a firewall offers total virus protection.
- Sometimes, especially when the firewall is down for maintenance.
- Yes. You should always be on guard, firewall or no firewall.

5.

What is the main difference between application firewalls and network firewalls?

- They focus on different parts of the TCP/IP packets.
- One protects applications, the other protects networks.
- There is no difference at all.

6.

A data packet can be divided into...?

- Firewall rules
- Layers
- Applications

7.

A company´s security policy should be known and understood by...?

- The company´s decision makers.
- The company´s decision makers and selected members of the technical staff.
- All employees using the internal network.

8.

The firewall treats its firewall rules hierarchically, but in which sequence?

- It checks its firewall rules in alphabetic order.
- The more general rules are looked at first.
- The more specific rules are looked at first.

9.

What does a firewall examine and, if acceptable, relay?

- Hackers
- Packet filters
- Data packets

10.

Generally, there are seven layers to a data packet, but how many layers do TCP/IP data packets make use of?

- Four.
- Three.
- Five.

Evaluate Reset