
Firewall systems

Joaquin Garcia-Alfaro

Recommended minimum reading time: 3 hours



First edition: February 2023
© of this edition, Fundació Universitat Oberta de Catalunya (FUOC)
Av. Tibidabo, 39-43, 08035 Barcelona
Authorship: Joaquín García-Alfaro
Production: FUOC

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.3 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. The terms of the license can be consulted in <http://www.gnu.org/licenses/fdl-1.3.html>.

Contents

Introduction	5
Objectives	6
1. Introduction to firewall systems	7
2. Evolution of firewall systems	9
2.1. First generation: packet filtering at network level	9
2.1.1. Advantages and disadvantages of first generation firewalls	14
2.2. Second generation: stateful inspection at the transport level	14
2.2.1. Stateful protocol filtering	15
2.2.2. Advantages and disadvantages of second generation firewalls	16
2.3. Third generation: application layer data processing	18
2.3.1. Advantages and disadvantages of third generation firewalls	19
3. Implementation of perimeter security through firewall systems	23
3.1. Single point architectures	23
3.2. Architectures with perimeter networks	24
Summary	29
Glossary	31
Bibliography	32

Introduction

Firewall systems represent an effective element in preventing cyber attacks. Already consolidated as indispensable elements to guarantee the protection of computers and computer networks, firewall systems represent a practical implementation of the concept of *access control*, both at the system level (for example, to prevent attacks against personal computers at the scale of application) and at the network level (to prevent attacks from hostile networks against personal or corporate networks).

This module assumes that you have a basic knowledge of how computer networks operate and some notions of cyber security. More precisely, and although there is no need for extensive knowledge, we consider that you are already familiar with the Open Systems Interconnection (OSI) reference model, as well as with the TCP/IP family of protocols.

In this module, we present a general introduction to firewall systems. One of the main concepts related to these systems is *packet filtering*. The first firewall systems (currently known as *first generation firewalls*) are actually routers using filtering rules to build an entry barrier and thus give way to what is known as *perimeter security*. The aim is to separate vulnerable environments, generally private networks, from hostile environments (for example, public networks, such as the Internet). Thus, in this module, we will first deal with firewall systems based on packet filtering, responsible for processing and inspecting traffic at the network layer level (level 3 of the OSI reference model). Then, we will also see configuration examples of the second and third generation firewall systems, responsible for filtering traffic with a more thorough inspection (levels 4 and 7 of the OSI reference model). Finally, we will discuss other aspects related to firewall systems that may be of interest to you, such as the most used current architectures, as well as the advantages and limitations of implementing perimeter security and its deployment in different types of networks using firewall systems.

Objectives

The aims to be achieved with this material are the following:

- 1.** To know what a firewall system is, and to understand how it can be used to provide protection to a computer network, thus preventing computer attacks.
- 2.** To know and understand the limitations of firewall systems.
- 3.** To know and understand some strategies and architectures related to firewall systems, starting with information filtering.
- 4.** To understand what perimeter security policies are and their deployment on the network using firewall system architectures.

1. Introduction to firewall systems

Firewall systems are hardware or software components that control traffic in and out of a system. In general, they are usually used to provide an access control mechanism on computer networks and they allow separating an internal part (in which the computers involved are considered as trusted) from other computers located in the outside (potentially hostile).

A **firewall system** is responsible for separating computer networks and controlling the traffic that circulates there. The control consists of allowing, denying or redirecting communications from one of the networks to the other, through an access control or use of the associated network protocols.

Therefore, a firewall system serves as a barrier in a network. It can be used to block incoming or outgoing traffic, to prevent unauthorized access, etc. In this context, the concept of a security policy is important. In other words, a firewall system makes it possible to implement the security policy associated with the network.

A firewall system is one of the possible security mechanisms that allow implementing the rules of a security policy. More specifically, the rules relating to access control at the network level and which are related to the perimeter security of the network.

When installing and configuring a firewall system, the following must be kept in mind:

- All traffic leaving or entering the network must pass through the firewall system. This can be achieved by physically blocking all access within the network.
- Only authorized traffic, defined in the system's local security policies, will be able to bypass the block.
- The firewall itself must be protected against possible attacks or intrusions.

Complementary reading

The following article provides a detailed introduction to firewall systems, from an academic perspective: **Kenneth Ingham; Stephanie Forrest** (2002). "A history and survey of network firewalls", The University of New Mexico, Computer Science Department, Technical Report 2002-37. <<http://www.cs.unm.edu/~treport/tr/02-12/firewall.pdf>>

Local firewalls

Although in this module we refer to firewall systems as elements to control traffic on a computer network, the same concept can be applied to attack prevention by installing a local firewall to directly protect computers and devices.

Firewall systems, as we know them today, appeared at the end of the eighties, developed by the DEC and AT&T companies. In 1991, the first commercial firewall, the DEC SEAL, appeared. Today, firewall systems are a very important element not only in network devices, but even in personal computers. There are different technologies for implementing firewalls, and above all, there are many architectures or ways to configure firewalls in a network. In this module we will see some of the most prominent ones. It is important to note that we will focus on the use of firewalls in TCP/IP networks, although the use of firewalls is not exclusive to these particular protocols.

Likewise, it is important to always keep in mind what we want to achieve with a firewall system. The most common use is to control incoming and outgoing traffic from one network to another. Generally, it is about protecting the internal network of an organization against a hostile network, such as the Internet. In the internal network we can find personal computer equipment, printers, mobile devices, smartphones, servers, etc. The presence of servers, if they offer services to the external network, may require special treatment in firewall systems. We refer, for example, to web, email or file sharing servers that the organization wants to offer on the Internet. In general, these servers are not treated like personal computers when designing their protection.

As we will see in the following sections, although firewalls provide many security measures, it should be noted that they are not a definitive or unique solution to the problem of network security. There are many threats that cannot be covered by firewall systems. In this sense, a very important aspect, as we will see, is that it is difficult to protect against an internal attacker with a firewall. The firewall itself, like any computer system, can present zero-day vulnerabilities and be a victim of malicious programs within the operating system in which it runs. In addition, firewall systems can have a certain degree of bad press among network users, as they often see them as a trade-off for their convenience or ease of use of network services.

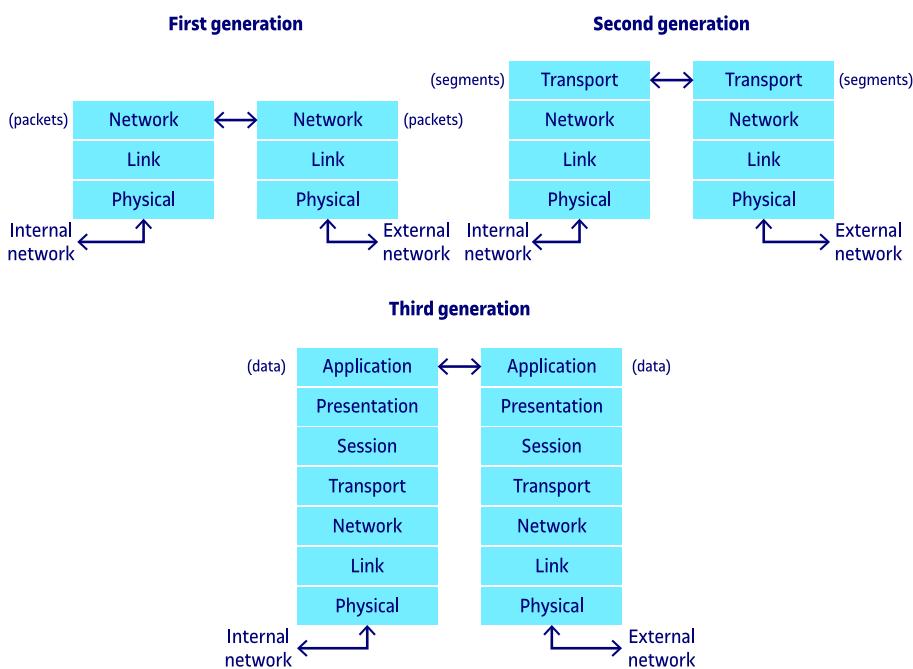
Complementary reading

The following article (available online) discusses the origins of firewall systems in more detail: **Frédéric Avolio** (June 1999). "Firewalls and Internet Security", The Internet Protocol Journal (volume 2, number 2, pages 24-32). <<http://ipj.dreamhosters.com/wp-content/uploads/issues/1999/ipj02-2.pdf>>

2. Evolution of firewall systems

A simple way of referring to firewall systems is in relation to the order of appearance; for example, first, second, and third generation firewall systems, respectively. The main difference lies in the level at which traffic filtering is done, that is, at the layer on which they act. Figure 1 shows this evolution, with respect to generations and performance layers, using the OSI reference model for the layers.

Figure 1. Evolution of first, second and third generation firewall systems regarding the analysis and action layers, using the OSI reference model.



Next Generation Firewall

Other names, such as last generation firewalls, or new generation firewalls (NGFW) are sometimes used in the commercial field, to refer to additional functions of third generation firewalls, such as support for encrypted traffic inspection with protocols such as TSL and SSH, active directory integration within the firewall, malware filtering, built-in network intrusion prevention, etc. Later, in the section on third generation firewalls, we will discuss some of these additional features.

In the first subsection of this section we will deal with the first generation firewall systems, based on the use of routers with packet filtering. In the following subsections, we will also discuss second generation firewalls (known as *stateful inspection firewalls*) and third generation firewalls (known as *application layer* or *layer 7 firewalls*, in relation to the OSI reference model). We will see representative examples, as well as the advantages and limitations for each type.

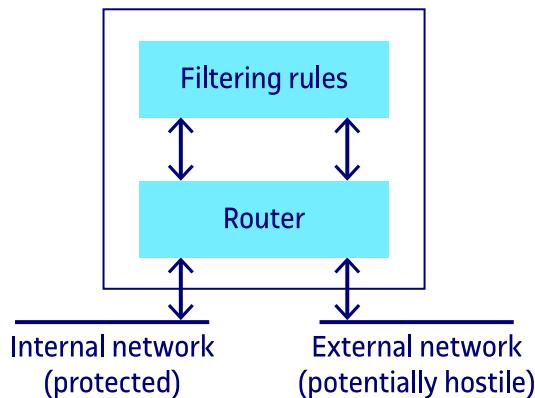
2.1. First generation: packet filtering at network level

As shown in figure 2, a router with packet filtering is a device that routes and inspects traffic at the network level (IP packets, for example). The router will decide whether or not to let the traffic pass according to filtering rules associated with a security policy.

Packet filtering firewalls

The router that filters traffic is also called a *screening router*.

Figure 2. First generation firewall: packet filtering firewall



Filtering rules are responsible for determining whether a packet is allowed to pass from the internal part of the network to the external part, and vice versa, by checking for legitimate data traffic between both parties.

Filtering rules associated with packet-filtering routers typically use information present in network packets traversing the firewall system. That is, they indicate which packets may or may not pass by looking at the headers of the associated protocols (for example, protocols such as ARP, IP or ICMP). This associated information can be:

- the source and destination addresses of the packets,
- the protocol type associated with the packets,
- the ports of origin and destination,
- the message type (at the network level),
- the contents of the packets (at the network level),
- the size of the packet,
- etc.

Note that although first-generation firewalls tend to use source and destination ports in their filtering rules in addition to protocol type, this does not mean that they perform a thorough inspection of the connections in the transport layer (layer 4, in relation to the OSI reference model). We will see later that second generation firewall systems, in addition to consulting this source and destination port information, can additionally check the connection, which we do not expect from a router acting as a firewall at the network level.

If the router uses a first-match filtering strategy, each packet arriving at the device will be compared with the filtering rules, starting at the beginning of the list until the first match is found.

If there is a match, then the action indicated by the rule is triggered (for example, deny the packet, accept it, or reroute it).

If no match is found, the *default policy* will be consulted to determine what action to take (for example, let the packet pass or discard it). If it is, for example, a deny by default policy, if there is no match with the packet, it is discarded.

A *deny by default* policy tends to be more expensive to maintain, as it will require the administrator to explicitly state all services to be kept open (others will all be denied by default). A *default acceptance* policy seems simpler to administer, but it increases the risk of receiving attacks against the network, since it requires the administrator to explicitly indicate which packets to discard (the rest, by default, will be accepted in their entirety).

Default policies

A deny by default policy is also called a *deny all* or *closed policy*, while a default acceptance policy is also called an *allow all* or *open policy*.

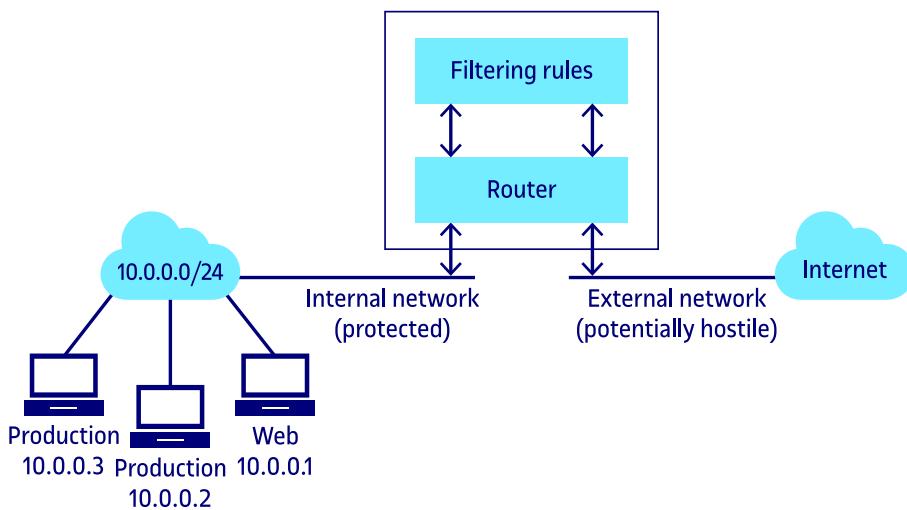
Most of the time, a default denial policy is chosen as a security measure. This strategy is sometimes called *the fail-safe security principle*.

A firewall system fulfils the **fail-safe** principle if it rejects an unexpected event, such as a packet for a new service.

Figure 3 shows an example diagram for the following security policy (very simplified, for ease of understanding):

- We assume a deny by default policy.
- All systems on the internal network (with network address 10.0.0.0/24) can access any service on the external network (Internet).
- External systems cannot connect to any internal system except the web server (computer with IP address 10.0.0.1).

Figure 3. Example of packet filtering with a first generation firewall.



The security policy configuration mentioned above, once applied by the router shown in figure 3, will allow all packets that have an internal network IP address as source and an external network (Internet) IP address as destina-

tion to pass through. It will also allow part of the traffic destined for the internal network, so that the computer with the 10.0.0.1 IP address can respond to requests from outside. Finally, since a deny by default policy is applied, the firewall will not allow packets that do not comply with the above rules to pass through. In addition, the rules corresponding to the response traffic must also be added to the router configuration, that is, rules representing the following two cases:

- allow traffic that enters the internal network and comes from web services (tcp source port 80)
- allow traffic that leaves the internal network and comes from the web server (10.0.0.1 source port tcp 80 and source IP address).

For the sake of simplicity, in this example we have not considered other web-related ports, such as port 443 related to TLS within HTTP (i.e., HTTPS), or the use of the Domain Name System (DNS) that would be needed in a real scenario.

So, the above example shows us that packet filtering is based on using the information available in the headers of a protocol such as IP, stateless, using data such as, for example, the source address, the destination address, source and destination ports, etc.

With this kind of data it is easy to specify rules of the type “accept all outgoing traffic intended for web services” (the way to do this is to consider tcp 80 as the destination port). On the contrary, it does not allow expressing rules of the type “accept all HTTP traffic only if it is not being used to download music”. To be able to solve this case, it would be necessary to use a gateway-based firewall at the application layer, so that the firewall would analyze the traffic at the application layer level and detect whether it is being used to download music or not.

As stated before, each filtering rule has an action associated with it. This action determines what the firewall system should do with each packet that meets the conditions associated with the filtering rule. Examples of used actions are those that indicate that the packet can be accepted or rejected.

If a packet is rejected, there is the possibility of generating error messages. Typically, this will involve using ICMP-type traffic to notify the device that originated the packet of the firewall’s decision to reject it. More specifically, ICMP type 3 (Destination Unreachable) messages, with codes like those shown in table 1. A second possibility is to silently reject the packet, without creating an ICMP message to inform the device that originated the packet.

ICMP Error Codes

Codes 9 and 10 were specially added to the ICMP specification for use with filtering systems. However, many firewall systems continue to use only codes 0 and 1, which were originally intended for other purposes.

Table 1. ICMP codes that a firewall can send when it rejects a packet

Type	Code	Description
3	0	Destination network unreachable
3	1	Destination host unreachable
3	9	Network administratively prohibited
3	10	Host administratively prohibited

Complementary readings

The following article provides more information about filter rule ordering and organization strategies for conflict resolution and configuration issues: *Garcia-Alfaro et al.* (2007). "Management of Exceptions on Access Control Policies", *22nd IFIP TC-11 International Information Security Conference, New Approaches for Security, Privacy and Trust in Complex Environments*, 97-108, Springer Nature. <http://dx.doi.org/10.1007/978-0-387-72367-9_9>

On the other hand, the following article provides information on troubleshooting strategies for configuring firewall systems: *Garcia-Alfaro et al.* (2008). "Complete Analysis of Configuration Rules to Guarantee Reliable Network Security Policies", *International Journal of Information Security*, 7(2):103-122, April 2008, Springer Nature. <<http://dx.doi.org/10.1007/s10207-007-0045-7>>

Generating ICMP messages or not has advantages and disadvantages. On the one hand, sending the informative ICMP message means that the source can close the connection immediately, without needing to waste time and without trying to retransmit the rejected packets. But it adds the following problem: the ICMP message generated by the firewall system can be interpreted in different ways by the source equipment that receives it. It may even result in a performance penalty for the firewall system or provide information that potential attackers can later use against the system. So, for many authors, it is considered safer to not send informational ICMP messages and simply reject packets silently.

Finally, the order in which filtering rules are processed is a very important parameter to consider, as it can be used as a mechanism to resolve conflicts and configuration issues. For example, if two conflicting rules are found (for example, one that accepts the packet and one that rejects it, a *first-match* strategy will give priority to the first rule that matches the conditions of the treated packet. This first rule will therefore decide the action to be executed for the packet in question. In contrast, using a *last-match* strategy will give priority to the last rule that matches the conditions of the packet, perhaps with a completely different action from the rules that precede it in the order.

In general, it is the firewall system administrator who will decide the resolution strategies regarding the selection and order of the rules. These decisions can affect the efficiency of the firewall system, at the expense of simplifying its configuration. While a configuration based on first- or last-match strategies simplifies configuration, it can also have a negative impact on filtering efficiency, as it can place a burden on network packet processing. In contrast, grouping rules by types or tables, later applying jumps according to those types

or tables, can improve packet processing efficiency (although it can also hinder the expressiveness associated with firewall system configuration). All these decisions can be very relevant when it comes to protecting high-speed networks (gigabit and terabit networks).

2.1.1. Advantages and disadvantages of first generation firewalls

Building a firewall system using a router with packet filtering is really cheap, as it is usually done with hardware that is already available. In addition, it offers high performance in networks with a high traffic load. A packet filtering router example can be easily implemented from the routing systems of GNU/Linux-based operating systems, along with the associated Netfilter modules and system commands (iptables or nftables applications, depending on the version of Linux used).

Additionally, this technology allows the implementation of most of the necessary security policies.

Despite these advantages, network routers with packet filtering can have some shortcomings, such as:

- Many of the routers in use may be vulnerable to existing attacks (although most providers have appropriate update packages to address this). On the other hand, they usually do not have logging capabilities. This makes it difficult for the administrator to know if the router itself is being attacked.
- Its performance can deteriorate due to the use of excessively strict filtering and also make the device management process more difficult if this number of rules becomes very high.
- Filtering rules can become very complicated and sometimes cause possible distractions in their configuration to be exploited by an attacker to commit a security policy violation.

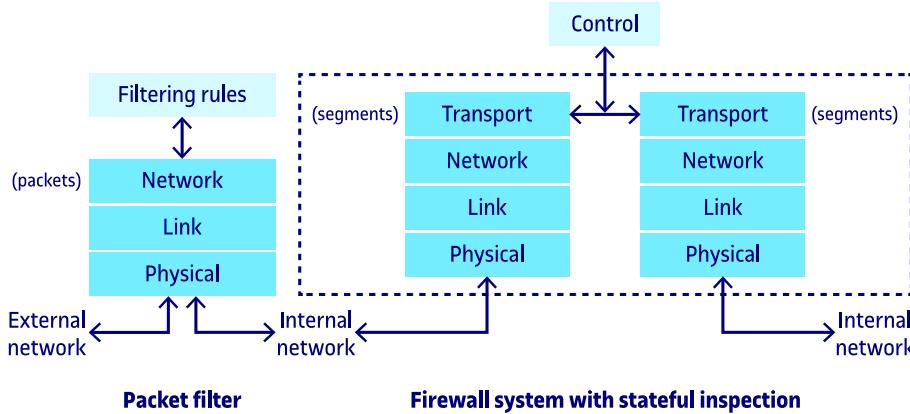
2.2. Second generation: stateful inspection at the transport level

Second generation firewall systems, also known as *gateways at the circuit level*, act as relays of traffic segments at the transport level. These devices filter content at the transport level (level 4) of the OSI layers reference model.

As we see in figure 4, these devices can supplement the packet processing done by a router with filtering rules, inspecting and deciding based on the states associated with TCP transport traffic segments, for example. That is, they can

hold network traffic until they get enough information about the end states of a communication and decide whether to allow or deny connections and associated data.

Figure 4. Second generation firewall, with stateful inspection, completing the packet filtering done by a first generation firewall



The **second generation firewall systems** are also known as *stateful packet inspection firewalls*.

Stateful packet inspection firewall systems work at level 4 of the OSI reference model, that is, on the transport layer between two ends. Therefore, they can monitor and decide based on the connection with each of these two ends and decide whether to relay the associated segments or not. Apart from all the information already available at the network level, they will also be able to use information at the transport level, such as special flags, associated with the state of a connection at the TCP transport level, for example.

Note that although first generation firewalls can use the port number (located at the transport layer) in packet filtering rules, the systems do not perform stateful inspection of the connection.

2.2.1. Stateful protocol filtering

As we already anticipated in the introduction of this section, second generation firewall systems allow stateful filtering using information about the state of the connections or sessions associated with the traffic.

This information can be used to add greater richness to the filtering rules and allow packets to be accepted or rejected based on their membership in specific sessions or specific states of a protocol with information at the transport level. Packet processing performed by the firewall system will need to keep track of the state of the transactions associated with the packets or the behaviour of the traffic passing through them.

Differences between stateful and stateless protocol filtering

In packet filtering associated with a stateless protocol, such as IP-type network-level traffic, each packet can be processed independently of other packets (i.e., regardless of the order or priority of packets). In contrast, when filtering traffic associated with a stateful protocol, such as traffic segments at the transport layer of the UDP or TCP type, the firewall system will require storage and monitoring of the sequence of packets, the flags associated with connections, etc. In this way, the firewall can make the filtering decision once the packets are placed in the order defined by the logic of the protocol. This implies greater complexity in tracking connection states, the origin and destination of packets, use of ports associated with applications, etc. So, in the case of stateless protocol filtering it is usually called *static filtering*, while in the case of stateful protocol filtering it is usually called *dynamic filtering*.

Thanks to dynamic filtering, by performing stateful and previous connection inspection, second generation firewalls can establish much more compact filtering rules and reduce the size of the filter rule set (in relation to the static filtering of the first generation firewalls).

Assume, for example, the rule “accept packets received in response to a previous request, originating on the internal network”, along with the UDP packets listed in table 2.

Table 2. Example of UDP packets for the dynamic filtering example

1	Origin IP Origin port	230.0.113.1 43321	Destination IP Destination port	192.0.2.1 7
2	Origin IP Origin port	192.0.2.1 7	Destination IP Destination port	230.0.113.1 43321
3	Origin IP Origin port	192.0.2.1 7	Destination IP Destination port	230.0.113.1 34511

In the above example, we are considering an internal network with the network address 230.0.113.0/24. A firewall system configured with the rule described above will accept packets 1 and 2, indicated in table 2, but will reject packet 3, since it is a packet that does not correspond to any request originating within the same network. The firewall will consider this third packet as a violation of the network security policy.

2.2.2. Advantages and disadvantages of second generation firewalls

The dynamic filtering performed by second generation firewalls provides certain advantages over the static filtering of the first generation firewalls.

To begin with, and as we have seen with the example of the previous subsection (in relation to the UDP packets represented in table 2), this new type of filtering allows reducing the set of rules to implement the same policy as a first generation firewall, as there is no need to anticipate response packets. In addition, the new rules can also be easily implemented from devices running GNU/Linux and the Netfilter framework (iptables and nftables applications, also depending on the Linux version used).

A second advantage of second generation firewall systems is the expansion of the perimeter security concept. As we will see later, in the section on implementing architectures, second generation firewalls are usually used to allow the implementation of demilitarized zones (DMZ). In this case, the firewall can use specific security protocols to route from a protected area to an unprotected area. A specific example is the use of second generation firewalls together with the SOCKS (SOCKet Secure) protocol. This protocol consists of a client and a server. The server can run inside the firewall, while the client can run on internal computers in the DMZ. The firewall will evaluate connection requests and decide based on a security policy whether the connection should be allowed or not. If the connection is to be allowed, the functionality of the SOCKS protocol will be used to relay the traffic and allow input or output at the transport level.

However, dynamic filtering is not perfect either. First of all, the rules defined in a firewall with dynamic filtering can be easily evaded, through IP spoofing attacks. Indeed, an attacker who succeeds in spoofing the IP address or the source port will be able to easily camouflage himself through the network and pass off his traffic as responses to previous requests, thereby fooling the connection tracker of the firewall. To avoid this, it will be necessary to think about more complex rules, which leads us to other drawbacks associated with this new type of filtering: the complexity of dynamic rules, the loss of efficiency in the treatment of traffic in congested networks, etc.

Complexity and loss of efficiency in dynamic filtering.

Dynamic filtering of stateful protocols increases the complexity of the rules, as well as a possible loss of efficiency. To be able to handle all possible cases associated with evasion or configuration errors, a firewall system with dynamic packet filtering will need additional memory resources to ensure connection and session tracking. If this is not possible, it will be difficult to guarantee correct processing of the state of the traffic flowing through the system. On the other hand, the increased complexity in packet processing, as well as the loss of efficiency can also open the possibility of suffering denial of service attacks. For example, adding dynamic rules to deal with potential IP spoofing attacks tends to lead to denial of service situations on firewall systems with limited storage and processing memory. A possible solution to these problems will be to add other defence devices within the system that needs to be protected. For example, the addition of attack and intrusion detection devices.

Use of the SOCKS protocol by second generation firewall

The SOCKS protocol, defined in RFC 1928, is considered a de facto standard for implementing second generation firewalls with circuit-level gateways.

The combination of dynamic filtering, together with the relaying of UDP or TCP traffic through clients and servers of the SOCKS protocol, allows the definition of DMZ-type architectures, hiding IP addresses of protected equipment from potentially hostile traffic from the outside.

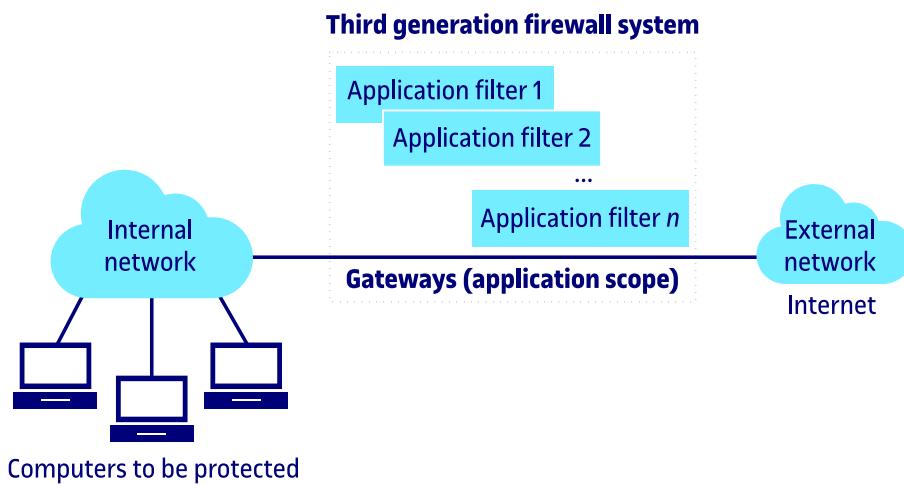
2.3. Third generation: application layer data processing

Third generation firewalls can enter to inspect (and modify) application layer information (level 7 of the OSI reference model). So they are built as application level gateways (proxy servers). In addition to being able to route packets at the network level or relay segments at the transport level, they act as gateways with access to application level data.

A proxy server is responsible for making the requested connections with the outside, and when it receives a response, it is responsible for retransmitting it to the equipment that initiated the connection. Thus, the proxy server running on the gateway applies the security policy to decide whether to accept or reject the connection request.

In fact, layer 7 firewalls typically inspect, modify, or eliminate traffic using specialized filters for a predetermined set of applications. Figure 5 shows this idea. Each filter can be considered as an intermediary processor, capable of modifying content at the application layer level. Packets corresponding to an application protocol known by the firewall are directed to the specific filter, which will be responsible for inspecting and processing the received data. This allows not only to filter contents, but also to alter part of the data. For example, a layer 7 firewall can perform HTTP (Hypertext Transfer Protocol, RFC 8740) traffic inspection. As a result of this inspection, and depending on the security policy rules associated with the firewall, it may reject the traffic, or allow it to pass, but making modifications to the associated data (for example, the filter may remove or rewrite JavaScript code found in HTTP data, rewrite headers or data based on known URLs, etc.)

Figure 5. Third generation firewall performing filtering or content modification at the application level, by selecting specific filters for each type of application



As also shown in figure 5, the gateway found within the third generation firewall system separates the internal network (with computers to be protected) from the external (potentially hostile) network. But, unlike first and second

generation firewall systems, this separation is done at the application layer level. This allows additional protection in the area of users or application data, in parallel with the security analysis associated with each filter installed on the firewall. These filters can be updated frequently, adjusting the analysis to modifications and service updates in the application layer, without the need to make changes to the lower layers of the firewall system.

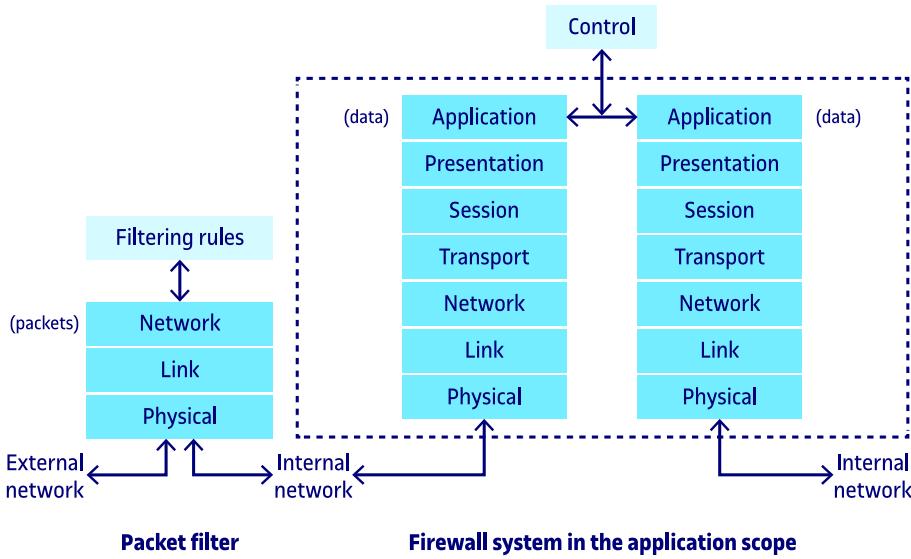
2.3.1. Advantages and disadvantages of third generation firewalls

Third generation firewalls offer advanced functionality with respect to packet filtering at the network level, or segment management at the transport level. They present a much higher range of possibilities. On the contrary, they also introduce a penalty to the traffic processing service, since they have to do a deep inspection of the data. In the case of congested networks, where the traffic load is high, the use of this kind of firewall can greatly penalize network performance and latencies.

A traditional way to solve this performance problem is to use cache systems that keep a local copy of the previous data received by the firewall, to be reused on subsequent connections (if this is possible). Despite this, the dynamism of current services, in which encrypted content changes continuously, also hinders the improvements provided by this cache-based solution.

A second way to deal with performance issues is to combine third generation firewalls with packet (first generation) or transport (second generation) pre-filtering. Thus, the simplest cases can be handled initially with traditional packet filtering, or stateful inspection if necessary, before moving on to application-level data inspection or modification. Figure 6 shows this last idea and illustrates the combination of packet filtering only, with application level analysis. This way of combining both systems, in addition to helping to reduce the loss of performance, also helps to provide flexibility at the configuration level.

Figure 6. Packet filtering and data gateway at the application level



The use of gateways within third generation firewall systems, with specific filters at the application level, provides other benefits. For example, it can enable effective identification as well as subsequent filtering of misuse of applications prohibited by the security policy. This last example is used to detect users trying to escape the ban of online gaming services or the use of P2P applications to download illegal material. Although this can also be done with traditional filtering at the network or transport level, this new identification can be based on the analysis of the data processed directly at the gateway level. That is, instead of making an identification according to ports or protocols indicated at lower levels, an identification of improper uses will be made based on the filters of the application level.

Another benefit of having the gateway at the application level is that the application protocol can also differentiate between specific situations when filtering. For example, the firewall can be configured to allow the use of P2P applications only for downloading operating systems or service updates, but continue to filter the download of illegal material (copyrighted films or music) within the same traffic flow. This feature would not be possible solely with network or transport level filters within first or second generation firewalls.

On the other hand, third generation firewall systems, despite offering more control over monitored services, still present certain drawbacks. A first drawback to highlight is the need to configure a filter for each service to be monitored. Creating filters for traditional services is not a problem, i.e., services with known flows for tracking sessions like TELNET, FTP, HTTP, etc. However, creating filters for little-known services or proprietary protocols without documentation can be a difficult problem to solve. The creation of these filters requires very detailed knowledge at the level of specifications. This tends to make it difficult to create specific filters for old protocols, for example, indus-

trial protocols, without standardized specifications. Email-related traffic can also lead to problems and produce spam filtering, or the specific removal of macros or executables within messages, incorrectly.

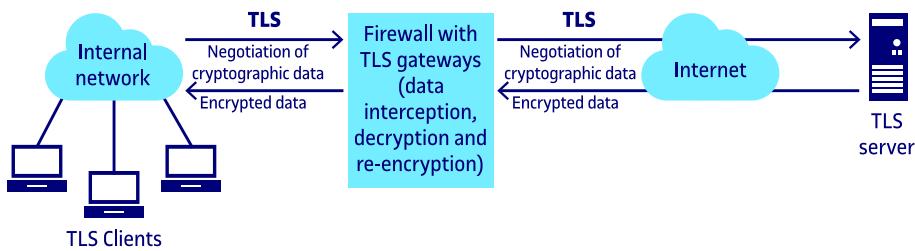
In fact, many of these drawbacks are also found in intruder detection systems, since the same concept of filters for the identification of improper uses and modification of known flows at the application level is the basis of these cyber defence tools.

A final drawback with analyzing data at the application layer is that this data is often encrypted because it is using transport layer security. The solutions to this kind of problem are varied and are beyond the scope of this material. Even so, as already mentioned briefly at the beginning of this section, it is one of the functions that are usually used as a commercial claim for Next Generation Firewalls (or NGFW).

These next generation firewalls incorporate new functionalities over gateways at the application layer of third generation firewalls, to perform (among others) threat detection, integrated network intrusion prevention, and malware filtering at system level and, in the case which concerns us here, encrypted traffic inspection.

This last functionality is usually used to carry out TLS inspection. Both TLS traffic inspection and other cases (such as SSH traffic inspection or similar) assume the incorporation of gateways at the application level, as well as the necessary filters to perform interception tasks, certificate negotiation, key and cryptographic data negotiation, online decryption, inspection and re-encryption. Figure 7 shows a simplified example of this technique, using a TLS gateway within the firewall system.

Figure 7. TLS traffic inspection by a next generation firewall.



Inspection of encrypted traffic by next generation firewall systems is a controversial practice. Apart from offering and using common techniques in interception attacks (man-in-the-middle), as well as possible violations of end-to-end encryption (responsible for protecting communications from origin to the final recipient), it also brings ethical problems when used to inspect corporate traffic, since the decryption of private data will not always be viewed favourably by the users of the organization, which leads to a possible intrusion into their privacy. Even so, it is a functionality increasingly demanded to solve the problem of evasion of hidden attacks in encrypted traffic (both to escape the filtering of a firewall system, and for the detection or prevention of intruders by detection systems).

Recommended readings

The following articles (both available online) provide more information on inspecting TLS-encrypted traffic and the potential consequences in the area of security or malpractice.

O'Neill *et al.* (2017). "TLS Inspection: How often and who cares". *IEEE Internet Computing*, IEEE Computer Society. <<https://doi.org/10.1109/MIC.2017.265102655>>

Durumeric *et al.* (2017). "The Security Impact of HTTPS Interception", NDSS Symposium. <<http://dx.doi.org/10.14722/ndss.2017.23456>>

3. Implementation of perimeter security through firewall systems

An effective implementation of the security policies associated with a system will largely depend on the architecture of the system to be protected. It is also very important to bear in mind the type of firewall systems that will implement this security policy, the location on the network or networks of the system, the combination with other equipment, etc.

Remember that, in general, the goal of implementing a security policy through firewall systems is to protect internal networks from the outside. Usually, this outside part tends to be the Internet. However, in other cases, firewall systems can be used to separate parts of the same internal network, such as the workstations of an industrial system or a test laboratory.

Numerous architectures and types of strategies exist in the firewall system literature. Although there is no sufficiently agreed-upon classification, we can simplify and divide the main filtering architectures or strategies into the following two types:

- single point architectures
- perimeter network architectures

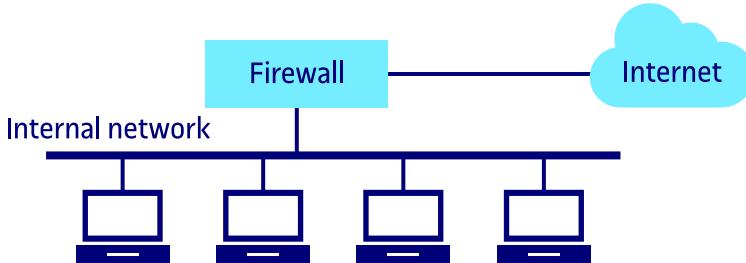
Next, we will present these two types of architecture in more detail.

3.1. Single point architectures

This first architecture is the simplest to implement. It consists of separating the network that we want to protect from the outside with a single protection device, as shown in figure 8. The device represents a single point of configuration. This makes the resulting filtering system simpler to implement and administer. At the same time, it also makes this single configuration item the critical point of the system. If an attacker manages to compromise any of the services behind this single configuration point, the computers associated with the protected system can be attacked without restriction.

The device labelled as a firewall in figure 8 can be based simply on a router with packet filtering (i.e., a first generation firewall system) or a device with more filtering capabilities, along with gateways inside to establish accepted communications (i.e., a second or third generation firewall system). From an architectural point of view, we will label this second case (second or third generation firewalls with gateways inside) with the concept of bastion host.

Figure 8. Single point architecture



A **bastion host** is a highly protected system prepared to withstand attacks from a hostile place (in this case, the Internet) and which usually acts as a point of communication between the inside and the outside of a network.

The bastion host is usually configured with two or more network interfaces, with the routing service disabled by default. Thus, traffic from one end of the network (the hostile side) will not be routed to the other side (the protected side) by default. Only if explicitly accepted, a gateway installed in the bastion host will be responsible for making the connections on behalf of these two parties. It will also allow the redirection of traffic to other networks, to carry out a more detailed analysis.

Bastion host

The name *bastion host* comes from the heavily protected walls that separated medieval castles from the outside.

However, the use of these bastion hosts with second and third generation firewalls including intermediate gateways leads to a loss of efficiency as the main drawback, especially if the traffic that travels through it is congested. To solve this problem, it will be more efficient to diversify and introduce multiple computers, combining routing and filtering at different levels (network, transport or application), as we will see in the next subsection.

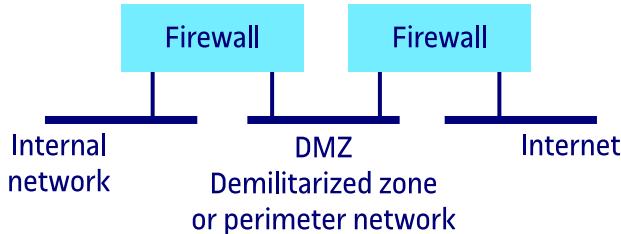
3.2. Architectures with perimeter networks

As already indicated above, regarding figure 8 of the previous subsection, we will differentiate with the firewall label the case of a first generation system, that is, a router with packet filtering, with respect to systems of second and third generation, which will be identified in the figures in this subsection with the bastion host label. With this proposed notation, and only from an architectural point of view, we will try to differentiate the use of filtering options with economical and efficient firewalls (with a traffic analysis at the network layer level), with more expensive schemes of filtering but offering more protection, assuming in addition the required gateway services to establish the accepted communications.

A first way to make it more flexible and add more security to the single-point architectures discussed earlier is the use of architectures with perimeter networks. In this case, we will add a subnet between the internal and external

network to act as a barrier against possible attacks and intrusions. This perimeter network (or the set of multiple perimeter networks) is also known as a demilitarized zone(s) (DMZ). Figure 9 shows a very simplified example.

Figure 9. Incorporation of a demilitarized zone or perimeter network



DMZ, perimeter networks and bastion host

The terms DMZ and perimeter network are used synonymously by many authors.

In general, a DMZ is considered to be the set of perimeter networks. In other words, a DMZ can consist of one or more perimeter networks.

Within each perimeter network, we will find bastion host (one or more), which usually refer to the devices that are continuously exposed to attacks.

Bastion hosts can provide public services of the organization (web, email, DNS, etc.), in addition to protection services (among others, traffic filtering at the transport or application level). That is why, in this section, we will use the concept of bastion host to refer to devices that provide second and third generation firewall services, to differentiate them from the input and output firewalls of a DMZ, generally implemented with first generation firewalls (i.e., routers with packet filtering at the network level).

Each perimeter network usually has one or more bastion hosts inside. These bastion hosts can offer, among other things, the filtering and gateway services discussed above. This will provide, in addition to a higher level of security, complete separation from the internal network. In addition, the bastion computers tend to give (or establish) access to internal network services that must be accessible from the outside.

If an attacker manages to bypass the security of the first firewall (or external firewall) and enter the perimeter network, he will not be able to immediately attack the internal network equipment, since they are protected by the second firewall (or internal firewall).

In figure 10 we can see in more detail a first perimeter network architecture. To simplify, we differentiate between two first generation firewall systems (outer and internal firewalls), as well as a bastion host (with a second or third generation firewall system inside). So let's assume that the external and internal firewalls are routers with packet filtering, and that the bastion host is configured with a minimum of two network interfaces, with routing disabled by default, but with gateways in charge of filtering tasks at the transport or application level, as well as the establishment of the necessary connections. Although the example in figure 10 shows only one bastion host, there could be more, as we will see later.

Figure 10. Simplified example of an architecture with perimeter network

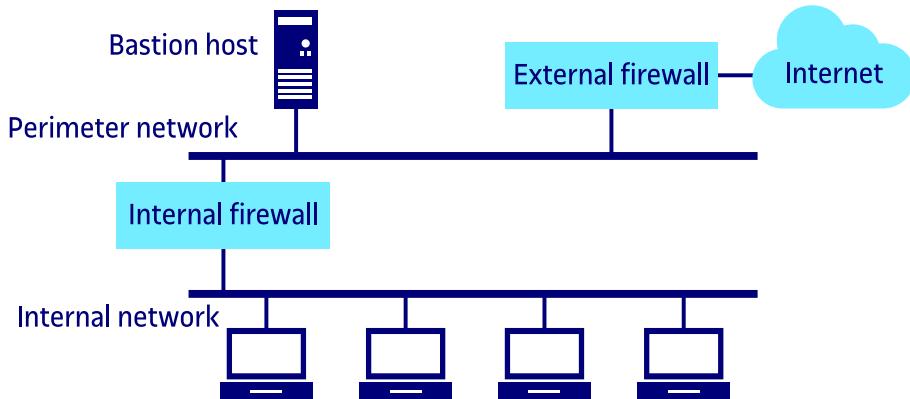


Figure based on the work of Zwicky, Cooper and Chapman (2000).

Note that the function of the two first generation firewalls shown in figure 10 is different. On the one hand, the firewall labelled as an internal firewall protects the internal network from the external network, but also from the perimeter network. We can see it as a router with packet filtering, to eliminate dangerous traffic (both input and output) to the internal network, from the outside world. That is, the internal firewall also controls the traffic between the internal network and the bastion host, thus ensuring that the traffic between the internal network equipment and the services of the bastion host is extremely limited, to prevent the commitment of the bastion host entailing a possibility of attacking the teams of the internal network. In addition, the bastion host can also offer traffic filtering at the transport and application level.

On the other hand, the external firewall protects both the internal network and the perimeter network. Again, we can see it as a router with packet filtering, but with less restrictive filtering rules. In fact, their rules will be specially designed to protect the bastion host from the outside. This external firewall may even be controlled by an external organization (for example, an Internet service provider).

Note that the initial example in figure 10 could also be configured more compactly, using a single packet filtering at the network layer, thus placing a single first generation firewall, routing and filtering tasks between the inner and outer part of the system, as we show in figure 11.

Figure 11. Example of architecture with a perimeter network with a single point of input and output.

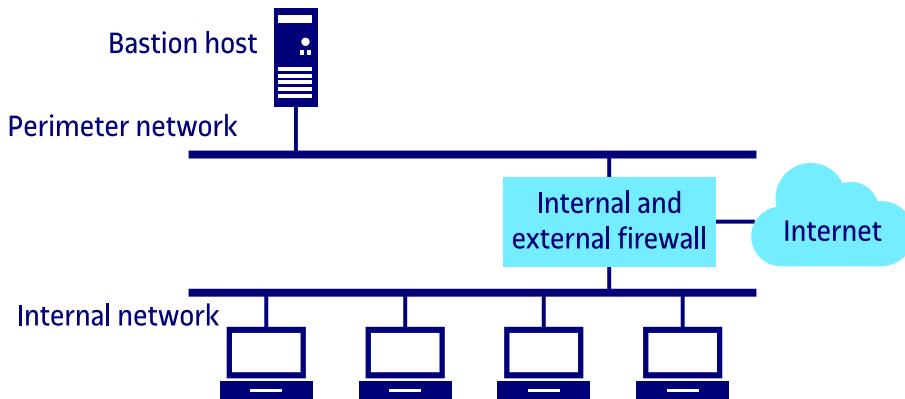


Figure based on the work of Zwicky, Cooper and Chapman (2000).

In any case, these two architectures with a single perimeter network allow us to see how different types of firewall systems can be combined. The main idea is to establish a first barrier of entry from a simple firewall, later combined with one or more system(s) configured with more advanced filtering services, together with the necessary gateways to establish the final communications, in case the traffic is accepted. These first two architectures shown in figures 10 and 11 aim to ensure a first level of protection, as well as addressing the performance issues associated with each type of filtering technology.

The architectures shown in figures 10 and 11 can be further generalized, for example, by splitting the initial perimeter network with other perimeter (sub)networks. Figure 12 shows a first example of this idea, increasing the number of perimeter networks and bastion equipment by one unit. A typical rationale for this architecture shown in figure 12 is to add redundancy and diversification to the services offered by the initial bastion host (not only the protection services at the level of filtering or data processing, but also the own services hosted on the bastion host for administrative tasks, for example). Again, this architecture will increase the security or efficiency levels of internal network equipment already shown in previous architectures, but with higher configuration and management costs.

Figure 12. Example of architecture with two perimeter networks.

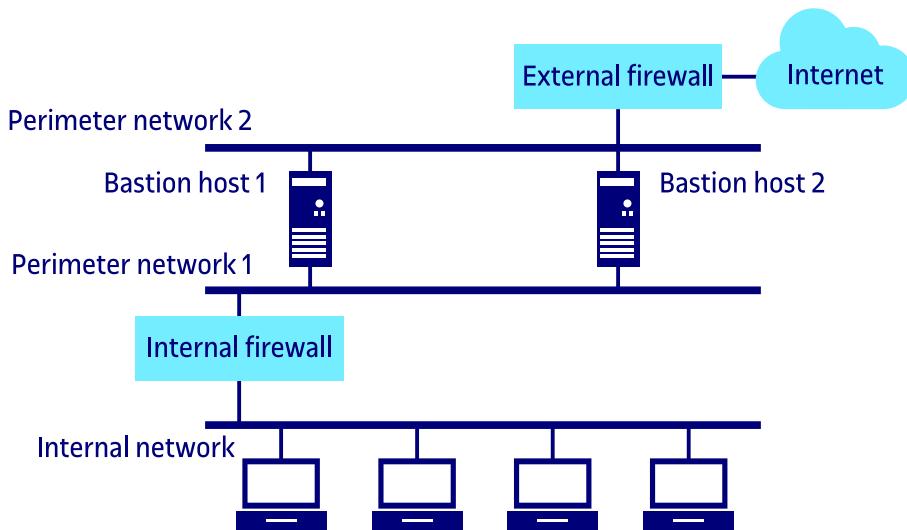


Figure based on the work of Zwicky, Cooper and Chapman (2000).

By evolving the previous architectures, we can reach much more complex situations, such as the architecture shown in figure 13. In this case, we find multiple perimeter networks further increasing the final redundancy of the system and also adding independent accesses to the Internet. This last configuration can be used as a measure of traffic separation between multiple perimeter networks, with different degrees of confidentiality; or to separate incoming traffic on the organization's servers from outgoing traffic of the organization itself. As in the previous cases, despite offering a much higher degree of protection, as well as improving the efficiency in the treatment of traffic, it will have as its main drawback a much more complex administration and configuration than the previous architectures. This can again lead to configuration errors and leave vulnerable spaces unprotected by mistake. It will also require a broader security treatment using other tools, such as monitoring and intrusion detection tools.

Figure 13. Example of a complex architecture with multiple perimeter networks

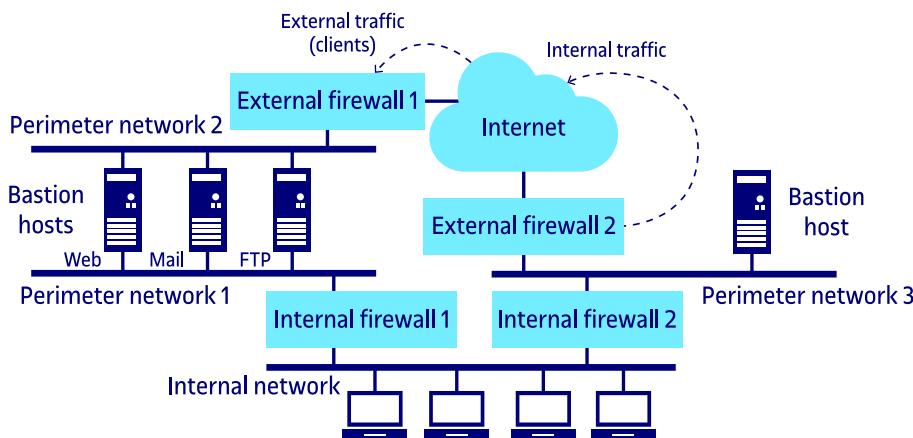


Figure based on the work of Zwicky, Cooper and Chapman (2000).

Summary

When a system is connected to a computer network, it is exposed to a set of threats that are always present. Furthermore, these systems are highly likely to have vulnerabilities which increase the likelihood of these threats taking place.

Firewall systems focus security decisions on a single point that is located where the greatest vulnerabilities exist and deny any connection that is not expressly allowed.

Through a packet filter configuration scenario in simple firewall systems, the decisions of a security policy defined by the organization can be technologically applied.

It is also possible to build firewall systems using proxy or gateway technologies, so that all received traffic can be interpreted at higher stack levels.

So, a firewall system is a control barrier that will keep the network protected from all unauthorized access and will act as a central point of control, simplifying the administration tasks.

On the other hand, due to the fact that they are located at a collision point, firewall systems offer other interesting security functions, such as the monitoring of network connections; content analysis (to search for viruses, for example); perform additional authentication checks; construction of virtual private networks; etc. They can also perform functions not directly related to network security, such as network address translation (NAT), network service management, bandwidth control, etc.

Finally, we must bear in mind that firewall systems are only prevention mechanisms and that they are not a single solution to solve all the security problems of a network connected to the Internet. These systems will never be able to protect the network from attacks that happen inside it and an external attacker may be helped by an internal (legitimate) user to collaborate in the attacks. Neither will they be able to prevent attacks against services with global access (in which anyone can access from anywhere) nor will they be able to protect the network against the transfer of malicious applications (viruses, worms, etc.). It would be impractical to use a device dedicated to analyzing all the traffic that circulates through it. This is why additional protection mechanisms are needed, such as intruder detection systems.

Glossary

Attack attack on the security of a system resulting from an intentional and deliberate act that violates the security policy of this system.

Bastion equipment see *Bastion host*.

Bastion host computer system that has been strongly protected to withstand attacks from a hostile location.

Demilitarized zone (DMZ) within a network protected by a firewall, an area separated from the public servers by a second firewall.

DMZ see *Demilitarized zone*.

DNS see *Domain Name System*.

Domain Name System hierarchical and distributed naming system that allows domain names to be associated with IP addresses.
acronym DNS

Dual-homed machine equipment with at least two network interfaces, each one associated with a network, that can act as a router between networks.

Firewall prevention element that will perform an access control to separate our network from outside (potentially hostile) equipment.

Gateway at circuit level device that acts as a gateway at the level of the transport layer between two ends. It establishes a connection with each one and relays data between the two connections.

ICMP see *Internet Control Message Protocol*.

Internet Control Message Protocol control protocol, mainly for sending TCP/IP error messages.
acronym ICMP

Internet Protocol protocol for interconnecting networks.
acronym IP

IP see *Internet Protocol*.

IP Address address that uses the IP protocol.

Perimeter security security based only on the integration of firewall systems and other prevention mechanisms into the network.

Proxy server software that will be responsible for making the requested connections with the outside and relaying them to the equipment that initiated the connection.

Router with packet filtering network device that routes TCP/IP traffic based on a series of filtering rules that decide which packets are routed through it and which ones are discarded.

Security policy set of rules and practices that define and regulate the security services of an organization or system with the purpose of protecting its critical and sensitive resources. In other words, it is a statement of what is permissible and what is not.

TCP see *Transmission Control Protocol*.

Threat potential violation of security based on circumstances, capabilities, actions or events that may cause a breach of security or cause damage to the system.

Transmission Control Protocol TCP/IP (end-to-end) transport protocol.
acronym TCP

UDP see *User Datagram Protocol*.

User Datagram Protocol TCP/IP (end-to-end) transport protocol.
acronym UDP

Bibliography

Avolio, F. (1999). "Firewalls and Internet Security, the Second Hundred (Internet) Years". *The Internet Protocol Journal* (vol. 2, no. 2).

Cheswick, W.; Bellovin, S.; Rubin, A. (2003). *Firewalls and Internet Security: Repelling the Wily Hacker* (2nd ed.). Boston, Massachusetts: Addison-Wesley Professional Computing.

Fraser, B. (1997). "Site Security Handbook". *RFC 2196, IETF. The Internet Society*.

Garcia-Alfaro, J. (2004). "Mecanismos de prevención". In: Herrera Joancomartí (coord.); Garcia-Alfaro; Perramón. *Seguridad en redes de computadores*. Barcelona: Fundació Universitat Oberta de Catalunya, 287 pp.

Microsoft (2010). *Windows Firewall with Advanced Security Getting Started Guide* (online). Microsoft TechNet Library.

Navarro Arribas, G. (2012). "Sistemes tallafoc". In: Herrera Joancomartí (coord.); Borrell, Garcia-Alfaro, Martínez, Navarro, Pérez, Perramón, Rifà, Robles. *Seguretat en xarxes*. Barcelona: Fundació Universitat Oberta de Catalunya.

Rash, M. (2007). *Linux Firewalls: Attack Detection and Response with iptables, psad, and fwsnort*. San Francisco, California: No Starch Press.

Shirey, R. (2000). "Internet Security Glossary". *RFC 2828, IETF. The Internet Society*.

Zwick, E.; Cooper, S.; Chapman, D. (2000). *Building Internet Firewalls* (2nd ed.). Sebastopol, California: O'Reilly Media.

Intrusion detection systems

PID_00293487

Joaquin Garcia-Alfaro

Recommended minimum reading time: 4 hours



Contents

Introduction	5
Objectives	6
1. Cyber defence and tools for intrusion detection	7
1.1. Background and current systems	8
1.1.1. Early solutions	10
1.1.2. Current intrusion detection systems	11
1.2. General architecture of a detection system	12
1.2.1. Event collectors	15
1.2.2. Event processors	17
1.2.3. Response units and remediation plans	21
2. Management of events, alerts and incidents	23
2.1. Configuration of event collectors	24
2.2. Information and collection policies	25
2.3. Normalisation of the collected information	26
2.4. Aggregation and fusion of information	27
2.5. Correlation of alerts and generation of reports	28
3. Detecting network intruders with Snort	32
3.1. Origin and architecture of Snort	33
3.1.1. Snort's basic architecture	34
3.2. Packet sniffer and preprocessor	36
3.3. Rules and detection engine	38
3.4. Notification system	40
4. Deception systems and techniques	42
Summary	45
Glossary	47
Bibliography	49

Introduction

The main consequence of the use of any digital technology is the exploitation of its weaknesses and, therefore, the execution of cyber attacks. Computer networks are no exception. Any network connected to the Internet is exposed to possible attacks. This is why it is very important to impose new security requirements beyond protection mechanisms (such as cryptography and firewall systems). Although deficiencies in these systems can be checked using conventional tools, they are not always corrected. In general, these weaknesses can put at risk network security and facilitate illegal entries into the system.

The area of cyber defence is precisely trying to solve this new problem with the use of mechanisms for monitoring and detecting attacks. Most current organisations have data protection mechanisms integrated into their networks. But, although these mechanisms must be considered essential, the security assumed by the organisation must carry on increasing.

Thus, a purely perimeter level of protection (based only on the integration of firewall systems and encryption systems into the network) is not enough. We must think that not all access to the network goes through the firewall or that all threats originate in the external area of the firewall. In addition, firewall systems, like other elements of the network, can also be attacked.

It is essential to install detection mechanisms, capable of alerting the network administrator at the time that these attacks take place, not only for the detection part, but also to remediate the attacks. An analogy that helps to understand the need to incorporate these elements could be the comparison between the security of a computer network and the security of a building: the entrance doors carry out a first level of access control, but we usually do not leave it at that; we will install motion detectors or surveillance cameras at key points in the building to detect the existence of unauthorised people, or people who may misuse resources and endanger security. In addition, there will be security guards, registration books in which it will be necessary to write down all the staff who access a certain department that we consider critical. All this information is processed from a security control office where the recording of the cameras is supervised and the registration books are kept. All these elements, projected into the digital world, shape what is known in the field of cybersecurity as *cyber defence elements*.

Objectives

The objectives you must achieve once you have worked through the materials of this module are the following:

1. To understand the origins of the concept of *intrusion detection* and see some of the classic techniques that can be used.
2. To learn to classify these systems according to various criteria, such as the place where the data analysis process is carried out or the specific detection mechanism.
3. To understand the limitations linked to the traditional processes of an intrusion detection system, including the problem of false positives and false negatives.
4. To know about the existence of additional management tools capable of normalising, merging and matching events and alerts collected in a distributed manner. These tools can consolidate the treatment of incidents detected through different event monitoring architectures.
5. To see a specific example of an intrusion detection system at the network level, using the Snort free open source tool.
6. To see complementary technologies to classical detection systems, with the specific example of deception systems and technologies.

1. Cyber defence and tools for intrusion detection

Cyber defence assumes that an attacker is able to violate the security policy associated with an information system, in order to misuse its resources or disrupt its operation.

The mechanisms associated with the **cyber defence** technologies are responsible for finding and reporting all kinds of types of malicious activity in the system or network, with the aim of reacting and stopping the attack in the most appropriate way.

In general, it is desirable to be able to identify the exact attack that is taking place in order to stop it and to recover from its effects.

In some particular situations, it will only be possible to detect and to report the suspicious activity that has been found, given the impossibility of really knowing what has happened and reporting how to stop it.

Generally, cyber defence mechanisms work under the premise of the *worst-case scenario*, assuming that an attacker has gained access to the system and is able to use or modify its resources.

The most notable elements within the category of cyber defence mechanisms are the Intrusion Detection Systems (**IDS**).

Origins of the term IDS

One of the first cyber defence communities comes from the world of security in database systems.

In these communities, rather than talking about attacks, they talk about intruders inside a database. Most of the references to the term *intruder* or *intrusion* are usually changed for *attacks* by other communities in the world of security (e.g., communities associated with the security of industrial systems).

A second community associated with the origins of the term IDS comes from the world of error detection and system failures. This community is interested in how to automate the detection of situations in which the different components of a system behave abnormally. In the case of an IDS, the aim is not to detect behaviours associated with random errors. On the contrary, the aim is to detect situations with marked intentionality. While errors have a probabilistic behaviour, attacks leave a trail associated with human behaviour, with malicious intent. For this reason, the term *intruder* is also used in these communities, to differentiate these two kinds of situations.

We will now introduce the main definitions that are used in the field of intrusion detection, with the aim of clarifying common terms that we will use later.

An **intrusion** is a sequence of actions taken by a malicious adversary with the ultimate goal of causing unauthorised access to a system or network of systems.

The intrusion is the sequence of steps taken by the intruder.

Each step taken represents a violation of the security policy of the system or the network. Therefore, the existence of a security policy indicating mandatory compliance actions, or with the malicious actions to be prevented, is a key requirement to detect an intrusion. In other words, a violation can only be detected when the actions observed can be compared to the set of rules defined in the security policy.

Intrusion detection is the process of identifying and responding to activities observed against the security policy of a system or network of systems.

This last definition introduces the notion of the intrusion detection process, which involves a whole series of technologies, users and tools needed to achieve success. Next, we will see these technologies in more detail, from their early origins to the current systems.

1.1. **Background and current systems**

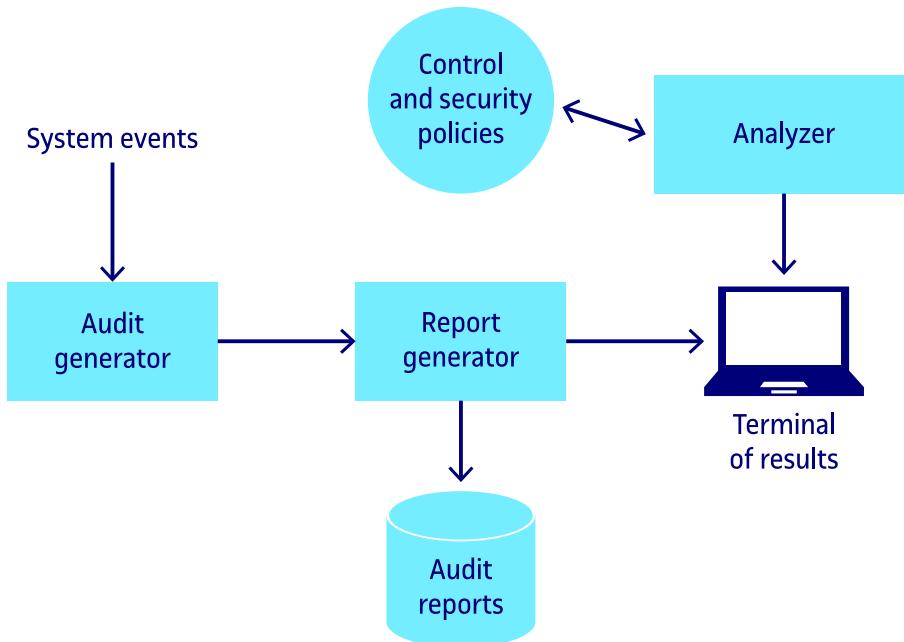
Current intrusion detection systems are a direct evolution of auditing systems. These systems were intended to measure the time that operators spent using the systems they monitored, with millisecond accuracy, and served, among other things, to bill for the service.

The first systems appeared in the 1950s, when the American company Bell Telephone System created a development group with the aim of analysing the use of computers in telephone companies. This team established the need to use audits through electronic data processing, unlike the previous system based on the preparation of paper reports.

This led the Bell System companies to embark on the first large-scale computer-controlled telephone billing system in the late 1950s.

Figure 1 shows a simple diagram of the operation of an audit system, in which the system events are captured by audit generators that will take the data to the element in charge of storing them in a report file.

Figure 1. Operation of a classic audit system



Trusted Computer System Evaluation Criteria (TCSEC)

Series of documents of the NSA (National Security Agency of the United States Department of Defense) on trust systems, also known as the *Rainbow series*, due to the colours of their covers. The main book in this series is known as *the Orange Book*. For more information, see the website: <http://www.fas.org/irp/nsa/rainbow.htm>.

Since the 1970s, the U.S. Department of Defense began to invest considerable resources for research into security policies, directives, and control guidelines. These efforts culminated in a security initiative in 1977, in which the concept of *trust systems* was defined.

Trust systems are systems that employ enough software and hardware resources to enable the simultaneous processing of a variety of confidential or classified information. These systems include different types of information distributed into levels, which correspond to their degree of confidentiality.

In the late 1970s, a section on audit mechanisms as a requirement for any trust system with a high level of security was included in the *Trusted Computer System Evaluation Criteria* (TSCSEC). This document, known as *the Tan Book*, lists the main objectives of an audit mechanism that we can summarise very briefly in the following points:

- To allow the review of access patterns (by an object or by a user) and the use of system protection mechanisms.

- To allow the discovery of both internal and external attempts to circumvent protective mechanisms.
- To allow the discovery of user transition when moving from regular to high level privileges (privilege elevation).
- To allow blocking attempts by users to bypass system protection mechanisms.
- To serve, in addition, as a guarantee to users that all the information collected about attacks and intrusions will be sufficient to control the possible damage caused to the system.

1.1.1. Early solutions

The *Intrusion Detection Expert System* (IDES) project, developed between 1984 and 1986 by Dorothy Denning and Peter Neumann, was one of the first real-time intrusion detection systems. This project, funded by the US Navy, proposed a correspondence between anomalous activity and abuse, or misuse (understanding strange or unusual activity in a statistical context as anomalous).

IDES used profiles to describe the subjects of the system (mainly users), and activity rules to define the actions that took place there (system events or CPU cycles). These elements allowed establishing, by statistical methods, the patterns of behaviour needed to detect possible anomalies.

A second real-time intrusion detection system to highlight was Discovery, which was able to detect and prevent security problems in databases. The novelty of the system was in the monitoring of applications rather than analysing an operating system in its entirety. By using statistical methods developed in COBOL, Discovery was able to identify possible abuses.

Other systems were also developed to help American officials find trademarks of internal attacks on the main computers of their airbases. These computers were mainly corporate servers that worked with unclassified but very confidential information.

One of the last important systems of this period was the Multics Intrusion Detection and Alerting System (MIDAS), created by the National Computer Security Center (NCSC). This detection system was implemented to monitor the NCSC Dockmaster system, where one of the most secure operating systems of the time¹ was executed. Like the IDES system, MIDAS used a hybrid system that combined both statistical detection of anomalies and the safety rules of

¹It is the Multics operating system, precursor of today's Unix systems.

an expert system. MIDAS used a progressive analysis process consisting of four levels of rules. In addition to these rules, it also had a database that it used for determining signs of abnormal behaviour.

MIDAS was one of the first detection systems for Internet-connected intrusions. It was published online in 1989, and it monitored the mainframe Dockmaster in 1990, to strengthen user authentication mechanisms.

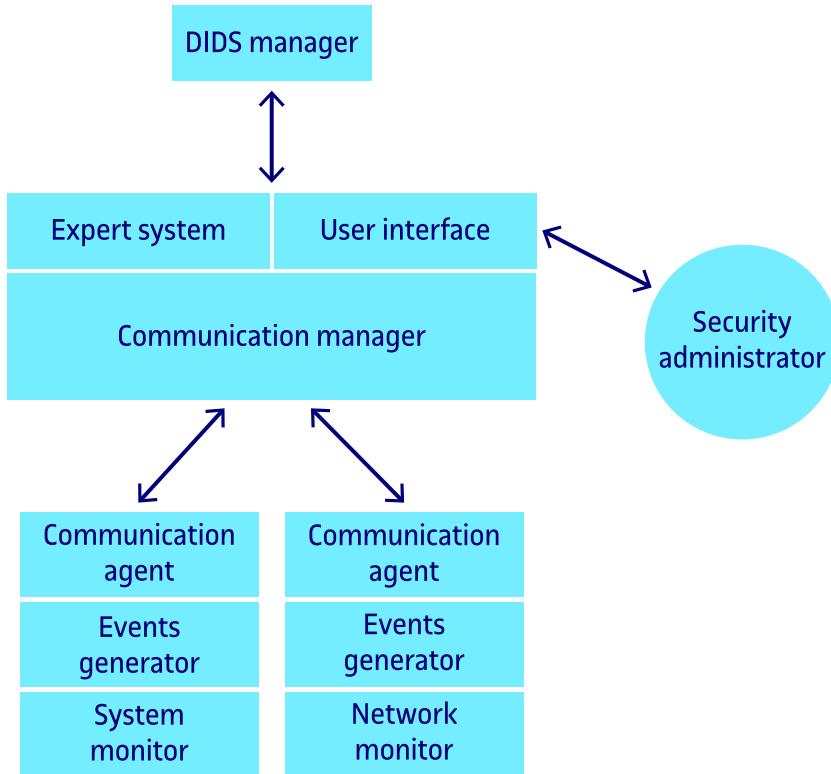
1.1.2. Current intrusion detection systems

Since 1990, the fast growth of computer networks led to the emergence of new intrusion detection models. As well as that, the damage caused by the famous Robert Morris worm (in 1988) helped to unite commercial and academic efforts in the search for security solutions in the field of cyber defence.

The first step was to merge the monitoring processes of the operating system with the monitoring systems of network traffic. This was the case of the Distributed Intrusion Detection System (DIDS), capable of enabling a security group to monitor security breaches through networks connected to the Internet.

The initial goal of DIDS was to provide means to centralise control and publication of results in a central controller. Figure 2 shows the operating diagram of the DIDS system.

Figure 2. Diagram of the operation of the Distributed Intrusion Detection System (DIDS)



At the same time, the first commercial intrusion detection programs began to appear. Some companies developed them to occupy a prominent position in the field of cyber security, although others did them to improve the levels of protection required by national security agencies, such as the NCSC.²

⁽²⁾It stands for National Computer Security Centre

Currently, a large number of intrusion detection systems to protect operating systems and complete networks are available. Many of these systems are commercial or reserved for military and research environments. There are also a large number of free solutions that can be used without any restrictions.

1.2. General architecture of a detection system

Since the beginning of the 1980s, many studies regarding the construction of intrusion detection systems have been carried out. In all these studies, different proposals and designs have been made in order to meet the following requirements (More information available at: <http://dx.doi.org/10.1007/BF02994844>):

- **Accuracy.** An intrusion detection system should not confuse legitimate actions with malicious actions at the time of detecting them. When legitimate actions are detected as malicious actions, the detection system may end up causing a denial of service against a legitimate user or system. These types of detections are known as *false positives*. The smaller the number of false positives, the more precision the intruder detection system will have.

- **Efficiency.** The detector must minimise the undetected malicious activity rate (known as *false negatives*). The lower the rate of false negatives, the greater the efficiency of the intrusion detection system. This is a complicated requirement, as it can sometimes become impossible to obtain all the necessary knowledge about past, current and future attacks.
- **Performance.** The performance offered by an intrusion detection system must be sufficient to manage real-time detection. Real-time detection must respond to intrusion detection before the attack causes damage to the system. According to studies, this time should be less than one minute.
- **Scalability.** As the network grows (both in size and speed), the number of events to be processed by the detection system will also increase. The detector must be able to withstand this increase in the number of events, without any loss of information. This requirement is of great relevance in distributed attack detection systems, in which events are launched on different equipment of the system and must be matched by the intruder detection system.
- **Fault tolerance.** The intrusion detection system must be able to continue to offer its service even if different elements of the system are attacked (including the situation where the system itself receives an attack or intrusion).

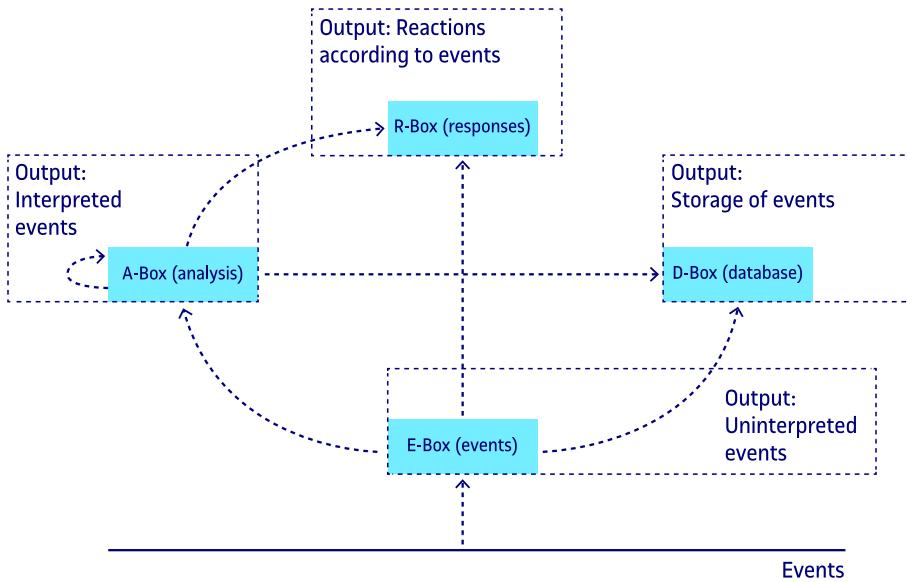
In addition, more recent criteria also include concepts such as knowledge *completeness*, implementation *facility*, maintenance *facility* and *explanation* of algorithms used for detection. These new criteria are considered essential for the future expansion of commercial IDS tools.

In order to normalise the heterogeneity of IDS architectures and tools, and achieve the aforementioned criteria, standardisation communities such as IETF³ have been working in recent decades on the conception of general-purpose IDS architectures. One of these, the architecture known as CIDF⁴ is shown in Figure 3.

(3) Stands for Internet Engineering Task Force

(4) Stands for Common Intrusion Detection Framework

Figure 3. CIDF architecture (The Common Intrusion Detection Framework)



Early IETF working groups, like the well-known IDWG,⁵ complement architectures such as the CIDF with the definition of development requirements, like the following:

- 1) The need to formalise the interaction between IDS tools with other security elements, such as prevention mechanisms (firewall systems, access control lists, etc.).

In this area, proposals for cooperation between the different security elements have been worked on through the exchange of messages. These mechanisms must ensure correct communication between the different elements of an IDS architecture and provide additional criteria such as the security of communications, authenticity and integrity of the information exchanged (alerts and events), etc.

- 2) The need to define the exchanged messages (events, alerts, etc.) among the system's elements. Therefore, the IETF has proposed formats for the exchange of messages between security devices, such as IDMEF⁶ and IODEF.⁷

By observing the above standardisation efforts, we can identify the following elements for the construction of an IDS-type general architecture:

- 1) event collectors
- 2) event processors
- 3) response units and remediation plans

⁽⁵⁾Stands for Intrusion Detection Working Group

⁽⁶⁾Stands for Intrusion Detection Message Exchange Format, RFC 4765

⁽⁷⁾Stands for Incident Object Description Exchange Format

Recommended link

The current version of the IDMEF and IODEF specifications is currently under review. We recommend seeing the www.secef.net website for more information.

We will now see some of these elements in more detail.

1.2.1. Event collectors

The event collectors, also known as **sensors**, are responsible for collecting information from equipment monitored by an IDS. The information collected is transformed into a sequence of events. These sequences will be analysed later by event processors.

The information stored in the events collected by the sensors is the decision base for the IDS detection. It is very important to be able to guarantee its integrity when facing potential modification attacks, when transmitting these events between the sensor that created them and the processing component that will treat them.

There are different ways to classify the possible implementations of these event collectors. The most common proposals are detailed next.

The first type of collectors, also known as **host-based sensors** or **system-based sensors**, are in charge of analysing and collecting information about events that have happened at operating system level (such as connection attempts and system calls).

These sensors are also known as **application-based sensors**, in the case of completing the information collected from the operating system, with events originated from applications.

In the second category we find sensors that collect information from events that have happened at the level of network traffic (for example, analysing the IP headers of datagrams that pass through the network interface). These types of components are known as **network-based sensors**.

The specific choice of the type of sensor will depend on the purpose of the desired detection. In fact, communities concerned about intrusion detection systems have thoroughly debated which sensors can offer better performance from a detection and effectiveness point of view. An ideal IDS should unify all possible options and offer a hybrid solution with the best of each option. Here are some additional details of each type:

1) Host-based sensors. These sensors are easy to configure, apart from being able to offer accurate information about events and potential attacks.

Snort, Suricata and Zeek

A widely used detection tool as a network-based sensor is Snort.

This tool is a network intrusion detector developed under the free software paradigm, capable of carrying out real-time traffic analysis, as well as logging packets in TCP/IP networks.

Two complementary tools in addition to Snort are Suricata and Zeek (formerly called Bro).

Check the websites www.snort.org, www.suricata.io and www.zeek.org for more information.

The data generated by these sensors can have a high density of information, such as the information reported by the file servers of an operating system log. They may also include a large amount of preprocessing information that later facilitates the work of other data processing components.

A disadvantage is that these sensors can have a significant impact on the efficiency of the system in which they are executed.

2) Network-based sensors. The main advantage of network-based sensors, compared to other solutions, is the possibility of being able to work in a non-intrusive way (i.e., in a passive manner). The collection of information does not affect the way the equipment works or the infrastructure itself. Since they do not necessarily reside in the equipment to be analysed, they are more resistant to attacks.

Moreover, most network-based sensors are also independent of the operating system and can obtain network-level information (such as the existence of fragmentation in IP datagrams) that could not be provided by host-based sensors.

Some network-based sensors are actually switches with an analysis capability that is transparent to the rest of the system.

As the main disadvantage of network-based sensors, it is worth noting the low scalability that this approach offers. In cases of networks with very high traffic loads, these sensors are likely to start losing packets, which means a loss in their ability to collect information.

These sensors would hardly be able to continue working normally on high-speed networks, such as Ethernet networks at speeds higher than a gigabit.

Another problem with network-based sensors is the increase in encrypted communications. The use of cryptography to protect communications makes the information to be collected incomprehensible to the sensors and thus reduces their detection capabilities. Some current solutions to deal with this problem include combining classical cryptography with homomorphic cryptography, to process part of the traffic even though it cannot decipher it.

Installation of event collectors

In addition to the sensor type, it is also important to determine the exact place where these components should be placed (from where to collect the information). The simplest to place are application-based sensors, generally

installed in the parts of the program in which debugging and generation of log files services are offered. But the situation is more difficult for the other variants.

When considering the installation of system-based sensors, the wide variety of operating systems there are, and the different facilities offered by each, poses a serious problem. In addition, it is not easy to determine what part of the large amount of information created by the kernel of an operating system should be relevant when analysing.

In the case of Unix-like operating systems, there is the *Orange Book* proposal (which we have already mentioned in this module), which shows some points of interest in which information should be analysed.

In the case of network-level sensors, segmentation through switches is a major drawback in choosing the right location to place these sensors. For example, a star topology causes packets to be directed only between the two parts of a communication, so the sensor should be placed at a point where it could analyse any exchange of information.

A first option would be to place the sensor on the link where all the equipment on the network joins. This option could mean the need to analyse such a high amount of data that the sensor would end up losing information.

The other option would be the placement of the sensor between the network link that separates the interior and the exterior, as if it were an additional perimeter protection system.

A variant of these two options would be the use of the tap port that many switches offer. It is a special port that reflects all the traffic that passes through the equipment. But this can easily overload the subsequent analysis capacity if the amount of traffic is very high. In addition, the internal bandwidth of the device is not always enough to deal with all active ports at once. If the traffic analysed begins to grow, the capacity of the tap port may be exceeded, with the corresponding loss of packets that this would entail.

1.2.2. Event processors

Event processors, also known as **analysers**, form the core of the detection system. They have the responsibility to operate on the information collected by the sensors in order to infer possible intrusions.

In order to infer intrusions, the analysers implement a specific detection scheme. Two of the most commonly used methods for detection are the misuse detection model and the anomaly detection model. We will now briefly comment on these two detection schemes.

Misuse-based detection

The detection of intrusions based on the misuse model has prior knowledge of malicious sequences and activities. Event processors that implement this scheme analyse events for known attack patterns or activity that attacks typical vulnerabilities.

The sequences or patterns described are known as *attack signatures* and could be compared to the virus signatures used by current antivirus programs.

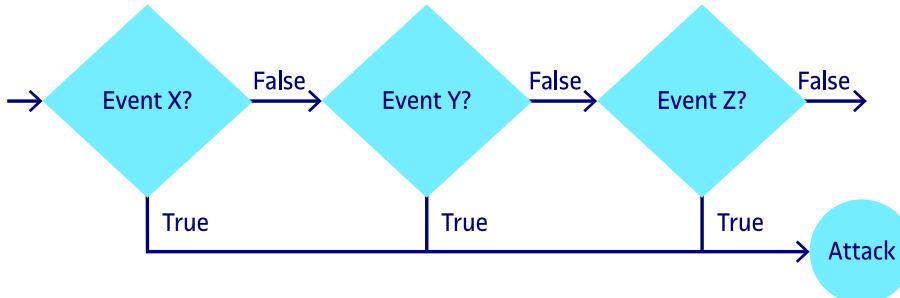
Thus, the detection components based on the misuse model will compare the events sent by the sensors with the attack signatures that they keep stored in their knowledge bases.

When an event or sequence of events matches an attack signature, the analyser will issue an alert.

When implementing a detection scheme based on misuse, two of the models most commonly used are analysers based on pattern recognition and analysers based on state transitions:

1) Pattern recognition analysers. By using *if-then-else* rules to examine data, these analysers process the information through internal functions in the system, in a completely transparent way to the user. Figure 4 shows the diagram of an *if-then-else* rule.

Figure 4. Example of an *if-then-else* rule

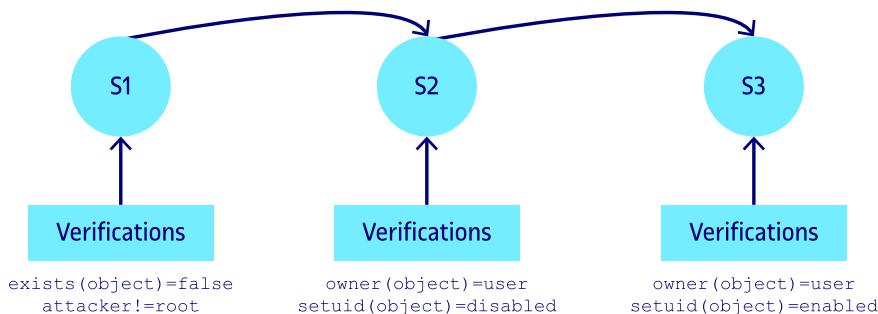


Although this model allows us to detect an intrusion from already known patterns, the main disadvantage is that these patterns do not define a sequential order of actions.

Detecting attacks composed of a sequence of events using this method can lead to great difficulties. The maintenance and updating of the pattern database are other critical points of this model.

2) State transitions. This model makes use of finite automata to represent the attacks, in which the nodes represent the states, and the arrows (arcs), the transitions. The use of transition diagrams (Figure 5) facilitates the association between states and the different steps taken by an intruder from the time it enters a system, with limited privileges, until it gains control.

Figure 5. Example of a state transition diagram



As the main advantages of this model, we can highlight that transition diagrams allow a high-level representation of penetration scenarios, and offer a way to identify a series of sequences that form an attack.

As well as this, these diagrams define very simply the attacks to be detected. The analysis engine could use different variants of the same diagram to identify similar attacks.

By contrast, transition diagrams, and therefore different steps in the sequence, must be created using specific languages that are often very limited and insufficient to recreate complex situations.

This limitation means that this model cannot detect some of the most common attacks, and that the use of additional analysis engines is necessary as a complement to this model. For example, engines with Bayesian inference or first-order logic, also used by deductive databases.

Anomaly-based detection

Event processors that base their detection on an anomaly scheme will try to identify suspicious activities by comparing the behaviour of a user, process or service, with profile behaviour classified as normal. A profile serves as a metric (measurement of a set of variables) of normal behaviours. Any deviation that exceeds a certain threshold with respect to the stored profile will be treated as evidence of an intrusion.

Disadvantages of anomaly-based detection

The drawbacks of the anomaly-based detection model make most of the commercial detection systems available today generally implement schemes based on the misuse model.

One of the requirements of this model is the need to initialise a default profile that gradually adapts to the behaviour of a normal (non suspicious) user, process or service. It is therefore necessary, to use heuristics and statistical descriptors that help to correctly model changes in behaviour as soon as they happen. Other proposals seek to incorporate artificial intelligence techniques to carry out these tasks (for example, use of machine learning with neural networks).

Anomaly-based detection has clear advantages over misuse-based detection. A first advantage is the possibility of detecting unknown situations (for example, *zero-day* attacks). This is possible since, regardless of the source of the intrusion, as soon as malicious activities start to deviate from the normal behaviour, the event processor will launch an alert.

Anomaly-based detection also has disadvantages. One of the first disadvantages is the lack of guarantee in the detection process. Stealthy intruders can slowly perpetrate malicious actions to evade detection.

As a second drawback, we can highlight the difficulty in explaining and accurately describing the attacks detected by anomaly-based analysers. Generally, an analyser not only has to trigger an alert, but also must specify where the attack comes from, what changes the system has undergone, etc.

Finally, the high rate of false positives and false negatives that can occur using this detection scheme is a serious drawback, since a deviation from the expected profile will not always coincide with an intrusion attempt. In the case of processors in which the events come from network-based sensors, it is possible that the number of triggered alerts easily get unmanageable. This often causes network administrators to ignore alerts issued by the detection system, or even deactivate the system altogether.

All of these drawbacks make most of the commercial detection systems available today implement their analysers by using the misuse detection model.

Stored data

In most situations, the volume of events collected by the IDS, including the alerts triggered by the analysers, becomes so high that a storage process outside the detector is necessary. Let us suppose, for example, that all packets in a high-speed network must be inspected by the analysers of the detection system. In this case, it is necessary to consider an external storage hierarchy that reduces the volume of information without penalising the possibilities of analysis.

One possibility is the classification of information in terms of short and medium term analysis. In the case of a **short term** analysis, the information is stored directly in the same components of the IDS (in internal *buffers*), so that after processing the data, and transforming it into an event format, it is trans-

Complementary reading

The following article (available online) provides more information about possible improvements needed in the world of anomaly-based detection, to change the current commercial situation: Seng, S; Garcia-Alfaro, J; Laarouchi, Y. (2021). "Why anomaly-based intrusion detection systems have not yet conquered the industrial market?", *14th International Symposium on Foundations and Practice of Security*, Springer Nature <https://doi.org/10.1007/978-3-031-08147-7_23>.

mitted to other processing elements. In the case of **medium term** information, the pre-processed data is stored on secondary devices (with the appropriate format) rather than being transmitted to the internal system processors.

The storage time for medium-term information can be two or three days, with the aim of being consulted by the system processors if the analysis process requires it.

Eventually, and after a compression process (to reduce the size), some of the medium-term information may continue to be stored for long periods of time (around months or even years) waiting to be consulted by **the long-term** detection processes.

1.2.3. Response units and remediation plans

The response units are responsible for initiating or providing decision aid information, in relation to possible repair plans that respond to a detected intrusion. Very rarely will response actions be automatic (**active response**). Most of them simply provide information and require human interaction (**passive response**), in order to assist with additional information before activating the final response by a security expert (or the administrator himself).

Responses and **remediation plans** aim to act against an intrusion, trying to neutralise the tasks already performed by the intruder before reaching the final objective. An active response, right at the moment when the intruder is detected, is usually not recommended. In general, response units will attempt to simulate and provide decision aid information, while the intrusion is still ongoing. Thus, administrators will be able to try to anticipate and take actions that neutralise the objectives of the intruder. An example could be to suggest new configuration rules to the firewall system associated with the affected components, with the aim of blocking future network connections such as the one that the intruder originated. Another example of a remediation plan could be the execution of additional tools to track the next phases or attacks associated with the intruder, to be able to perform the appropriate analysis later.

In most cases, we will have rather passive responses, which are limited to launching an external alert, to inform and describe the attack detected to the system administrators.

Most response components offer different ways to send this information to administrators, such as email, SMS messages, WhatsApp, GLPI tickets, etc.

In the case of automating the remediation plan, without first going through the system administrators, there could be collateral damage. For example, an active response (without human intervention) could lead to a denial of service against legitimate users or systems. It is very likely that some of the alerts thrown by the processors may handle are incorrect (**problem of false positives**).

For this reason, if a response unit immediately cut off the connection that originated this false positive, it could imply the loss of work of a legitimate user or service.

In most systems (e.g., e-commerce servers) such errors can lead to customer loss, which is inadmissible. Hence, most companies in the electronic commerce sector opt for hiring specialists who manually analyse the reports generated by the detection system in order to determine whether an active response to this guidance is necessary or not.

Finally, it should be noted that, like sensors, response units can also be classified into different categories according to the point of action. Both categories are again **host-based response units** (in charge of acting at the operating system level, for example, by blocking users, process completion, etc.) and **network-based response units** (in charge of acting at the network level, for example, dropping out connection attempts, applying network address filtering, etc.).

Problem of false positives and false negatives

A false positive occurs in those situations where an IDS characterises legitimate traffic, which is not part of any attack, as malicious; it is, therefore, an event detected by mistake. Contrarily, a false negative takes place in those situations in which malicious traffic is discarded and considered legitimate traffic by mistake; it is, therefore, an event that should be detected, but which escapes the detection process.

2. Management of events, alerts and incidents

In the previous section we introduced intrusion detection systems as isolated solutions, with components implemented in general architectures that allow the detection of elementary actions. These actions usually correspond to previous stages of an intrusion or the final objectives of an intrusion, that is, actions coordinated with complex objectives, beyond mere data theft. We are talking about coordinated actions that can seek long-term attacks, such as disrupting industrial or financial processes, with periods that go beyond months and even years. In order to anticipate these situations and prepare remediation plans accordingly, additional elements must be used to complete and conclude the detection process made by the components of the traditional IDS that we have seen in the previous section.

In this section, we will deal with these additional elements, which we summarise with the concept of *NMS* and *SIEM*⁸ type of platforms. Although there are many differences between both concepts (the NMS are more focused on error diagnosis, and SIEM are more focused on the detection and activation of remediation plans), in this section, we will refer only to SIEM platforms, understanding it as a combination of both concepts.

⁽⁸⁾NMS stands for *Network Management System* and SIEM stands for *Security Information & Event Management*

Coordinated actions and distributed attacks

Situations that cannot be identified by looking for patterns in isolation, but must be detected from the combination of multiple indications found at different points in a system.

Recommended links

Some products (with commercial and free software versions), related to the NMS and SIEM concept, can be consulted at the following links: www.nagios.org (NMS), www.vigilo-nms.org (NMS), www.prelude-siem.org (SIEM), <https://cybersecurity.att.com/products/ossim> (SIEM).

SIEM platforms are absolutely essential nowadays. Although they have existed for more than three decades, the expansion of the Internet in corporations and institutions has led to the recent developments and improvements that we summarise below:

- **Collection and normalisation of events.** A SIEM must be able to manage the collection of events from extremely heterogeneous sources. Therefore, a process of data normalisation must be carried out, not only at the syntactic level, but also at the semantic level. At the syntactic level, the existence of formats and standards, such as the IDMEF⁹ and IODEF¹⁰ formats can help in the tasks of syntactic normalisation of alerts from heterogeneous collectors or detection sensors. As for semantic normalisation, the use of ontologies, modelling languages such as the UML (*Unified Modeling Language*), or any other mechanism for the formal representation of alerts, can help in this matter.

⁽⁹⁾Stands for *The Intrusion Detection Message Exchange Format*.

⁽¹⁰⁾Stands for *Incident Object Description and Exchange Format*.

- **Consolidation of the monitoring functions of the detection tools.** The second vital function attributed to a SIEM is the consolidation of low-level alerts produced by the security components of a network, among which firewall systems, IDS, antivirus and vulnerability detection systems can be highlighted. As we have seen in the case of IDS, network security tools are vulnerable to problems related to false positives and false negatives. For this reason, what is expected from the use of a SIEM is the management of the processes of fusion, aggregation and correlation of alerts from previous hosts, and to reduce the rate of false positives. In addition, they will improve the diagnosis of errors to also reduce false negatives.
- **Preparation of remediation plans.** In general, most IDSs are usually configured as purely passive mechanisms. However, solutions available nowadays provide the technology needed to turn them into active or semi-active solutions (waiting for confirmation from an operator before activating the reaction process), with the aim of being able to react and neutralise the detected activities or actions. As we have seen in the previous section, this functionality must be analysed and activated with caution. The normalisation and correlation tasks of a SIEM should improve this point and reduce the potential collateral damage of remediation plans.

We will now detail some of the required tasks in order to carry out the functionalities that are expected of a SIEM.

2.1. Configuration of event collectors

In addition to information from an IDS, the data collected by a SIEM may come from any other component with the capability of creating log files, such as:

- **Firewall systems** (or any other type of component for packet filtering through network-level access control lists). Although we have often seen these components as prevention mechanisms, which aim to block traffic considered dangerous to the system, these components can be configured to report the information observed through their network interfaces. Their audit files are therefore of great interest to complement the process of aggregation and correlation of events of the SIEM.
- **Email servers** (based on protocols such as SMTP, POP or IMAP). Again, the audit files generated by these servers will offer a wealth of information to characterise and discover malicious activities, such as the spread of worms, *zero-day* attacks, or to control traffic associated with other illegal activities.
- **Traffic management servers** (based, for example, on the SNMP protocol). The information reported by the components will also contain informa-

tion associated with actions related to the violation of the internal security policies of the organisation where they are installed.

Another thing to keep in mind during the configuration of the SIEM detection sensors will be the data, events and alerts provided by an IPS (Intrusion Prevention System). An IPS is primarily based on research into the vulnerabilities of equipment and systems. Most IPS-type solutions actually combine intrusion detection techniques with traditional access control mechanisms. The border between IDS and IPS is currently difficult to define and can be complemented with many other preventive solutions, such as vulnerability detection systems and antivirus systems.

Vulnerability detection systems try to analyse the configuration of systems deployed in network, with the aim of discovering poorly configured parts that potentially present vulnerabilities that could be the target of attacks. These systems also include detection of software defects (i.e., programming errors or *bugs*), conception errors in the topological configuration of a network, hardware errors, etc. However, antivirus systems are designed to protect workstations and servers against known malware. Most of these systems will use a database of antivirus signatures that identifies the known malicious software. From the prevention point of view, it is expected that both a vulnerability detection system and an antivirus will be able to correct vulnerabilities and disinfect the equipment that are malware victims. Thus, the installation and combination of an IDS with other prevention mechanisms has the ultimate aim of detecting and reacting to the general concept of *intrusion*.

2.2. Information and collection policies

Since the purpose of a SIEM is to be able to provide system operators with centralised management capabilities, it is common for the configuration of the collection components associated with the SIEM (IDS and the components of previous examples) to be done through the SIEM user interface. That is why a global information collection policy is usually used, administered by the SIEM, and which can then be refined for the local configuration of each of the equipment associated with the SIEM (IDS, firewall, mail servers, etc.).

This policy, or its associated sub-policies, is usually defined by rules, based on the use of regular expressions, the search for traffic patterns, signalling protocols, etc. The components may also incorporate the possibility of dealing with a collection of events through configurations and based on the recognition of anomalous activities. Consequently, the policy must offer the required syntax and semantics to be able to deal with the use of statistical models, collected through data mining, expert systems, Bayesian networks, etc.

SIEM's event collection and log information policy may also involve executing a set of actions and preparing the data before passing them to the next stage. For example, it may involve data preprocessing from the execution of regular expressions, with the aim of conducting a search for specific patterns (search for signs of a particular virus, or for a specific port scan action with specific system components), as well as replacing specific parts of the events to improve the identification of a particular component or to enrich the data (for example, adding computer names as a complement to their IP addresses).

Some other technical criteria that will need to be considered in a SIEM's event collection policy will be to determine precisely the location of the collection components that must be controlled in the system. This aspect will determine the coverage and the global vision modules that the system will incorporate to ensure that the subsequent process of aggregation and correlation of information guarantees quality management. Therefore, it will also be necessary to specify, within the collection policies, the physical topology and the logical description of the system that the SIEM must monitor. It is important to be able to treat and structure this system in terms of subnetworks, so that the parts of the system that have a greater need for surveillance can be identified, and thus the correlation functions can recognise the tasks associated with specific incidents for each subnetwork of the system. For example, if the system contains DMZ-type areas (demilitarized zones), it should be possible to create specific rules to closely monitor the actions being taken on this part of the network. The use of subnetwork masks (which is specified in terms of CIDR¹¹ or interdomain routing without classes, for example) is common in the log file collection policies of most SIEMs.

¹¹Stands for *Classless Inter-Domain Routing*.

2.3. Normalisation of the collected information

Once the information is collected, and before moving on to subsequent processes for detecting attack scenarios, the SIEM must ensure that it is possible to put into correspondence all the data coming from the detection of the same event (such as malicious, suspicious or anomalous events). Thus, there must be a set of normalisation processes to preprocess the collected data and anticipate problems that could hinder the correspondence of data related to events derived from the same network traffic flow (such as the origin of the traffic, the destination, the ports, etc.) and that correspond to the supervision tasks of different monitoring sensors installed in different places of the system (either deployed in the same domain or in different domains).

The normalisation process must ensure, for example, that data associated with network traffic have the same format with regard to the classification of the origin of traffic and also metadata associated with the time at which the event was detected, associated protocols and services, origin and destination addresses, content of the packets associated with traffic, etc. Since the nature of the sensors is potentially heterogeneous, the normalisation process must

guarantee, either with SIEM proprietary formats, or with efforts from existing standards, that there are no interoperability problems that limit the expressiveness of the information that the SIEM correlation module will have to add.

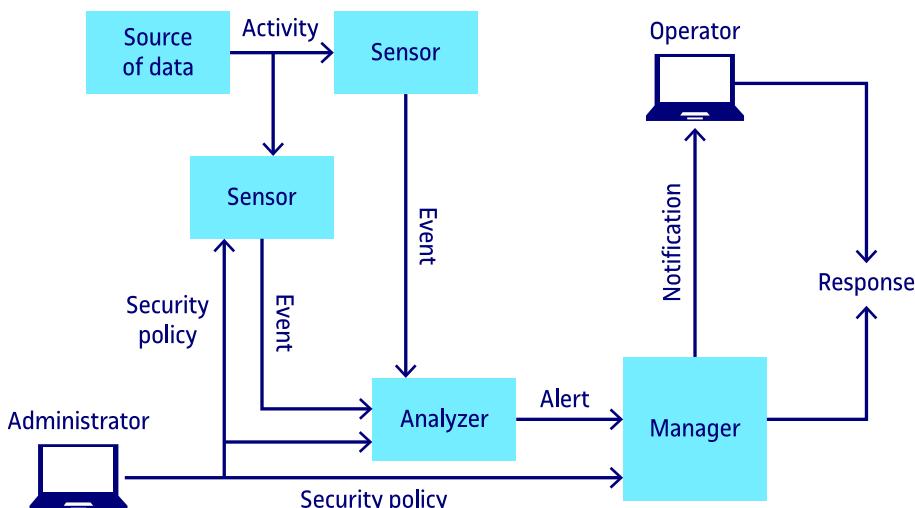
Examples of normalisation efforts

Some examples of efforts regarding the problem of normalisation are the following:

- CIDF (Common Intrusion Detection Framework)
- IDWG (Intrusion Detection Working Group)
- IDMEF (Intrusion Detection Message Exchange Format)
- INCH (Extended Incident Handling)
- FINE (Format for Incident Report Exchange)
- IDIP (Intrusion Detection and Isolation Protocol)
- OSEC (Open Security Evaluation Criteria)
- IODEF (Incident Object Description and Exchange Format)

Each of the above examples has tried to define common languages to specify the description of events and the activities associated with the events that security components must exchange. Apart from languages (both syntactic and semantic) to guarantee homogeneity in data exchange, it is also a priority to allow the description of a common process that specifies the precise protocol for the exchange of data between the different sensors configured in a SIEM. Of the examples listed above, two of the formats and procedures that have the support of working groups of the Internet Engineering Task Force (IETF) are the *IDMEF*¹², RFC 4765 and the *IODEF*¹³ (summarised in Figure 6).

Figure 6. Components associated with efforts such as IDMEF and IODEF



⁽¹²⁾Stands for *Intrusion Detection Message Exchange Format*.

⁽¹³⁾Stands for *Incident Object Description Exchange Format*.

Link of interest

The last version of the IDMEF and IODEF specifications is now being revised. For more information, we recommend seeing the <http://www.secef.net/> website.

Complementary reading

You will find more information about the model associated to IDMEF in the following book:

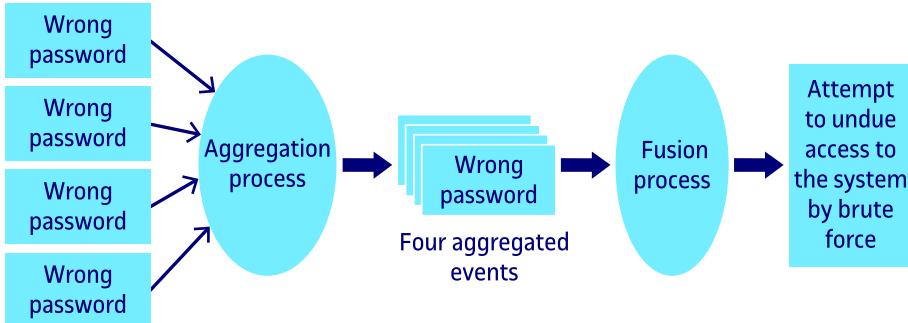
Stallings (2017). *Cryptography and Network Security - Principles and Practice* (7th edition), Pearson Education, Inc., Hoboken, NJ. All rights reserved. See chapter 22 for more information.

2.4. Aggregation and fusion of information

Data aggregation and fusion functions are used to intelligently reduce large volumes of data that are likely to contain redundant events (repetitions and congruent information). Both functions must be applied before matching events detected by different system components. First, the aggregation process will have to be responsible for grouping the data resulting from the detection of the same event, reported by one same sensor or by different ones. Once the grouping of this information is done, the process of merging the information

will take place, with the aim of summarising and offering a single datum that characterises the detected event. Figure 7 shows the difference between the aggregation and fusion functions with a simple example.

Figure 7. The difference between aggregation and fusion



During the aggregation process, the information corresponding to detected events will be grouped into sessions and will attempt to unify the data of the same event so that they can be used later during the fusion process, such as the source address, the destination address, the ports, the protocols, etc. In this way, the different data associated with an attack on a specific element of the system will be grouped into a single session and a single identifier. The rest of the data reported by other system sensors will be linked to the same reference, so that during the final fusion process it will be possible to generate an identifier alert. This alert will contain all the information observed by the different sensors configured by the SIEM. The alerts generated from the fusion process will therefore contain a synthesis of all SIEM knowledge about each one of the basic attacks observed by the monitoring system. As a result, the amount of data needed to be stored in the system is reduced without any loss of information. Once all the information reported by the system has been merged, the alerts will be communicated to the last SIEM component, responsible for managing and matching (correlating) the flow of alerts.

2.5. Correlation of alerts and generation of reports

In general, we can define the *alert correlation process* as the conceptual interpretation of multiple alerts with the aim of providing a semantic improvement and reducing the overall amount of alerts in an intrusion detection system.

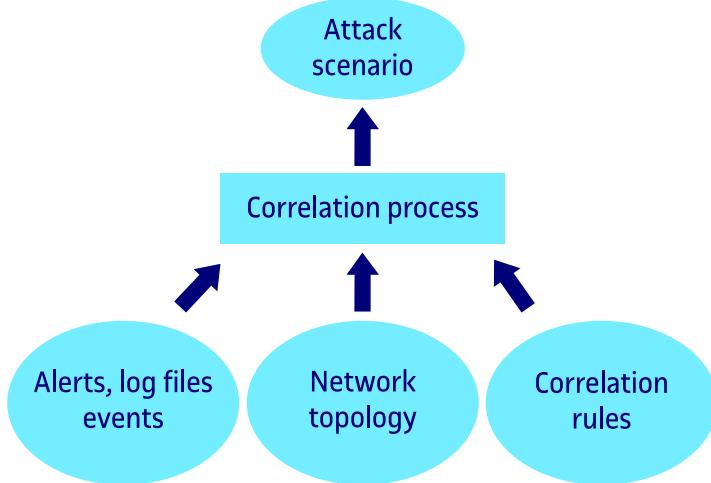
The correlation of alerts is therefore considered one of the keys in the evolution of intrusion detection systems, as it tries to solve the most relevant drawbacks of these systems (i.e., excessive alerts), improve them semantically and reduce false positives and negatives. The studies related to the correlation of

alerts in the field of intrusion detection are relatively recent. Most of them deal with observations and experiments made in current systems. The theoretical part in this field is still in the process of evolution.

Most commercial SIEMs still do not have real functionalities for conducting a full alert correlation. Although most of the solutions that exist talk about correlation services, the vast majority are limited to storing logical links between alerts stored in the same relational database that the responsible administrator can then consult with a control console. In contrast, within the field of academic research on intrusion detection systems, there are a large number of proposals for the inclusion of alert correlation techniques in new generation systems.

Most of these proposals essentially exploit the information contained within the alerts and, in addition, try to make explicit references to annexed knowledge, necessary for the final process of matching all the events observed in the same system. Therefore, the correlation in these systems is not limited only to the information contained within the alerts generated by SIEM probes, but they also make use of a prior knowledge of the surveillance state of the systems, on the attacks that can be made, and even on the topology of the system and the correlation rules generated by the operators and interpreted by expert systems. Figure 8 summarises these types of correlation systems.

Figure 8. Correlation function



As we can see in Figure 8, the first entry in the correlation process corresponds to the events already aggregated and fused by the lower SIEM modules. The objective of the correlation process is precisely to match these groupings of events and reconstruct attack scenarios to which the observed actions could belong. To do this, it will be necessary to combine the received alerts with the physical and logical information about the system (for example, the topology of the system and the structuring of the IP addressing of the equipment) and the predefined knowledge of malicious activities. The latter is represented in Figure 8 in the form of correlation rules.

The physical and logical properties of the network must be specified in a topological database, either manually (by system administrators), or through the use of automatic discovery and assisted creation services of network topologies. The latter must be responsible for the creation of topological maps with the configurations of the devices and the security policies associated with the protected information system.

The awareness of malicious activities must be configured in the system by means of correlation rules. Each correlation rule will be defined by characterising sets of actions corresponding to the same intrusion scenario, and will specify the required conditions to carry out an action and the consequences on the system after the execution of each action. As with the detection rules of an IDS based on misuse, these correlation rules will be the basic element to ensure the detection of attack scenarios formed by the actions detected in the system. Therefore, to ensure the effectiveness of the correlation process, it will be of paramount importance to ensure a correct configuration of the relevant correlation rules for each system, as well as its correct updating and commissioning. The elimination or corruption of a single rule containing information about multiple incidents may result in the loss of detection of a large number of scenarios. The configuration of this part of the SIEM will therefore be very prone to errors and requires a deep knowledge by the operators in charge of configuring the system.

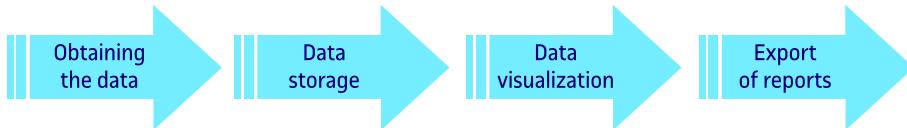
The detection of an incident or intrusion scenario is derived, during the correlation process, from the series of events indicated by the set of correlation rules preconfigured in the system. From a small set of correlation rules, it is possible to define many intrusion scenarios. Thus, the aim of the correlation process is to reduce the excess of information that the system operator will have to manage. Instead of asking the administrator to analyse thousands of events, the correlation process provides the generation of incident reports. Each report will contain the representation of scenarios that, with a high probability, could have been developed in the system from the events (i.e., primitive actions) detected by the SIEM probes.

With the same aim, the correlation process can also be configured to reduce the number of false alerts that are analysed (false positives). To do this, the system can be configured to execute internal verification actions triggered after the generation of each intrusion scenario and verify the certainty that this incident has occurred in the system or not. In this way, incidents that can be discarded with a high probability will be eliminated from the final report that the operator must prepare. This internal verification may be limited, for example, in the execution of an analysis of vulnerabilities of the system that serves to decide if a specific incident can be discarded, if the associated vulnerabilities are not present in the system represented in the system's topological database; that is, an incident involving the exploitation of vulnerabilities or services not deployed in the system may be discarded with a high probability and thus, avoid overloading the system operator's analytical capabilities.

Additionally, these incidents may be reported as false positives rather than alerts or intrusion scenarios. All this will facilitate the ordering of the response mechanisms in the system and also the optimisation of the remediation actions that will have to be deployed in the system on the real incidents that have been detected.

Finally, as a graphic complement to the alert correlation, most SIEMs offer graphical interfaces (of the *dashboard* type) to manage the generation of reports and provide command control to the end user. Figure 9 summarises this process, concluding the presentation of the typical components of a SIEM.

Figure 9. Stages associated with the generation of SIEM reports



3. Detecting network intruders with Snort

Snort is a very complete open source security tool for creating intrusion detection systems in network environments. It is very popular among the community of network and service administrators. Thanks to its ability to capture and record packets in TCP/IP networks, Snort can be used to implement from a simple packet sniffer for traffic monitoring of a small network to a complete intrusion detection system in real time.

As a network monitor, Snort behaves like a true vacuum cleaner (hence its name) of IP datagrams, which offers different possibilities in terms of their treatment: from acting as a simple passive network monitor that is responsible for detecting malicious traffic that circulates through the network to the possibility of sending all the captured traffic to log file servers or database servers.

But, apart from excellent features such as packet detector and alert generator, Snort has many other features that have allowed it to become one of the most complete software solutions for the construction of detection systems in network environments based on pattern recognition. Snort is defined as a *Lightweight* Network Intrusion Detection System sensor. This *lightweight* qualification means that, as an IDS, its design and implementation allow it to operate under different operating systems and that its functions as a detection mechanism may be part of different (even commercial) security products. The latest versions of Snort also define themselves as an IPS, that is, an Intrusion Prevention System, since they offer the possibility of blocking traffic, in active response unit mode, as a complement to the detection of intrusion attempts. However, in this section we will discuss Snort only as a collector of network events, that is, as a network sensor for the detection of intruders, since this is its initial, and best known, function.

In fact, Snort's popularity has increased in recent years in parallel with the increasing popularity of open source operating systems, such as the Linux and BSD family of systems. However, its nature as an open source product does not limit it to being available only for such operating systems. Snort can run on commercial solutions, such as Microsoft Windows.

From the point of view of its detection engine, Snort would be part of the detection category based on misuse. Using signature recognition, Snort will match all captured traffic against its detection rules.

A Snort detection rule is nothing more than a set of requirements that will allow it to activate an alert, if they are met. An example would be a Snort rule that would verify the use of P2P applications for file sharing over the Internet, while verifying the use of the GET chain in services other than the traditional HTTP protocol port. If a packet captured by Snort matches this simple rule, its notification system will issue an alert to indicate the facts. Once the alert is released, it can be stored in different ways and with different formats, such as a single system log file, an entry into an alert database, a SNMP event, etc.

We will now see the origins of Snort and an analysis of its architecture and some of its most remarkable features.

3.1. Origin and architecture of Snort

Very briefly, we can define Snort as a packet sniffer with additional functionalities for packet logging, alert generation and a misuse-based detection engine.

Snort was developed in 1998 under the name APE. Its developer, Marty Roesch, was trying to implement a multiplatform packet detector (although the initial development was made for the GNU/Linux operating system) that had different options for classifying and displaying the captured packets. Marty implemented Snort as an application based on the libcap library (for the development of packet capture), which ensured great portability in both the capture and in the collected traffic format.

Snort began to be distributed through the *Packet Storm* website (<http://www.packetstormsecurity.com>) on the 22nd of December, 1998, with only one thousand six hundred lines of code and a total of two source files. At that time, the main use given to it by its author was an analyser of his network connections with a cable modem and as a debugger of the network applications he was implementing.

The first signature analyser developed for Snort (also known as a rule analyser by the Snort development community) was added as a new application functionality in January 1999. This new functionality allowed Snort to start being used as an intrusion detector.

In December 1999, Snort version 1.5 was released. In this version, its author decided on a new architecture based on plugins, which is still preserved in current versions. After this version, Marty Roesch left the company where he worked and began to dedicate himself full-time to the task of adding new functionalities to improve the configuration capabilities and facilitate the use of Snort in more professional environments. Thanks to Snort's great acceptance among the community of administrators, Marty thought it was a good

Commercial version of Snort

Although Snort is available under the GPLv2+ (GNU Public Licence, version 2) licence, it also offers alternative licences and commercial products. Most alternative products are based directly on Snort and distributed by the Cisco company.

Cisco acquired the rights to Snort in 2013, along with the Sourcefire company, founded by the creator of Snort, Marty Roesch. The analysis of commercial versions of Snort is beyond the scope of this didactic module. For more information, you can visit <http://cisco.com> and search for information about these products.

More than just a detector

The author of Snort was trying to indicate with the name Snort that it was more than just a detector. The term *snort* means inhaling obsessively and violently. In addition, Marty said at the time that he already had too many applications called *a.out* and that all popular names for detectors, called *TCP-something*, were already occupied.

time to offer his product with a support for companies and obtained the necessary funding to establish Sourcefire. Cisco acquired the Sourcefire rights in July 2013.

However, Snort remains free code and promises to do so forever. Version 3 of Snort features more than 100,000 lines of code and involves a complete restructuring of the original design of its initial architecture. The free version of Snort, under the GNU licence, can be freely downloaded from Snort's website (<https://snort.org>) and allows any user to have support for the latest available versions and the latest updates to the rule files for those versions.

Currently, Snort has a large repertoire of accessories that allow reporting notifications to different database managers and a large number of traffic preprocessors that allow analysing remote calls and port scanning before these are contrasted with the set of rules associated with the search for new incidents.

Snort's rule sets have also evolved as the application has grown. The size of the rules sets for the latest Snort version available for download is increased in a similar way to the speed of new releases of exploits. These rule files are currently classified into different categories such as P2P, denial-of-service attacks, attacks against web services, viruses, pornographic traffic, etc.

Each of these rules is associated with a unique identifier (Sensor ID, SID) that allows recognising and finding information about the detected attack or misuse. For example, the SID for the SSH *banner attack* is 1838. In addition, thanks to Snort's extensive use among the network administrator community, other intruder detection tools have adopted the Snort rules format and also the coding used for the dumping of the captured packets (based on *libcap*).

Support for these rule files is increasing every day. In this way, any user of Snort, or any other network detection tool with a compatible rule format, could create its own rules as new attacks appear and collaborate with the Snort development community to keep its signature base perfectly updated.

3.1.1. Snort's basic architecture

Snort provides a set of features that make it a very powerful security tool, among which are the capture of network traffic, the analysis and registration of captured packets and the detection of malicious or dishonest traffic. Before seeing Snort's outstanding features in more detail, it is important to know and understand its architecture.

Snort consists of a set of components, most of which are developed as plugins that allow the customisation of Snort. Among these components, the preprocessors stand out, which allow Snort to manipulate the content of the

Recommended links

Version 3 of Snort is available at <https://github.com/snort3/snort3>.

The main differences between version 2 and version 3 of Snort can be quickly consulted on the following website: <https://blog.snort.org/2020/08/snort-3-2-differences.html>.

packets more efficiently before moving it to the detection element; its notification system is also outstanding, which allow the reported information to be sent and stored in different formats and following different methods.

Snort plugins

The term *plugin* refers to the software modules of an application, developed independently of the general kernel of the application. The aim is to add additional features without affecting the core source code or other components. To do this, the application must provide an API (*Application Programming Interface*), that allows the development and compilation of such complements. Thus, a Snort plugin is a component developed according to the Snort plugin API, which will be used alongside the core of the Snort code, but separated, so that a change in the component code does not affect the kernel or other components.

Snort's central architecture is based on the following four components:

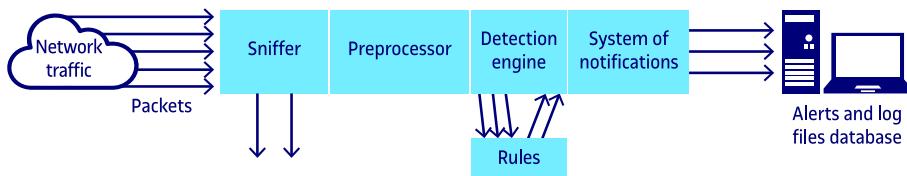
- Packet detector
- Preprocessor
- Detection engine
- System of notifications

According to this structure, Snort will allow the capture and preprocessing of network traffic through the first two components (packet detector and preprocessor) and will then check them through the detection engine (according to the set of activated rules) and will generate, through the last of the components, the relevant notifications.

Figure 10 shows Snort's basic architecture that we have just commented on. Looking at the figure, we can draw an analogy between Snort and a mechanical machine for automatic coin sorting:

- 1) It takes all the coins (network packets collected by the packet detector).
- 2) Each coin will be dropped down a ramp to determine which group of coins it belongs to (packet preprocessor).
- 3) It sorts the coins according to each coin type and packs them into cannons according to the category (detection engine).
- 4) Finally, the administrator will decide what to do with each one of the ordered coin cannons (notification system).

Figure 10. Snort's basic architecture



Snort's preprocessor, detection engine and notification system are also implemented in the form of independent components. We will now examine in more detail each of the basic components of Snort that we have just seen.

3.2. Packet sniffer and preprocessor

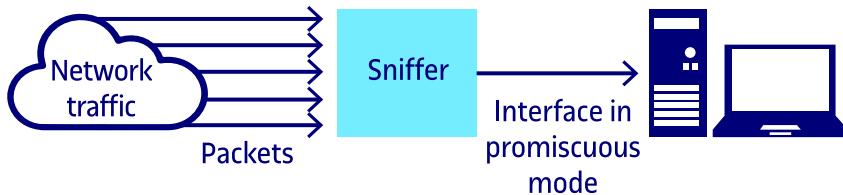
A sniffer is a device (software or hardware) that is used to capture packets travelling through the network to which it is associated.

In the case of TCP/IP networks, this traffic is usually IP datagram traffic, although it is also possible to have traffic of different types, such as UPnP traffic or (Apple) Bonjour traffic. In addition, since IP traffic also consists of different types of protocols, such as TCP, UDP, ICMP, routing protocols, IPSec, etc., many sniffers will need to know beforehand the type of traffic to be able to later interpret the packets that are being collected, and display them in a language that a network administrator can understand.

Like many other tools related to network security, sniffers can be used for more or less dishonest purposes. Among the different uses that can be given to a sniffer, we can think of traffic analysis for the solution of congestion and network problems, improvement and study of the performance of resources, passive capture of sensitive information (passwords, user names, etc.).

Thus, like the rest of traditional detectors, the Snort packet decoder will be the element in charge of picking up the packets, which the other components will examine and classify later. To do this, the packet sniffer must be able to capture all the traffic that it can, to then pass it to the next component (the preprocessor), which will be responsible for detecting what type of traffic has been collected. Figure 11 shows a diagram of how the Snort packet sniffer works.

Figure 11. How the Snort packet sniffer works


Network interface in promiscuous mode

The promiscuous mode is the way in which the network card of a computer connected to a network (whether a wired network or a wireless network) will allow the capture of all the traffic that flows through this network interface.

As the Snort packet sniffer collects the traffic passing through the network, it will deliver it to the preprocessing element. This element will adapt the captured packets and deliver them to the detection engine.

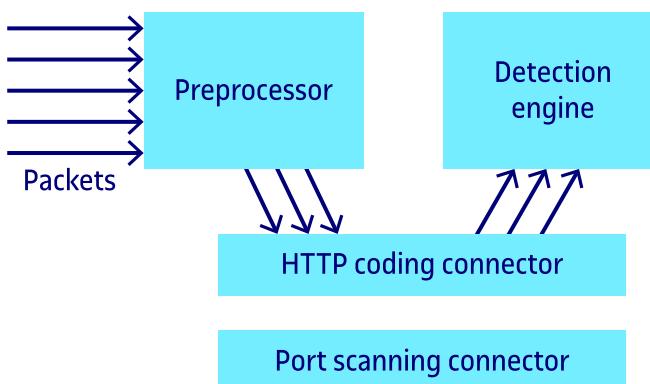
Thus, the preprocessor will get raw packets and will verify them through a set of plugins. For example, it will use a plugin for the treatment of packets related to RPC¹⁴ type traffic, or a plugin for the treatment of packets related to port scans. These plugins will verify packets in search of certain behaviours that allow Snort to determine their type. Once the type is determined, the packet will be sent to the detection engine.

⁽¹⁴⁾Stands for *Remote Process Calls*.

This preprocessing feature is really important for a detection tool, as it is possible to use third-party applications that can be activated and disabled according to the needs of the preprocessing level. For example, if a network administrator is not concerned about the RPC traffic that enters and leaves their network (and therefore does not need to analyse it), for whatever reason, it will be sufficient to disable the RPC plugin and continue using the rest.

Figure 12 shows a diagram in which the Snort preprocessor uses two of its plugins to verify the type of packets it receives and decide whether to move them to the detection engine or not.

Figure 12. Snort preprocessor that combines two data processing plugins



Snort preprocessors offer great flexibility for the implementation of different traffic processing algorithms.

3.3. Rules and detection engine

As we have already mentioned, Snort bases its detection on the misuse model. For this reason, Snort must be configured using a set of rules that will use the detection module to perform the recognition of attacks and intrusion signatures. Snort rules are usually grouped into signature sets that categorise incidents. Thus, we will find sets of rules associated with the detection of Trojans, the detection of intermediate buffer overflow attacks, and so on.

Each rule can be divided into two parts. First, we have the header of the rule, in which we indicate the action associated with this rule in case it is fulfilled (generation of a log file or generation of an alert); the type of packet (TCP, UDP, ICMP, etc.); the address of origin and destination of the packet, etc. Second, we have the option field of the rule, in which we will find the information that must be contained in the packet (in the data part, for example) so that the action associated with the rule is activated.

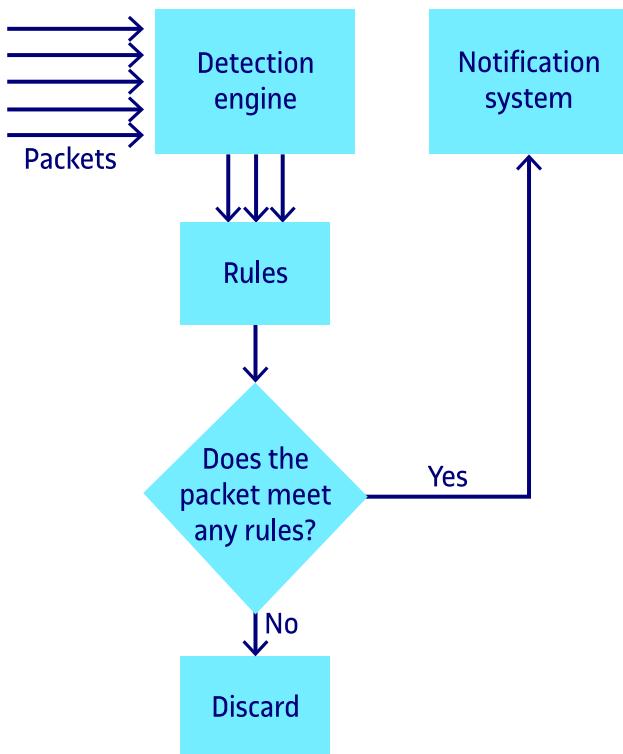
The action defined in a Snort rule can be chosen from five basic actions:

- **Alert.** Generation of an alert that also contains the log information corresponding to the packet that activates the rule.
- **Log.** Generation of only the log information associated with the content of the packet associated with the activation of the rule.
- **Pass.** Pass the packet associated with the rule without having to log or alert about the event.
- **Activate.** Generation of an alert together with the activation of a dynamic rule (see Dynamic action).
- **Dynamic.** Rule that remains inactive and is activated through the activate action, for example, to trigger the generation of logs associated with the packets associated with the initial rule to obtain additional information regarding traffic subsequent to the packet that triggered the rule (and with a given time duration).

Using Snort in IPS mode can provide additional actions, such as reactive actions to block detected traffic. In this case, the use of additional actions (of *drop*, *sdrop* and *rejected* types, as in firewall systems), allows to directly reject the packets identified by Snort as suspicious.

In Snort's detection engine we find the required treatment algorithms to complete the detection process. Based on the information provided by the pre-processor and its associated plugins, the detection engine will contrast this data with the rules base defined by the operator. If any of the rules matches the information obtained, the detection engine will be responsible for warning the notification system, indicating the rule that it has skipped. Figure 13 shows a simple outline of the overall behaviour of Snort's detection engine.

Figure 13. Snort's detection engine



Of all the elements we have seen, the detection engine and the syntax used by the detection rules are the most complicated and difficult parts to understand when studying Snort's behaviour. However, once we start working with Snort and have minimally learned the syntax used, it is quite easy to get to customise and adjust the behaviour of Snort's detection functionality. In addition, the rule sets can be easily activated or deactivated, so that the desired detection behaviour can be defined according to the type of network in which Snort will be configured.

3.4. Notification system

Once the information captured by the Snort packet decoder is analysed by the detection engine, the results must be reported in some way. Through this component, this function can be performed and the results can be generated in different formats and to different equipment.

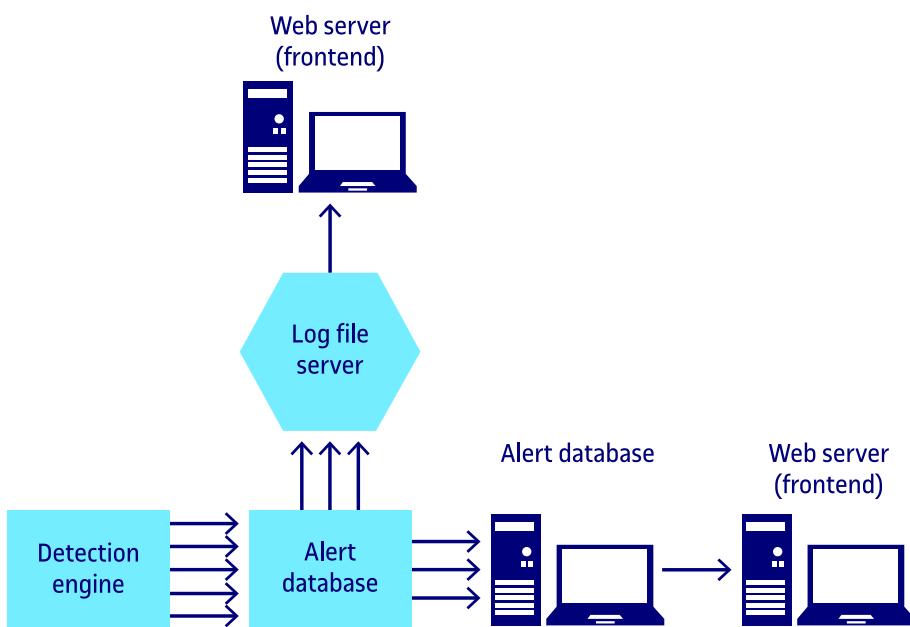
When the detection engine launches an alert, the result may involve the generation of a log file, sending the alert over the network through an SNMP message or even storing the information associated with the alert in a structured way by some database management system. Many of these tools are available and can be freely downloaded from the Internet.

As in the case of the detection engine and the preprocessor, Snort's notification system also uses a plugin system for the processing of information. Figure 14 shows a diagram for the operation of this notification system using plugins.

Figure 14. Snort notification system, through plugins

Snort alerts

If we try to run Snort in a congested network with a reasonable number of detection signatures activated, the amount of alerts launched by Snort can reach a hundred in a short time. This huge volume of information will not be easy to deal with using a simple browser of log files.



The purpose of installing Snort on a network is not only to obtain information (intrusion attempts), but to analyse this information and to be able to take the necessary actions based on the data obtained. If the number of active rules is high and the network traffic increases easily, it will not be very easy to analyse the information reported by Snort unless we use some other complementary tools, such as correlations and visualisation tools.

Furthermore, the interesting aspect of a detection system like Snort is not simply to log events, but to be able to react to intrusion attempts in a reasonably short period of time. Therefore, it will be necessary to use second applications that help consolidate and analyse the information reported by Snort in order to alert network administrators of the intrusion attempts analysed.

Currently, there are a large number of utilities to work with the information generated by Snort. Some of these tools are maintained by the same community of Snort developers, although we can also find applications developed by third parties or commercial applications.

4. Deception systems and techniques

As we have just seen in the previous sections, the detection of intruders assumes that an attacker is able to violate the security policies of a system. Thus, tools are used to inform system administrators, prepare reports and remediation plans, with the ultimate goal of reacting to attacks by an intruder in the appropriate manner. We have also seen that classic intrusion detection tools have evolved towards complete incident detection and management platforms, with the aim of facilitating the management of large volumes of information created by detection tools, not only tools for intrusion detection, but also events generated from other technologies, such as vulnerability scanners, or decoy and deception systems.

Specifically, these decoy and deception systems can be seen as complementary technologies. They can provide additional information to the management, information fusion and alert correlation platforms that we have seen before. Instead of addressing the problem of cyber defence from a preventive point of view, these systems try to change the rules of the game, offering the defender (e.g., the network administrator) the possibility of taking the initiative.

Decoy and deception systems, rather than neutralising the actions of a possible intruder, try to use similar techniques, but with the aim of learning from the intruder's offensive actions.

We will now summarise two strategies that can be used when building such deception systems and techniques.

First of all, we have a strategy known as the *installation of honeypot systems*. It is about installing network equipment that tries to attract traffic from an intruder, imitating as much as possible the real behaviour of the operating systems of the network in question. Thus, the network administrator will be able to see the attempts of a possible intruder before it enters the real computers. The aim can also be to see how the security elements implemented in real equipment behave. In this case, emulators, co-simulators and digital twins are usually used for creating the deception equipment.

Recommended reading

The following article (available online) provides more information about possible improvements in implementations of the honeypot concept from the use of digital twins: **Eckhart and Ekelhart** (2019). "Digital Twins for cyber-physical security: State of the art and Outlook", *Security and Quality in Cyber-Physical Systems Engineering*, pp. 383-412, Springer Nature. <http://dx.doi.org/10.1007/978-3-030-25312-7_14>

Decoy and deception equipment, also known as ***honeypots***, is network equipment that attempts to attract traffic from one or more attackers. In this way, their administrators can see attack attempts to enter the system and how the security elements implemented in the network behave.

Another aim is to obtain information about the tools and knowledge needed to make an intrusion in network environments such as those we intend to protect. All this information will end up serving to stop future attacks on the rest of the production equipment.

The conceptual idea of a *decoy and deception system* has existed for several decades. As a first approximation, we could define it as a network resource designed so that different attackers can introduce themselves in a simple way. These computers are usually designed to mimic the behaviour of production computers in order to be of interest to a community of attackers.

They usually have prevention mechanisms so that an intrusion to the decoy and deception computer does not end up causing an intrusion to the real computers of the network. Naturally, if an intruder manages to attack the computer, he does not have to realise that he is being monitored or deceived.

Thus, these computers should be installed behind configured firewall systems to allow incoming connections to the decoy computer, but limiting output connections.

In the second place, we have a strategy known as *building an entire network segment* which consists solely of decoy and deception equipment, all prepared to deceive the intruders (allowing access without too many difficulties). This is what is known as building a ***honeynet***.

The equipment in this segment will have to offer configured services to attract the attention of a whole community of intruders with the aim of recording all their movements through the decoy equipment.

The network can be real, with physical equipment and gateways, joined to the production network protected with traditional elements of monitoring, detection and prevention. The network can also be completely or partially virtual, combining the interaction between the physical equipment and gateways, with the use of emulation and virtualisation technologies. Here we find, for example, the use of virtual machines with virtualisation products of the Oracle/virtualbox, VMware/vSphere, Microsoft Hyper-V, etc. type, as well as

the use of Docker-type software containers or alternative solutions. Also, emulation and virtualisation of network elements, with the use of software-defined networking and other traditional elements of cloud computing.

Instead of routing at the network level (i.e., at level 3 of the OSI reference model), the deception system could also function as a network bridge at the link level (i.e., level 2 of the OSI reference model), so that it could do without using IP addresses and reduce the chances of detection by the attackers.

All systems installed within a deception network should be systems of decoy and deception and as realistic as possible. That is, they should be real systems and applications, like those we can find in any production equipment.

Since we have not found simulated services in these systems of decoy and deception, all conclusions drawn from an investigation can be extrapolated directly to a real production network. Thus, all the deficiencies and weaknesses that are discovered within a deception network will be the same as in most organisations today.

The operation of the decoy and deception network is based on a single principle: all traffic that enters any of its equipment must be considered suspicious.

The monitoring process will be carried out through the detection mechanisms installed on the gateway, detecting attacks based on already known trends or statistics. However, the possibilities of investigating all the activity of a network of decoy and deception should also help detect unknown attacks.

Decoy and deception systems should be seen as research tools to improve the safety of production networks. They are a valuable solution, but require high dedication (time and storage resources).

Summary

Computer networks are exposed to cyber attacks so often that a large number of security requirements must be imposed for the protection of their resources.

Although the deficiencies of these systems can be checked by using conventional tools, they are not always corrected. In general, these weaknesses can cause a hole in network security and facilitate unauthorized or intrusive activities into the system.

In the same way that to ensure the physical safety of a building, motion detectors, surveillance cameras, log books, etc. are usually installed, a computer network needs equivalent components in the digital world to collect and create intrusion scenarios, process alerts and prevent intrusive activities.

In this didactic module, we have presented intrusion detection systems (IDS), the aim of which is precisely to provide these complementary elements to traditional security mechanisms and to be able to offer additional capabilities to warn and guide network administrators when there are attacks and computer intrusions.

We have also seen additional functionalities to facilitate normalisation, consolidation and correspondence of events collected by heterogeneous detection sensors. Some of these features have been presented in detail and we have made a first approach to Snort, an intruder detection system, based on open source and offered to the community of network administrators as a free software tool. Snort's main goal is to help network administrators continuously monitor network traffic in search of intrusion attempts or misuse. Finally, we have completed the module with a quick presentation on deception systems and techniques, as a complementary technology to learn from the offensive actions of the intruder.

Glossary

attack Aggression to the security of a system resulting from an intentional and deliberate act that violates its security policy.

bug Programming defect that can trigger a security deficiency.

CERT See *Computer emergency response team*.

common vulnerabilities and exposures Public standard for the identification of vulnerabilities. It associates a unique identifier to each different vulnerability.
acronym CVE

common vulnerability scoring system Common framework for the evaluation of the criticality of vulnerabilities.
acronym CVSS

computer emergency response team Team of responses to computer emergencies; one of its main tasks is vulnerability management.
acronym CERT

computer security incident response team Computer security incident response team; one of its main tasks is vulnerability management.
acronym CSIRT

CSIRT See *Computer security incident response team*.

CVE See *Common vulnerabilities and exposures*.

CVSS See *Common vulnerability scoring system*.

DDoS See *distributed denial of service*.

denial of service Attack that attempts to saturate the victim's resources, such as the memory or computing and processing capacity.
acronym DoS

distributed denial of service Denial of service that takes place from various connection points.
acronym DDoS

DoS See *denial of service*.

exploit Program or script that allows one or more vulnerabilities to be exploited; that is, a program that allows an attack to be made by taking advantage of the vulnerability.

exploration of ports Technique used to identify the services offered by a particular system or piece of equipment.

malware Program with malicious purposes.

risk A loss expectation expressed as the probability that a specific threat will exploit a particular vulnerability with especially harmful results.

rootkit Program that allows privileged access to a computer and manages to hide its presence from the administrator. It usually uses several vulnerabilities to install itself and achieve its purpose.

security policy Set of rules and practices that define and regulate the security services of an organisation or system with the purpose of protecting its critical and sensitive resources. In other words, it is the declaration of what is allowed and what is not.

security vulnerability A failure or weakness in the design, implementation, operation or management of a system, which can be exploited to violate its security policy.

sniffer Application that intercepts all the information that passes through the network interface to which it is associated.

threat Potential breach of security that exists on the basis of circumstances, capabilities, actions or events that may cause a security breach or some damage to the system.

Trojan horse A program, apparently harmless, that contains an attack on an uncorrected vulnerability.
syn . trojan

vulnerability scanner An application that allows checking whether a system is vulnerable to a set of security deficiencies.

zero-day vulnerability Vulnerability that, at the time of being exploited, is not previously known.

Bibliography

Beale, J.; Foster, J. C.; Posluns J.; Caswell, B. (2003). *Snort 2.0 Intrusion Detection*. Oxford: Syngress Publishing.

Garcia-Alfaro, J. (2004). "Mecanismos para la detección de ataques e intrusiones". In: Herrera, J.; Garcia-Alfaro, J.; Perramon, X. *Security in computer networks*. Barcelona: Fundació Universitat Oberta de Catalunya.

Garcia-Alfaro, J. (2007). "Detección de ataques en red con Snort". In: Herrera, J.; Garcia-Alfaro, J.; Perramon, X. *Advanced Security Aspects in Networks*. Barcelona: Fundació Universitat Oberta de Catalunya.

Koziol, J. (2003). *Intrusion Detection with Snort*. Indianapolis: Sams Publishing.

Kruegel, C.; Valeur, F.; Vigna, G. (2004). *Intrusion Detection and Correlation: Challenges and Solutions*. Springer-Verlag.

Northcutt, S.; Novak, J. (2002). *The Network Intrusion Detection* (3rd ed.). New Riders.

Miller, D. R; Harris, S.; Harper, A. A.; Vandyke, S.; Blask C. (2011). *Security Information and Event Management (SIEM)*. Mc Graw Hill.

Proctor, P. E. (2001). *The practical intrusion detection handbook*. New Jersey: Prentice-Hall.

Rehman, R. (2003). *Intrusion Detection Systems with Snort. Advanced IDS Techniques Using Snort, Apache, MySQL, PHP, and ACID*. New Jersey: Prentice Hall PTR.

